

**A Novel Constraint-Based
Data Fusion System for
Limited-Angle Computed Tomography**

by
Jeffrey Edwin Boyd

Technical Report 94-7
March 1994

Department of Computer Science
University of British Columbia
Rm 201 - 2366 Main Mall
Vancouver, B.C.
CANADA V6T 1Z4

Telephone: (604) 822-3061
Fax: (604) 822-5485

A NOVEL CONSTRAINT-BASED DATA FUSION SYSTEM
FOR LIMITED-ANGLE COMPUTED TOMOGRAPHY

by

JEFFREY EDWIN BOYD

B.Sc., The University of Calgary, 1983
M.Sc., The University of Calgary, 1986

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

in

THE FACULTY OF GRADUATE STUDIES

Department of Computer Science

We accept this thesis as conforming to the required standard

Dr. James Little

Dr. Robert Woodham

Dr. Uri Ascher

Dr. David Lowe

Dr. Jean Meloche

THE UNIVERSITY OF BRITISH COLUMBIA

March 1994

© Jeffrey Edwin Boyd, 1994

Abstract

Computed tomography (CT) is a non-destructive evaluation technique that reconstructs the cross section of a specimen from x-ray raysum measurements. Whereas CT reconstruction is an ill-posed inverse problem that is easily solved, limited-angle CT, where raysum data are missing for a range of angles, is more severely ill-posed and more difficult to solve. In the limited-angle case, *a priori* assumptions are necessary to constrain the problem. Specimens wider than the x-ray source to sensor spacing require limited-angle CT. Furthermore, if the specimen is a sandwich structure, i.e., some core material surrounded by load-bearing face sheets, then the face sheets must lie in the null space. Components in the null space do not appear in the raysum data and thus confound CT reconstruction because there is no basis for interpolation. This thesis proposes a novel constraint-based data fusion method for limited-angle CT reconstruction of sandwich structures. The method reduces the reliance of limited-angle CT on assumptions by using range and ultrasound measurements to constrain the solution. Fusion of the data sources results in a problem with a much smaller null space that no longer includes the face sheets. The reduction of the null space in a manner consistent with the specimen yields a more accurate tomographic reconstruction. Synthetic and real data experiments show marked improvement in reconstruction accuracy achieved by using the fusion system.

Table of Contents

Abstract		ii
Table of Contents		iii
List of Tables		vi
List of Figures		vii
Acknowledgments		xii
Chapter 1	Introduction	1
1.1	Background	5
1.2	Limitations of Limited-Angle CT and Proposed Solution	7
1.3	Numerical Methods	9
1.4	Experimentation and Results	12
1.5	Outline	16
Chapter 2	Background	18
2.1.	Constraint-Based Data Fusion	19
2.2.	Computed Tomography	23
2.2.1	X-ray Absorption	24
2.2.2.	Basic Apparatus for CT Data Acquisition	25
2.2.3.	The Radon Transform and the Fourier Slice Theorem	26
2.2.4.	CT Reconstruction	27
2.3.	Limited Angle Computed Tomography	30
2.3.1.	Ill-Posed Nature of Limited Angle Computed Tomography	31
2.3.2.	Reconstruction from Limited-Angle Data	33
2.4.	Ultrasound	37
2.4.1.	Forward Models	38
2.4.2.	Ultrasound Inversion	40
2.4.3.	Minimal Ultrasound for Fusion with Computed Tomography	42
2.5	Chapter Summary	43

Chapter 3	Limited-Angle Computed Tomography for Sandwich Structures	
	Using Data Fusion	45
3.1.	Null Space of the Limited Angle Radon Transform	46
3.2.	Data Acquisition for Limited-Angle CT Data Fusion System	49
	3.2.1 Raysum Data Acquisition	50
	3.2.2 Range Data	51
	3.2.3 Thickness Data	52
3.3	The Fusion System	53
	3.3.1 Data Flow	54
	3.3.2. Mathematical Formulation	55
3.4	Chapter Summary	58
Chapter 4	Numerical Methods for Limited-Angle Tomography System	60
4.1	Singular Value Decomposition	61
	4.1.1 Properties of the Decomposition	61
	4.1.2 Least-Squares Minimum-Norm Solution	63
	4.1.3 Application to Limited-Angle CT	64
4.2	Regularization and Conjugate Gradient Method	66
	4.2.1 Regularization	66
	4.2.2 Parameter Selection	69
	4.2.3 The Conjugate Gradient Method	70
	4.2.4 Application to Limited-Angle CT	72
4.3	Projection onto Convex Sets	73
	4.3.1 General Description of POCS	73
	4.3.2 Constraint Sets for Limited-Angle CT	76
	4.3.3 Constraint and Parameter Selection	79
	4.3.4 Application to Limited-Angle CT	80
4.4	Chapter Summary	81

Chapter 5	Experimentation and Results	83
5.1	Description of Experiments	84
5.1.1	Implementation of Algorithms	85
	Raysum Operator	85
	Implementation of SVD	87
	Implementation of R/CG	88
	Implementation of POCS	90
5.1.2	Synthetic Data	91
5.1.3	Real Data	94
5.2	Results	97
5.2.1	SVD Trials	98
5.2.2	R/CG Trials	102
5.2.3	POCS Trials	109
5.3	Chapter Summary	117
Chapter 6	Conclusions and Discussion	119
6.1	Conclusions	119
6.1.1	Novel Contributions	119
6.1.2	Generalization of Results	121
6.2	Discussion	122
6.2.1	Practical Application	122
6.2.2	CAD Model as a Source of Constraints	124
6.2.3	More Sophisticated Constraints	125
6.2.4	Compton Scatter and Ultrasound Fusion	130
Bibliography	133

List of Tables

Table 4.1	Summary of numerical methods considered for limited-angle computed tomography with data fusion.	82
Table 5.1	SVD trial results.	99
Table 5.2	R/CG synthetic data trial results.	102
Table 5.3	R/CG real data trial results.	103
Table 5.4	POCS synthetic data trial results.	108
Table 5.5	POCS trial 6 (real data) results.	110
Table 6.1	Summary of successive minimizations used to produce reconstruction in Figure 6.1.	123
Table 6.2	Summary of successive POCS applications used to produce reconstruction in Figure 6.2.	124

List of Figures

Figure 1.1	Synthetic cross section image for SVD trials.	13
Figure 1.2	Synthetic cross section image for R/CG and POCS trials.	13
Figure 1.3	Cross section of plexiglass phantom simulating a sandwich structure with graphite/epoxy composite face sheets and aluminum honeycomb core.	14
Figure 1.4	Reconstructions from synthetic data with range and ultrasound constraints for SVD, R/CG, and POCS.	15
Figure 1.5	Reconstructions of the plexiglass phantom using R/CG and POCS.	16
Figure 2.1	Model of computational vision.	19
Figure 2.2	Clark and Yuille taxonomy of data fusion systems.	20
Figure 2.3	Classes of weakly-coupled data fusion systems.	22
Figure 2.4	Strongly-coupled data fusion systems.	23
Figure 2.5	Schematic of x-ray absorption.	25
Figure 2.6	Schematic of basic computed tomography apparatus.	26
Figure 2.7	Example of CT showing lines of projection on a two-by-two grid.	29
Figure 2.8	Example cases for which only limited-angle raysums are available: (a) object is very long in one direction, and (b) other objects obstruct scanning.	31
Figure 2.9	Effect of limited-angle Radon transform on Fourier domain sampling: (a) range of angles sampled (shaded region) and (b) effective sampling of Fourier domain (shaded region is sampled).	32

Figure 2.10	Triangular region of spatial support from Sato <i>et al.</i> [31] method for limited-angle CT reconstruction of the letter 'A'.	36
Figure 2.11	Schematic of a lossless layered ultrasonic medium: (a) top layer and (b) internal layers.	39
Figure 3.1	Function $s_{r,\phi,a,d}(x,y)$; shaded region has value a and zero everywhere else.	48
Figure 3.2	Sandwich specimen with face sheets at angle ϕ . The width of the specimen prevents raysum acquisition outside the range $-\Theta \leq \theta \leq \Theta$. Note that $\Theta < (\pi/2) - \phi$, so the face sheets lie almost entirely within the limited-angle Radon transform null space.	49
Figure 3.3	Raysum data acquisition system.	50
Figure 3.4	Scan method for limited-angle raysum acquisition: (a) parallel with vertical and (b) at an angle to vertical.	51
Figure 3.5	Scanning method modified to include laser range finder for bounding region measurement.	52
Figure 3.6	Apparatus for ultrasound data acquisition.	53
Figure 3.7	Block diagram of strongly-coupled feed forward data fusion system for limited-angle CT of sandwich structures.	54
Figure 4.1	Example of POCS convergence.	75
Figure 4.2	Demonstration of the variability of POCS convergence. POCS converges to three distinct but correct solutions depending on the order of projection operators.	75

Figure 5.1	Steps to a compact raysum operator: (a) hypothetical grid for reconstruction, (b) parallel rays perpendicular to upper surface of reconstruction, (c) parallel rays at an angle to reconstruction, and (d) set of rays, one at each scan angle.	85
Figure 5.2	Convolution kernel for Laplacian operator.	89
Figure 5.3	Convolution kernel for Laplacian operator with different coefficients for x and y directions.	89
Figure 5.4	Synthetic cross section image for SVD trials.	91
Figure 5.5	Synthetic cross section image for R/CG and POCS trials.	92
Figure 5.6	Synthetic fusion data vectors: for spatial support data only (a) known region and (b) known linear attenuation, and for spatial support and face sheet data (c) known region and (d) known linear attenuation.	93
Figure 5.7	Cross section of plexiglass phantom simulating a sandwich structure with graphite/epoxy composite face sheets and aluminum honeycomb core.	94
Figure 5.8	Apparatus for limited-angle CT raysum data acquisition system.	95
Figure 5.9	Reconstructed images for SVD trials: (a) trial 1 reconstruction (max = 0.48, min = -0.01), (b) trial 1 error vector (max = 0.30, min = 0.0), (c) trial 2 reconstruction (max = 0.55, min = 0.0), (d) trial 2 error vector (max = 0.23, min = 0.0), (e) trial 3 reconstruction (max = 0.42, min = -0.04), (f) trial 3 error vector (max = 0.04, min = 0.0), and (g) rescaled trial 3 error vector (max = 0.04, min = 0.0).	100

Figure 5.10	(a) Weakly-coupled reconstruction using SVD (max = 0.48, min = -0.01), and (b) the corresponding absolute error vector (max = 0.15, min = 0.0).	102
Figure 5.11	Reconstructed images for R/CG synthetic data trials: (a) trial 1 reconstruction (max = 0.49, min = -0.06), (b) trial 1 error vector (max = 0.40, min = 0.0), (c) trial 2 reconstruction (max = 0.55, min = -0.05), (d) trial 2 error vector (max = 0.43, min = 0.0), (e) trial 3 reconstruction (max = 0.45, min = -0.07), and (f) trial 3 error vector.	103
Figure 5.12	(a) Weakly-coupled reconstruction using R/CG (max = 0.49, min = -0.05), and (b) the corresponding absolute error vector (max = 0.22, min = 0.0).	105
Figure 5.13	Reconstructed images for R/CG real data trials: (a) $\alpha^2 = 0.001$, (b) $\alpha^2 = 0.01$, (c) $\alpha^2 = 0.1$, (d) $\alpha^2 = 1.0$, (e) $\alpha^2 = 10.0$, (f) $\alpha^2 = 100.0$, and (g) $\alpha_1^2 = 0.001$ and $\alpha_2^2 = 0.1$. See Table 5.3 for maxima and minima.	107
Figure 5.14	Reconstructed images for POCS synthetic data trials: (a) trial 1 reconstruction, (b) trial 1 error vector (max = 0.4, min = 0.0), (c) trial 2 reconstruction, (d) trial 2 error vector (max = 0.4, min = 0.0), (e) trial 3 reconstruction, (f) trial 3 error vector (max = 0.22, min = 0.0), (g) trial 4 reconstruction, and (h) trial 4 error vector (max = 0.25, min = 0.0).	110
Figure 5.15	(a) Weakly-coupled reconstruction using POCS, and (b) the corresponding absolute error vector (max = 0.24, min = 0.0).	113
Figure 5.16	Reconstructed cross section for POCS real data trial 5.	114

Figure 5.17	Reconstructed images from POCS real data trial 6: (a) $\epsilon_R = 0.1$, and (b) $\epsilon_R = 0.001$	114
Figure 5.18	Plots of convergence for POCS: (a) trial 6(a) - $\epsilon_R = 0.1$ and (b) trial 6(b) - $\epsilon_R = 0.001$	116
Figure 6.1	Reconstruction from successive R/CG iterations incorporating the assumption that pixels are either air ($\mu = 0cm^{-1}$) or plexiglass ($\mu = 0.45cm^{-1}$).	127
Figure 6.2	Reconstruction from successive POCS iterations incorporating the assumption that pixels are either air ($\mu = 0cm^{-1}$) or plexiglass ($\mu = 0.45cm^{-1}$).	129

Acknowledgments

Many people have contributed to the success of my research. I owe Jim Little, my supervisor, a special debt of gratitude. The time he committed to me and his genuine interest in my work helped me greatly. I am specially grateful for the opportunity he gave me to prove myself in computer science. The other members of my committee, Bob Woodham, Uri Ascher, David Lowe and Jean Meloche, have also earned my gratitude both as instructors in graduate courses and in guiding my thesis research.

UBC is a special place. It is not special because of the buildings or the scenery, but because of the people who live and work there. I take this opportunity to thank my fellow students for their friendship, conversation, and companionship that made my stay at UBC worthwhile.

I thank the Department of National Defence for its financial support. At Defence Research Establishment Pacific, several people lent me a hand. The assistance of Bill Ramsbottom, of Galatea Research Inc., with data acquisition was invaluable. Bill Sturrock contributed his wisdom, guidance and encouragement. Ken McRae made sure that I had the best available advice with respect to ultrasound technology.

During the course of my research I had the occasion to meet Heinz Bauschke and his colleagues in the Centre for Experimental and Constructive Mathematics at Simon Fraser University. Heinz's help in understanding convex projections was not only extremely helpful to my project, but also inspired me to learn more about applied mathematics.

It is a long road to complete a PhD program, so I always strove to enjoy what I was doing. Fortunately for me there is my wife, Brenda. She not only supports my academic endeavours, but is ever vigilant in assuring that my life is full and happy, both at work and at home.

Chapter 1

Introduction

Non-destructive evaluation (NDE) is the field of endeavour that evaluates the structural integrity of an object without destroying the object. Measurements of a plethora of physical phenomena, including x-ray radiation, neutron radiation, ultrasound, eddy currents, infrared radiation, and fluorescence, indicate the structural integrity of NDE specimens. Although often obscured by detail, the ultimate goal of an NDE technique is to use such measurements to predict the failure of a specimen and prevent expensive or catastrophic failure.

Radiographs, images produced by projecting x-rays through a specimen onto a sensor, are commonplace in NDE. Each pixel in a radiograph is a *raysum* measurement, i.e., the integral of linear attenuation coefficients of the specimen along the ray from the x-ray source to the sensor. A single radiograph presents the three-dimensional structure of a specimen projected onto a two-dimensional image. Proper inspection requires images from multiple viewing angles. For example, in real-time radiography, viewing images over continually changing angles hints at the underlying three-dimensional structure.

Computed tomography (CT) is an NDE technique that reconstructs a cross section of a specimen from its raysums. Adjacent cross sections combine to form a volumetric image, so CT solves the problem of determining three-dimensional structure from raysums. The proliferation of CT scanners in medicine, in spite of cost and patient exposure to radiation, attests to the diagnostic value of reconstructed cross sections.

The Radon transform of a cross section is the set of all possible raysums through the cross section. CT reconstruction from the continuous Radon transform is an ill-posed inverse problem; the solution does not vary continuously with the data. The practical problem of reconstruction from a discrete sampling of the Radon transform is also ill-posed, but the ill-posedness is overcome by assuming that the reconstruction is band-limited.

Often it is not possible to measure raysums over the full range of angles. In such cases, reconstruction is more ill-posed* because of the missing data. Techniques that deal with this type of limited data are called collectively *limited-angle* CT and have received considerable attention. To compensate for missing data, limited-angle CT methods rely on constraints to compute reconstructions.

Specimens wider than the x-ray source-to-sensor spacing prevent raysum measurement at some angles and require limited-angle CT for reconstruction. Normally the specimen is part of some mechanical structure and must be rigid in the presence of bending moments. For economy of material and weight the structures use a thin load-bearing material at the outer surfaces with some core between, i.e., a sandwich structure. Common examples are *I*-beams, windsurfers, skis, and aircraft control surfaces. This thesis focuses on aircraft control surfaces that use two face sheets (typically aluminum, but some modern aircraft use graphite/epoxy composite) to carry loads on either side of

* We say that a problem is more ill-posed than another if it has a larger null space, i.e., the null space has a greater number of dimensions.

an aluminum honeycomb core. The core provides a rigid separation between the face sheets.

Not only do wide specimens restrict raysum acquisition, but, as this thesis shows, limited-angle CT cannot accurately reconstruct the cross section in the presence of face sheets. The face sheet structure lies almost entirely within the null space of the limited-angle Radon transform. In short:

1. wide specimens require limited-angle CT,
2. wide specimens are often sandwich structures with face sheets, and
3. limited-angle CT does not work in the presence of face sheets.

Typical constraints that make limited-angle CT possible cannot account for the face sheets but instead lead to erroneous reconstructions based on general assumptions. However, the reliance of limited-angle CT on constraints makes it an excellent candidate for constraint-based data fusion. Constraint-based data fusion uses both *a priori* assumptions and data from other sources to constrain the solution of an inverse problem.

Ultrasound offers a source of additional data that complements the raysums. Whereas raysums are oblivious to discontinuities along the rays (the very property that prevents them from dealing with the face sheets), ultrasound responds precisely to such discontinuities and is capable of measuring the face sheet structure. This thesis proposes a novel method for limited-angle CT using data fusion. The method capitalizes on the complementary nature of x-rays and ultrasound and is specifically applicable to sandwich structures. In the context of this thesis we consider the method as a means of producing accurate reconstructions only. For the sake of development, we ignore the ultimate application, inspection to predict failure.

The method fuses laser range data and ultrasound data with raysums to compute an accurate reconstruction. Laser range data give the spatial support of the specimen. Ultrasound data, combined with range data, segment the reconstruction into exterior, face sheet, and interior regions. These data then restrict the null space of the forward operator

so that it does not include the face sheet structures, leading to accurate reconstructions. Data fusion alone cannot guarantee better solutions to inverse problems, but here the complementary nature of ultrasound and x-ray yields superior reconstructions. This thesis shows that the proposed limited-angle CT system computes accurate reconstructions of sandwich structures where an accurate reconstruction is not possible otherwise.

The following system of linear equations combines the raysum data with the fusion data:

$$\begin{bmatrix} R \\ W \end{bmatrix} \mathbf{x} = \begin{bmatrix} \mathbf{y} \\ W\mathbf{x}_f \end{bmatrix},$$

where R is the forward raysum operator, \mathbf{x} is the cross section image vector, W is a diagonal matrix, \mathbf{y} is the raysum data and \mathbf{x}_f is a vector of linear attenuation coefficients based on the fusion data. An image vector of weights indicating whether or not the linear attenuation coefficient for a pixel is known from the fusion data, forms the diagonal of W . This thesis considers three methods to solve for \mathbf{x} from the above equation:

1. singular value decomposition (SVD),
2. regularization and conjugate gradient method (R/CG), and
3. projection onto convex sets (POCS).

Experimental trials with synthetic and real data apply all three methods to compute reconstructions. Synthetic data trials show marked improvement in the reconstructions obtained using the fusion data constraints. In addition, the SVD trials show that the number of dimensions in the null space (as indicated by the number of singular values equal to zero) decreases by adding the fusion constraints. Real data trials show that the proposed fusion method works with real data, and, in general, experimental results support the contention that the proposed technique improves reconstruction.

SVD gives accurate reconstructions from synthetic data, but, because of computational limitations, it is only applicable to small problems (smaller than practical real data problems). R/CG gives excellent reconstructions from error free synthetic data

and is fast (allowing application to real data problems). However, in the presence of errors in the real data, regularization forces a tradeoff between image smoothness and fidelity to data, and no degree of compromise is totally acceptable. POCS yields the best results of the three numerical methods considered. It produces accurate reconstructions from error-free synthetic data, does not force a smoothness compromise in the presence of errors in real data, and converges quickly. Each of the numerical methods uses different constraints to arrive at a solution, which means the methods solve different problems. The results presented here serve only to demonstrate alternative methods for reconstruction and do not comprise a proper comparison of the methods.

This thesis concludes that the complementary nature of x-ray and ultrasound allows exploitation of constraint-based data fusion to improve limited-angle CT. This complementary nature means the constraints from the ultrasound data are not redundant but instead reduce the null space of the problem. In particular, limited-angle CT cannot deal with face sheets in sandwich structures. However, the ultrasound can measure the face sheets of the structure thus allowing accurate reconstruction when combined with the raysum data.

1.1 Background

The proposed limited-angle CT method uses constraint-based data fusion as described by Clark and Yuille [9]. They describe two types of data fusion systems for solving inverse problems. The systems are composed of modules where each module takes some sensory data as input and produces as output a solution or set of solutions to an inverse problem. The modules may or may not use *a priori* constraints.

In the first type of fusion system, the modules operate independently to produce solutions. A fusion module then combines these solutions to form a single solution. The independent operation of the modules characterizes these *weakly-coupled* systems. In the second type of fusion system, at least one module provides constraints to another. Interaction of modules distinguishes these *strongly-coupled* systems. Implementation of

the interaction via constraints to modules gives the alternative appellation *constraint-based*.

It is common practice to use *a priori* assumptions to constrain a module solving an inverse problem. The constraints serve to stabilize the solution to an ill-posed or ill-conditioned inverse problem (e.g., Poggio, Torre and Koch [30] and Bertero, Poggio and Torre [5]). However, the assumptions may be insufficient or invalid, thereby leading to erroneous solutions. In such cases, constraint-based data fusion is useful because it shifts reliance from poor assumptions to measurements from other sensory sources. Although in general the use of other sensory data to provide constraints does not necessarily improve the solutions of inverse problems, if the data sources complement each other, i.e., one provides information where the other cannot, then improved results are possible.

Computed tomography (CT) is an inverse problem in the field of non-destructive evaluation (NDE). CT computes a reconstruction of the cross section of a specimen from its raysums (line integrals measured by an x-ray system). The cross section is an image in which the pixel values are the linear attenuation coefficients of the specimen. The complete set of raysums for a cross section is its Radon transform. The Fourier slice theorem relates the Radon transform to the Fourier transform. It states that measuring raysums is equivalent to sampling in the Fourier domain. Inversion of the continuous Radon transform is an ill-posed problem; the solution does not vary continuously with the data. In practice, only measurement of the discrete Radon transform is possible. Reconstruction from the discrete Radon transform is also ill-posed since the null space is non-trivial and consists of high-frequency components. Requiring that the reconstruction be band-limited easily handles the ill-posedness.

Circumstances arise that prevent raysum measurement throughout a full range of angles leading to a special variation of the CT problem called limited-angle CT. Limited-angle CT is ill-posed because the missing raysum data leave the Fourier space under-sampled. Ill-posedness in the limited-angle case is more severe than for conventional CT

(i.e., the null space is much larger for limited-angle CT) and it is not easily handled. A variety of techniques proposed in the literature perform limited-angle CT. Without exception, they all rely on *a priori* assumptions to constrain the reconstruction and find a unique solution. In cases where *a priori* assumptions may be erroneous or insufficient, limited-angle CT can benefit from constraint-based data fusion by incorporating data from a complementary source.

Ultrasound is a source of data that complements x-rays very well. X-rays measure line integrals and are insensitive to discontinuities along the line of integration. The opposite is true for ultrasound which detects discontinuities in a layered system. The limited-angle CT system with data fusion proposed in this thesis exploits this property of ultrasound to measure the thickness of face sheets in sandwich structures. X-rays cannot measure this thickness, so ultrasound, by virtue of its complementary nature to x-rays, adds non-redundant data to the problem.

In summary, CT is an ill-posed inverse problem that relies on band-limiting constraints for solution. Limited-angle CT is further ill-posed because of incomplete sampling of the Fourier domain and requires additional constraints for solution. Situations may arise where *a priori* assumptions used as constraints for limited-angle CT are either erroneous or insufficient. Reliance on constraints for solution suggests the use of constraint-based data fusion. Data fusion requires another source of data to complement the raysum data. Ultrasound provides such a source because it can measure discontinuities that x-rays cannot.

1.2 Limitations of Limited-Angle CT and Proposed Solution

The null space of the limited-angle Radon transform is the set of all functions with Fourier transforms equal to zero within the region of the Fourier domain sampled by the limited-angle Radon transform. Physical structures that lie within this null space cannot be reliably reconstructed from the limited-angle raysums. Such structures contain long edges at angles beyond those sampled by the raysums.

Sandwich structures are commonplace and lie within the null space of the limited-angle Radon transform. They economize on weight and material by using some thin material at the outer surfaces of the structure with some core material sandwiched between. The outer surfaces carry the loads due to bending while the core serves primarily to separate the two surfaces. Commonplace examples are I-beams, windsurfers, skis, and aircraft control surfaces. In aircraft control surfaces, two face sheets (typically aluminum but some modern aircraft use graphite/epoxy composite) surround an aluminum honeycomb core. Face sheets carry most of the loads in the specimen while the honeycomb provides a rigid separation of the face sheets. An x-ray source-to-sensor spacing narrower than the face sheet width necessitates limited-angle CT. At the same time, the face sheet that necessitates limited-angle methods lies in the null space of the limited-angle Radon transform. Thus, the presence of face sheets in a sandwich structure confounds accurate CT reconstruction.

This thesis proposes a constraint-based data fusion system that deals with the inability of other limited-angle CT systems to reconstruct sandwich structures accurately. The system uses constraints derived from laser range data and ultrasound data to constrain the tomographic reconstruction. Here we consider the system as a means of producing accurate reconstructions only. Further evaluation as a technique for inspection and prediction of failure is ignored.

There are three modules in the system. The first module takes laser range data as input and computes a bounding box for the specimen, i.e., the spatial support of the specimen. The output of the module is a segmentation of the reconstruction into exterior and interior regions. Since the exterior regions must be air (which has zero linear attenuation) the segmentation provides a constraint for reconstruction.

The second module takes ultrasound thickness measurements as input, and, constrained by the spatial support data, further segments the reconstruction into exterior,

interior and face sheet regions. With the known linear attenuation of the face sheet material, the segmentation provides another constraint for reconstruction.

The last module takes limited-angle raysum data as input and, subject to spatial support and face sheet constraints, reconstructs a cross-section of the image as output. This module computes its output by solving the linear system:

$$\begin{bmatrix} R \\ W \end{bmatrix} \mathbf{x} = \begin{bmatrix} \mathbf{y} \\ W\mathbf{x}_F \end{bmatrix},$$

where R is the discrete limited-angle raysum operator, \mathbf{x} is the cross section image, and \mathbf{y} is the raysum data. The segmentation of the reconstruction provides a vector of weights that forms the diagonal of the diagonal matrix W . A weight of one on the diagonal indicates that the corresponding pixel has a known value while a zero indicates unknown. \mathbf{x}_F is a vector of linear attenuation coefficients derived from the constraints. Entries in \mathbf{x}_F that correspond to a weight of one in W have known linear attenuation and other entries are arbitrary.

Let $A = \begin{bmatrix} R \\ W \end{bmatrix}$ and $\mathbf{b} = \begin{bmatrix} \mathbf{y} \\ W\mathbf{x}_F \end{bmatrix}$ to give a generalized version of the above equation:

$$A\mathbf{x} = \mathbf{b}.$$

A is not square and is not full column rank. Solution for \mathbf{x} from the above is still an ill-posed problem; if a solution exists then there is an infinity of solutions. Some *a priori* assumptions are necessary to arrive at a unique solution. However, these assumptions do not lead to an erroneous reconstruction because A properly accounts for the face sheets due to the fusion constraints.

1.3 Numerical Methods

This thesis considers three numerical methods for solving the linear system of equations for the proposed limited-angle CT method. These are:

1. singular value decomposition (SVD),
2. regularization and conjugate gradient method (R/CG), and

3. projection onto convex sets (POCS).

This section introduces the methods which are explained in more detail in Chapter 4.

Table 4.1 summarizes the comparison of all three methods.

The first method, SVD, computes the decomposition:

$$A = U\Sigma V^T$$

where U , and V are orthogonal and Σ is diagonal. Given the decomposition, it is possible to condition the matrix A and solve for the minimum-norm least squares solution, $\hat{\mathbf{x}}$, using the pseudo-inverse of Σ , Σ^+ , i.e.:

$$\hat{\mathbf{x}} = V\Sigma^+U^T\mathbf{b}.$$

SVD has the beneficial property of determining the null space and range of A , and its singular values (the square roots of the eigenvalues of $A^T A$). The number of non-zero singular values gives the rank (row and column) of $A^T A$. Unfortunately, SVD has $O(mn^2 + n^3)$ complexity in time (where m is the number of rows in A , and n is the number of columns in A and the number of pixels in the reconstruction) and is too slow to apply to large problems.

The second method, R/CG, combines regularization and the conjugate gradient method. The conjugate gradient method (CG) solves the linear system $A\mathbf{x} = \mathbf{b}$ by minimizing the objective function:

$$E = \frac{1}{2} \mathbf{x}^T A \mathbf{x} - \mathbf{x}^T \mathbf{b} + c,$$

where c is a constant and A is positive definite. For a sparse matrix A , iterations are fast and, depending on the condition number of A , CG converges quickly to a solution. A for the data fusion system is neither positive definite nor square. Regularization solves this problem by creating the new objective function:

$$E = \left\| \begin{bmatrix} A \\ \alpha\Omega \end{bmatrix} \mathbf{x} - \begin{bmatrix} \mathbf{b} \\ \alpha\Omega\mathbf{x}_0 \end{bmatrix} \right\|^2,$$

where Ω makes $\begin{bmatrix} A \\ \alpha\Omega \end{bmatrix}$ full column rank, α is a constant and \mathbf{x}_0 is some vector to be specified. CG finds \mathbf{x} that minimizes E by solving the new positive definite system of equations:

$$(A^T A + \alpha^2 \Omega^T \Omega) \mathbf{x} = A^T \mathbf{b} + \alpha^2 \Omega^T \Omega \mathbf{x}_0.$$

α^2 selects a compromise between the original system of equations and the solution to $\Omega \mathbf{x} = \Omega \mathbf{x}_0$. With α^2 set too small the system is ill-conditioned and CG fails. A larger α^2 gives a system that is well-conditioned but biases the solution towards $\Omega \mathbf{x} = \Omega \mathbf{x}_0$.

There is a variety of options in selection of Ω and \mathbf{x}_0 . In experimental trials conducted for this thesis Ω is:

$$\Omega = \begin{bmatrix} D_1 \\ D_2 \end{bmatrix},$$

where D_1 and D_2 are discrete first derivative operators in the x and y directions respectively, and:

$$\mathbf{x}_0 = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

Thus, α^2 sets the degree of compromise between a solution that conforms to the original ill-conditioned and ill-posed problem $A \mathbf{x} = \mathbf{b}$, and a well-conditioned problem with the solution biased to be flat.

The third numerical method, POCS, uses a set of constraints where each constraint requires the solution to lie within a convex set. POCS also requires that the intersection of the constraint sets be non-empty. Starting from some arbitrary initial solution, POCS iterates towards an ultimate solution by sequentially projecting the solution onto each of the constraint sets. A non-empty intersection of the sets guarantees convergence. POCS subsumes the better known technique called ART (algebraic reconstruction technique). Whereas POCS applies to any convex constraint sets, ART is POCS using only convex sets that are hyperplanes. For the limited-angle CT problem at hand the convex sets are hyperslabs and spheres. In practice, with $\mathbf{0}$ as a starting point,

POCS finds a solution near the minimum norm solution. This is not guaranteed in general for POCS. POCS converges quickly for the problem of interest. So POCS is a fast method (comparable to CG for this problem), finds a solution like the minimum norm solution of SVD, and avoids the smoothing of R/CG. Also, the flexibility of POCS allows mixing of l_2 and l_∞ fits to the data, which is valuable in constraint-based data fusion.

Implementation of POCS for the proposed limited-angle CT data fusion system uses three constraint types:

1. conformity to raysum data,
2. conformity to fusion data, and
3. conformity to amplitude limits.

Previous work implements type 1 constraints, conformity to raysum data, as linear equalities, i.e., the convex set is a hyperplane within which the solution must lie. Such a constraint restricts the solution severely by requiring it to match the data precisely. In this thesis, the type 1 constraints are hyperslabs surrounding a hyperplane, i.e., the intersection of two half spaces centred about the hyperplane. Solutions in the hyper plane match the data exactly while other solutions in the slab only approximately match the data. A constant parameter determines the thickness of the hyperslabs. Type 2 constraints, conformity to fusion data, are a variation on previously published work using full reference images for reconstruction. In this thesis, the constraint does not require a full reference image, but instead uses a partial reference image from the fusion data. Type 3 constraints, conformity to amplitude limits, are also hyperslabs requiring that values in the solution lie between fixed limits.

1.4 Experimentation and Results

A series of experimental trials verifies the validity of the proposed limited-angle CT data fusion system. Each trial consists of application of one of the three numerical

algorithms to some data, to produce a reconstruction for analysis. The algorithm, its parameters, and the data uniquely define each trial.

Experiments use both synthetic and real data. Application of the forward raysum operator to a synthetic cross section image generates synthetic raysum data. The synthetic cross sections imitate the plexiglass phantom used in the real data trials, which itself imitates a graphite/epoxy sandwich structure with honeycomb core. The plexiglass phantom has thicker parts than aluminum honeycomb to avoid pushing resolution limits of the apparatus. Because SVD can handle only small problems, SVD trials are limited to synthetic data based on the low-resolution cross section of Figure 1.1. Synthetic data for R/CG and POCS trials are based on the high-resolution image of Figure 1.2.




Figure 1.1: Synthetic cross section image for SVD trials.




Figure 1.2: Synthetic cross section image for R/CG and POCS trials.

The plexiglass phantom used as the specimen for real data trials, shown in Figure 1.3, consists of two plexiglass face sheets (3mm thick) surrounding a core of plexiglass members (3mm thick) perpendicular to the face sheets. Some plexiglass inserts provide an anomaly in the structure similar to entrapped water. Chapter 5 describes the data acquisition system for the real data trials. It is a modification of an x-ray inspection

system intended to produce two-dimensional radiographs and is not ideally suited to CT. Nevertheless, the data it produces are adequate to show the validity of the method.





Figure 1.3: Cross section of plexiglass phantom simulating a sandwich structure with graphite/epoxy composite face sheets and aluminum honeycomb core.


For each numerical method, the first three trials use synthetic data with varying levels of fusion constraints. All methods exhibit a marked improvement in the accuracy of reconstruction when using spatial support and face sheet constraints. For SVD the error measure improves from 50.7% to 4.8% by using all constraints. Similarly, with R/CG the error measure drops from 64.0% to 6.7%, and with POCS the error measure drops from 62.6% to 6.0%. Furthermore, SVD trials show the change in the number of singularities in the problem. There is a reduction from 148 singularities of 300 singular values for raysum data to only 13 singularities for raysum data with spatial support and face sheet constraints. The character of the problem changes significantly with the constraints in such a way as to improve the accuracy of the reconstruction. Figure 1.4 shows sample reconstructions for all three numerical methods using all the constraints.

Real data trials with R/CG and POCS show the validity of the proposed system with real data. Although the quality of reconstruction is not on par with the synthetic


trials, this is largely due to errors in raysum measurements. Superior apparatus should lead to reconstructions of sufficient quality for practical application. Figure 1.5 shows reconstructions of the plexiglass phantom for R/CG and POCS.



(a)



(b)



(c)

Figure 1.4: Reconstructions from synthetic data with range and ultrasound constraints for (a) SVD, (b) R/CG, and (c) POCS.

Errors in the real data emphasize the effects of parameters used in reconstruction. A series of real data trials for both R/CG and POCS highlights the effects of the parameters. R/CG trials show the compromise selected by the parameter α^2 and POCS trials show the effect of changing constraint set size on convergence and reconstruction quality.




Figure 1.5: Reconstructions of the plexiglass phantom using (a) R/CG and (b) POCS.

1.5 Outline

Chapter 2 gives background material for the thesis starting with a review of the Clark and Yuille model for constraint-based data fusion. Following that, the chapter gives an introduction to computed tomography, pointing out that CT is an ill-posed problem that is easily managed. From CT follows the idea of limited-angle CT where reconstruction is done from incomplete data. Limited-angle CT is ill-posed to a greater degree and must rely on *a priori* assumptions for constraints, thus making it a good candidate for improvement by constraint-based data fusion. Chapter 2 describes ultrasound, which complements x-rays well, and, therefore, provides an excellent source of additional data for fusion.

Chapter 3 defines the null space of the limited-angle Radon transform and shows that face sheets in a sandwich structure lie mostly within this null space. The result is that limited-angle CT cannot reconstruct sandwich structures properly. A novel method for limited-angle CT using constraint-based data fusion proposed in the chapter addresses

this problem and is the focus of this thesis. Chapter 3 describes the method and its data sources, ending with the mathematical formulation of the method as a system of linear equations.

There are many ways to solve a system of linear equations. Chapter 4 describes the three numerical methods used in experimentation for this thesis, including application of each method to the proposed limited-angle CT method. A comparison at the end of Chapter 4 highlights the strengths and weaknesses of each numerical method.

Chapter 5 presents the results of experimentation that verify the validity of the proposed method. Synthetic data trials show a reduction in null space and a great improvement in accuracy of reconstruction with the incorporation of data fusion constraints. Real data trials show that the method works in practice and not just with error-free synthetic data. Experiments also show the effects of parameters on reconstruction quality.

Chapter 6 summarizes the conclusions of the thesis. It also discusses possibilities for further enhancements and variations of the proposed method, and its potential for application.

Chapter 2

Background

As a prelude to the proposal of a novel method for limited-angle CT in Chapter 3, this chapter presents relevant background material in four areas: constraint-based data fusion, CT, limited-angle CT, and ultrasound. Section 2.1 describes the constraint-based approach to data fusion. Constraint-based data fusion, so called because it uses the output from some inversion module as constraints for another, reduces or eliminates reliance on *a priori* assumptions for constraints. The value of the fusion emerges when the fusion constraints replace invalid assumptions.

CT is the ill-posed problem of reconstructing the cross section of an object from its x-ray raysums. Although the problem is ill-posed, the null space is small, consisting of high-frequency oscillations. Assuming *a priori* that the reconstructed cross section is band-limited is a sufficient constraint for reconstruction. Section 2.2 reviews conventional CT.

The shape of an object, or its environment, may prevent acquisition of raysum data through the full range of angles required for CT. The ensuing problem of limited-angle CT (reviewed in Section 2.3) is ill-posed too, but has a much larger null space. A

priori assumptions provide constraints beyond those of conventional CT. The reliance of limited-angle CT on assumptions that may not always hold suggests that it can benefit from constraint-based data fusion.

Ultrasound (reviewed in Section 2.4) is a versatile NDE technique providing data that are complementary to x-ray raysums. Whereas x-rays are insensitive to discontinuities along the path of radiation, ultrasound explicitly detects discontinuities in its path. It is this complementary nature that makes ultrasound potentially useful as a source of data for fusion with limited-angle CT.

2.1. Constraint-Based Data Fusion

As data fusion is a central part of this thesis, it is necessary to consider how to implement the data fusion. This thesis adopts the elegant *constraint-based* approach of Clark and Yuille [9], summarized in this section.

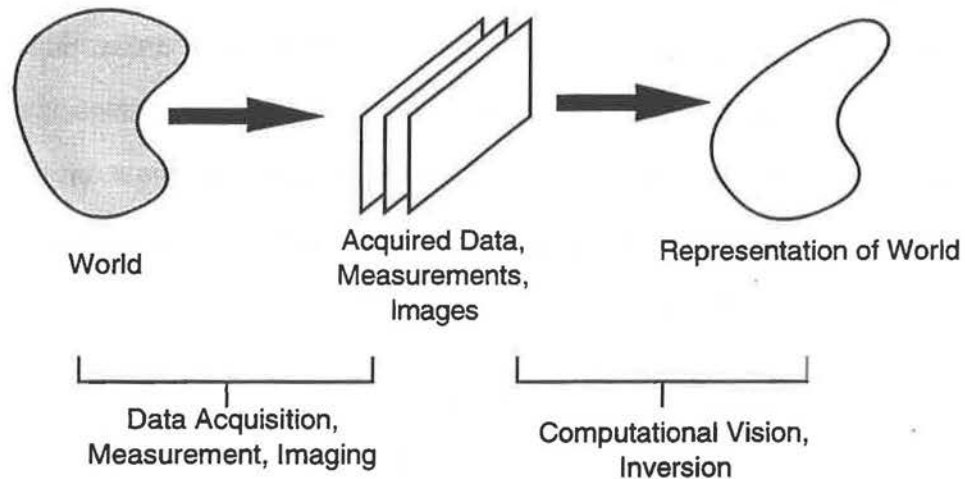


Figure 2.1: Model of computational vision.

Figure 2.1 shows a model of computational vision systems. On the left-hand side of the figure is a world that is the setting for a task. Some data acquisition process acquires measurements of the world. These measurements (or data) are images for visual systems, but they may be raysum data or ultrasound traces for NDE systems. The goal of the computational vision system (right-hand side of Figure 2.1) is to invert the acquired

data to a representation of the world. Thus, computational vision is the solution of inverse problems where the solution is a representation of the world, suitable for the task at hand.

Invariably, the inversion problem is ill-posed because it has no unique solution. One approach to dealing with this ill-posedness is to create a new problem by applying constraints to the original. The constraints restrict the set of solutions so that the modified problem is well-posed. Ill-posed problems occur often in computer vision. For examples see Poggio, Torre and Koch [30] or Bertero, Poggio and Torre [5].

When there is only one source of data, *a priori* assumptions about the nature of the world, and consequently about the nature of the solution, provide constraints. The assumptions are a convenient way of restricting solutions in the face of ambiguous data and may not always hold. Where assumptions are not valid, data fusion offers an alternative. Constraint-based data fusion systems acquire additional data from independent sources and use the resulting solutions to constrain the original inverse problem. The notion of fusing data from multiple sources as constraints is the essence of constraint-based data fusion. It strives to improve the quality of solutions to inverse problems by using independent data sources as constraints, without falling prey to erroneous *a priori* assumptions.

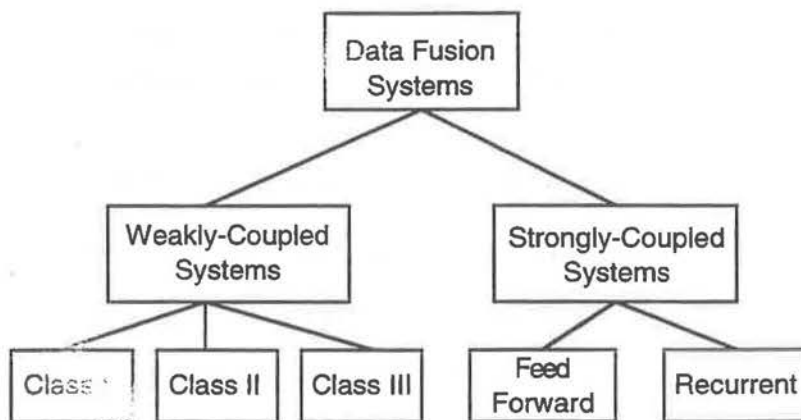


Figure 2.2: Mark and Yuille taxonomy of data fusion systems.

Clark and Yuille give a taxonomy of data fusion systems (shown in Figure 2.2). Fusion systems consist of a set of modules where each module accepts some input data and performs an inversion subject to some constraints. The first division in the taxonomy is between weakly-coupled and strongly-coupled systems. Weakly-coupled systems are distinguished by their independent modules, i.e., there is no sharing of information between modules. Strongly-coupled systems are the opposite; for a system to be strongly-coupled at least one module must constrain the output of another.

There are three classes of weakly-coupled systems, all shown in Figure 2.3. Modules in a class I system are stable and produce unique solutions independently. The ultimate solution produced by the system is a weighted sum of the solutions from the individual modules, with the weights determined by the relative reliability of the modules. Solutions from more reliable modules are weighted greater than those from less reliable ones. Class II weakly-coupled fusion modules do not produce a unique solution. It is the function of the fusion module to combine the module outputs into a unique solution (the fusion module may use *a priori* constraints to help it). An example is the measurement of electrical resistance by separate voltage and current measurements. Neither measurement yields a unique resistance value, but the combination of the two does. In this case the fusion module performs the calculation $R = \frac{V}{I}$. Class III fusion combines classes I and II. Modules in a class III system do not produce unique solutions so a fusion module combines the module outputs as in a class II system, while solutions from the modules are not equally reliable so they are weighted as in class I.

Figure 2.4 shows schematically feed-forward and recurrent strongly-coupled fusion systems. Their common feature is that at least one module constrains the output of another module. Feedforward is the simplest of the strongly-coupled systems. The system feeds forward the output of one module to provide constraints to a second module. Output from the second module is the output of the system. In the recurrent system the

output from each module feeds back to other modules in the form of constraints. The recurrent system is the most complex of the fusion methods.

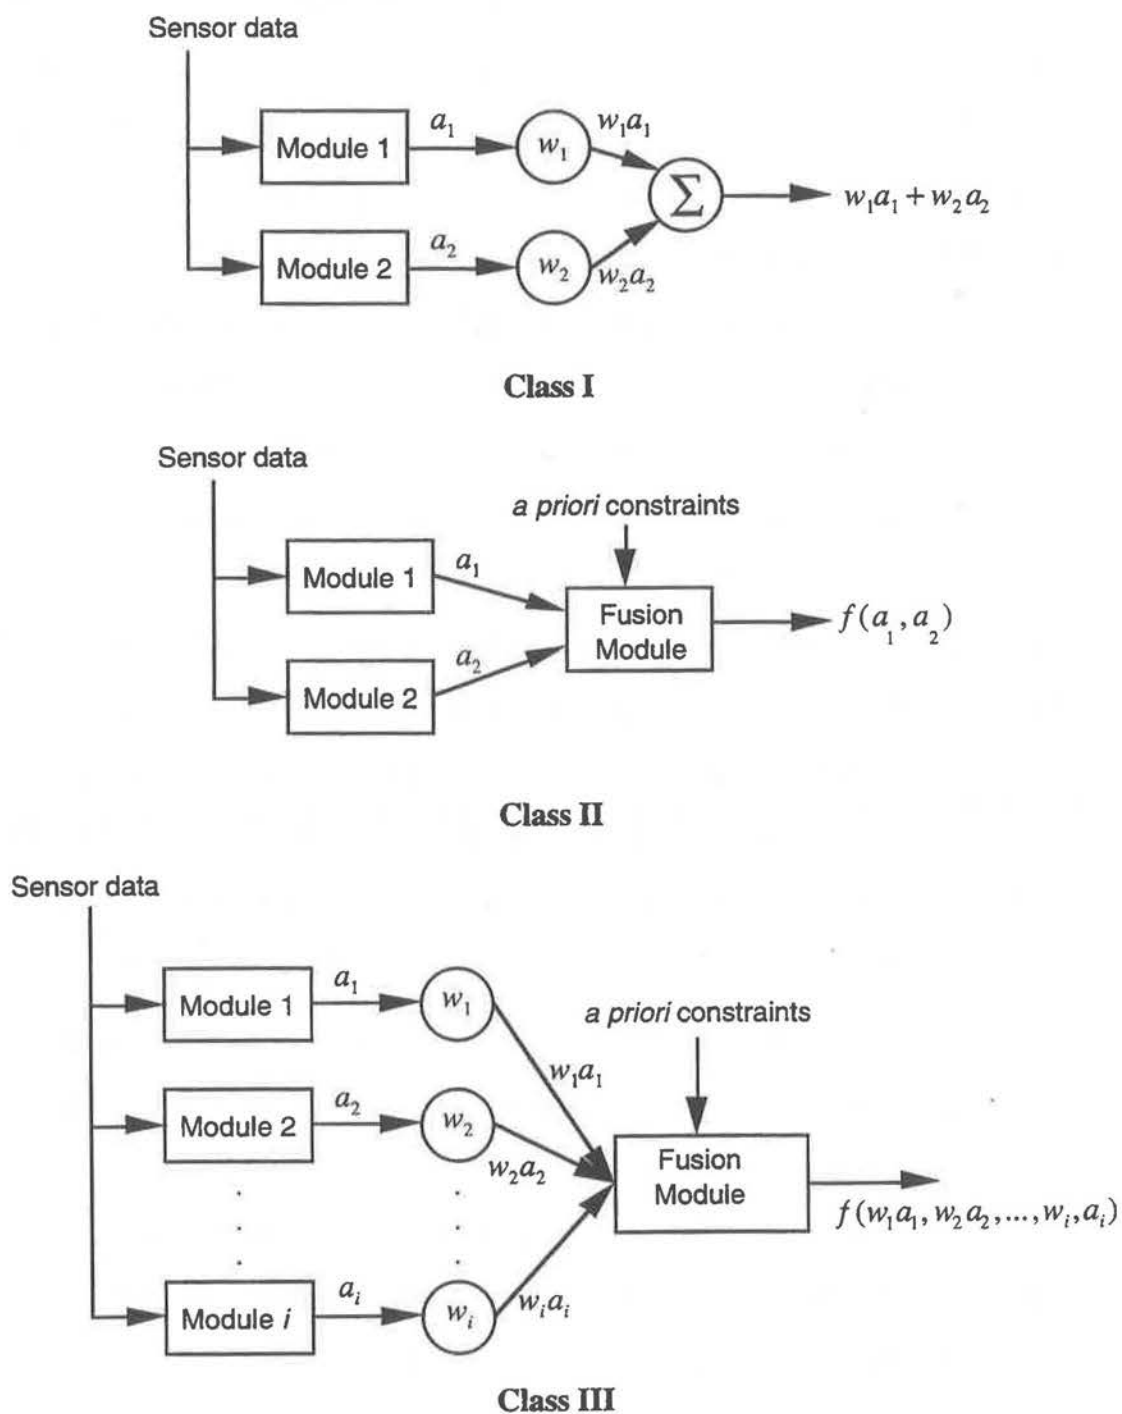


Figure 2.3: Classes of weakly-coupled data fusion systems.

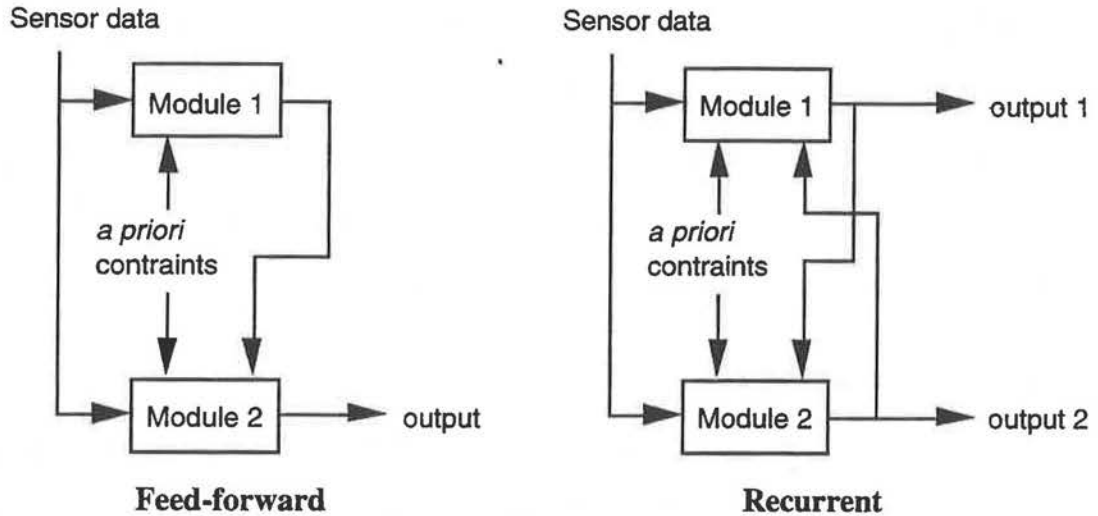


Figure 2.4: Strongly-coupled data fusion systems.

Strongly-coupled fusion is synonymous with constraint-based data fusion.

Whenever an inversion module relies on *a priori* assumptions to constrain the solution of an inverse problem, constraint-based data fusion can reduce reliance on assumptions. It does so by replacing constraining assumptions with constraints derived from other data sources. There is, of course, no guarantee that by using fusion the inverse problem will become well-conditioned or even well-posed. However, with careful attention to the nature of the data and its use as constraints, it is possible to improve the quality of solution by avoiding reliance on possibly erroneous *a priori* assumptions.

2.2. Computed Tomography

As x-ray radiation passes through an object the object attenuates it.

Measurements of the attenuation approximate raysums of the object. Computed tomography (CT) is the inverse problem of reconstructing a cross section of an object from a set of coplanar raysums. The solution, i.e., the cross section, is a distribution of x-ray linear attenuation coefficients.

This section presents the fundamentals of CT starting with a description of x-ray absorption and raysum measurement. The basic apparatus for CT illustrates the use of x-rays to acquire raysum data for reconstruction. These data are a discrete measurement of

the Radon transform of the object, the basis for CT reconstruction. Inversion of the continuous Radon transform is an ill-posed problem because the solution does not vary continuously with the data. In the case of discrete data, reconstruction is ill-posed because there is not a unique solution. In both cases the ill-posedness is easily handled by assuming that the solution must be band-limited.

2.2.1. X-ray Absorption

Figure 2.5 shows schematically x-ray radiation penetrating an object along a ray. The object absorbs some of the radiation, attenuating the intensity of the emerging radiation. The following equation models the attenuation process [19]:

$$I = I_0 e^{-\int_l \mu(x) dx}, \quad (2-1)$$

where I_0 is the incident radiation intensity, I is the emerging intensity, l is the path of the radiation, x is the distance along that path, and $\mu(x)$ is the x-ray linear attenuation which varies along l . Rearranging Equation (2-1) gives:

$$\ln\left(\frac{I}{I_0}\right) = -\int_l \mu(x) dx, \quad (2-2)$$

which shows the linear relationship between log attenuation, $\ln\left(\frac{I}{I_0}\right)$, and linear

attenuation. The integral on the right-hand side of Equation (2-2) is called a *raysum* or *projection* of $\mu(x)$ (raysum is used for the remainder of this thesis to avoid confusion with projections in the POCS method described in Chapter 4) and is the basis of data acquisition for CT reconstruction.

If the linear attenuation function, $\mu(x)$, is piece-wise continuous, Equation (2-2) becomes:

$$\ln\left(\frac{I}{I_0}\right) = -\sum_i \mu_i \Delta x_i, \quad (2-3)$$

where μ_i and Δx_i are the linear attenuation and the thickness of the i^{th} component of the object along the path of radiation. Thus x-ray raysums are linear functions of linear attenuation coefficients and a linear system adequately models x-ray absorption.




Figure 2.5: Schematic of x-ray absorption.

2.2.2. Basic Apparatus for CT Data Acquisition

Figure 2.6 shows a schematic of the basic apparatus for a CT data acquisition system [22] [34]. The acquisition procedure is:

1. Collect projection data for one orientation:
 - (a) Project a thin beam of radiation (from the source) through the specimen and measure the output intensity (with the sensor).
 - (b) Move the source and sensor in a direction perpendicular to the direction of the radiation and repeat 1(a).
2. Rotate the specimen and repeat step 1.
3. Repeat steps 1 and 2 to collect raysums through π radians.

Collimators (not shown in Figure 2.6) confine the x-rays to a thin beam so that raysums measure only a thin slice of the specimen. Successive adjacent slices sample the specimen in three dimensions if desired.

The above process measures discrete samples from the Radon transform of the specimen. This transform, described in the next section, is fundamental to CT.

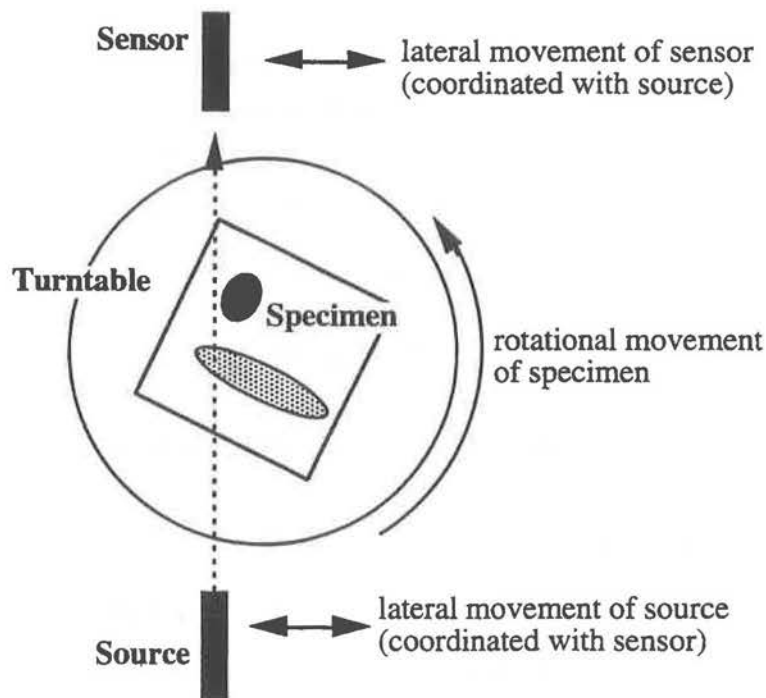


Figure 2.6: Schematic of basic computed tomography apparatus

2.2.3. The Radon Transform and the Fourier Slice Theorem

The *Radon transform*, first explored by Radon in 1917 [32], is fundamental to CT and forms the basis of the *Fourier slice theorem*. Together, the Radon transform and the Fourier slice theorem provide a mathematical basis for understanding CT reconstruction.

Define a Cartesian coordinate system with its origin at the centre of rotation in Figure 2.6. Based on this coordinate system, the Radon transform of a function $\mu(x, y)$, $[\mathcal{R}\mu](l, \theta)$, is [22]:

$$[\mathcal{R}\mu](l, \theta) = \int_{-\infty}^{\infty} \mu(l \cos \theta + \xi \sin \theta, \xi \cos \theta + l \sin \theta) d\xi. \quad (2-4)$$

where the integral on the right-hand side is a raysum of $\mu(x, y)$ along a ray at angle θ with the y axis and distance l from the origin. Thus, the Radon transform of an object is the set of all its raysums. The apparatus of Figure 2.6 measures a discrete approximation of the Radon transform. Measurements are approximate because raysums of Equation (2-

4) have infinitesimal thickness (not possible in practice), and because other practical problems in radiography such as scatter and beam hardening confound data acquisition.

The Fourier slice theorem (also known as the projection theorem) relates the Radon transform to the Fourier transform. A short intuitive explanation of the theorem follows; a proof is available from many sources, e.g., [22] or [34]. The separable nature of the Fourier transform allows separate computation in x and y directions. Suppose computation begins with one-dimensional transforms in the y direction, then the zero-frequency component of the resulting transforms are raysums for rays parallel to the y axis. Completion of the transform on the zero frequency data, i.e., raysum data, gives a slice of the Fourier transform of the specimen in the x direction. Expressed mathematically, the Fourier slice theorem is [22]:

$$[\mathcal{F}_l \mathcal{R}\mu](\omega_l, \theta) = [\mathcal{F}_{x,y} \mu](\omega_l \cos \theta, -\omega_l \sin \theta), \quad (2-5)$$

where \mathcal{F}_l is the one-dimensional Fourier transform operator in l , ω_l and θ are polar coordinates of the two-dimensional Fourier space, and $\mathcal{F}_{x,y}$ is the two-dimensional Fourier transform operator in Cartesian coordinates. The Fourier slice theorem says that measuring a set of parallel raysums of an object is equivalent to sampling a slice of the Fourier transform along a line through the origin, perpendicular to the direction of the raysums.

2.2.4. CT Reconstruction

The Fourier slice theorem allows a short and simple explanation of why CT is possible. According to the theorem, step one of the data acquisition procedure samples a slice of the specimen in Fourier space. Steps 2 and 3 sample slices through π radians, thus sampling the entire Fourier space. A complete sampling of the Fourier space means that all the data necessary for reconstruction are available. Although tomographic reconstruction by assembling the Fourier transforms from raysums is possible, this method is not popular. Of the many methods available, the convolution-back projection

(or some variant) is most popular, and the work presented here focuses on other methods. Nevertheless, the Fourier slice method gives a conceptual model useful for understanding how raysums sample a function, and the ill-posed nature of limited-angle CT.

Inversion of the Radon transform was first shown by Radon [32] and his result may be found in more recent literature, e.g. [22] and [34]. The solution to the inverse problem exists, but does not depend continuously on the data. Thus CT reconstruction is an ill-posed problem. In practice it is not possible to acquire the continuous Radon transform, only a finite number of samples, usually distributed uniformly throughout (l, θ) . The practical problem of tomographic reconstruction from discretely sampled raysum data is also ill-posed because there is no unique solution.

Commonplace use of CT suggests that its ill-posedness can be managed, as shown below. The following linear system models raysum acquisition:

$$R\mathbf{x} = \mathbf{y} \tag{2-6}$$

where \mathbf{y} is a vector of projection data, \mathbf{x} is a vector of linear attenuation coefficients, and R is a linear matrix operator that models x-ray raysum acquisition. (i.e. a discrete version of the Radon transform operator, \mathcal{R}). Assume that R samples the Radon transform uniformly. CT reconstruction is the inverse problem of solving for \mathbf{x} given raysum data \mathbf{y} .

The matrix R has a non-trivial null space, i.e., the set of vectors, \mathbf{n} , that satisfy $R\mathbf{n} = \mathbf{0}$ is not limited to the zero vector. Fortunately, although the null space is non-trivial, the non-trivial components of the null space are high-frequency oscillations. Assuming that the solution should be band-limited modifies the ill-posed problem into one that is well-posed. This is normally a reasonable assumption because it is known *a priori* that the oscillations do not exist or they are not of interest even if they do exist. For reconstruction from the continuous transform, the same band-limiting assumption yields a well-posed problem by ensuring that the solution varies continuously with the data.




Figure 2.7: Example of CT showing lines of projection on a two-by-two grid.

A small example illustrates the ill-posed behaviour of CT reconstruction from a discrete Radon transform. Consider the two-by-two discrete image shown in Figure 2.7. Arrows show the sampled raysums which represent a uniform sampling of the Radon transform on a two-by-two grid. R for this sampling scheme is:

$$R = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix}.$$

The square roots of the eigenvalues of $R^T R$ are the singular values of R (see Chapter 4 for a description of singular value decomposition, SVD) and are 2 , $\sqrt{2}$, $\sqrt{2}$, and 0 . The zero singular value indicates that R has a non-trivial null space as expected.

Suppose we have the following 2 by 2 cross section of a specimen with pixel values:

$$\mathbf{x} = \begin{bmatrix} 1 & 2 \\ 1 & 2 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 1 \\ 2 \end{bmatrix}.$$

Application of the forward operator, R , to \mathbf{x} gives the data:

$$\mathbf{y} = \begin{bmatrix} 3 \\ 3 \\ 2 \\ 4 \end{bmatrix}.$$

From this value of \mathbf{y} , the minimum-norm solution, $\hat{\mathbf{x}}$, (using SVD of R) is:

$$\hat{\mathbf{x}} = \begin{bmatrix} 1 \\ 2 \\ 1 \\ 2 \end{bmatrix} = \begin{bmatrix} 1 & 2 \\ 1 & 2 \end{bmatrix}.$$

SVD also gives the vectors in the null-space of R . In this case there is one singular value of zero so there is only one dimension in the null space in the direction of:

$$\mathbf{n} = \begin{bmatrix} -1/2 \\ 1/2 \\ 1/2 \\ -1/2 \end{bmatrix} = \begin{bmatrix} -1/2 & 1/2 \\ 1/2 & -1/2 \end{bmatrix}.$$

As expected, \mathbf{n} is an oscillatory image. The reconstructed image, $\hat{\mathbf{x}}$, is accurate because the image \mathbf{x} is perpendicular to \mathbf{n} , as the minimum-norm solution must also be. In this case the minimum-norm requirement selects a band-limited solution by rejecting the high-frequency components of the null space. If the original image, \mathbf{x} , were to contain some component of \mathbf{n} , then the minimum-norm solution, $\hat{\mathbf{x}}$, would be in error. CT reconstruction from a uniform sampling is only feasible if the cross section does not contain high-frequency components, or if those high-frequency components exist, they are not important.

2.3. Limited Angle Computed Tomography

The previous section showed that CT reconstruction is possible from a uniformly distributed subset of the continuous Radon transform. Reconstruction is an ill-posed problem which becomes soluble by assuming the solution is band-limited. This section presents the more difficult problem of reconstruction from a set of raysums which is incomplete not simply because of discrete sampling, but also because raysums are only available for a limited range of angles. This problem arises because physical constraints imposed on the CT apparatus, either by the specimen or its environment, prevent raysum acquisition throughout the required range of $-\frac{\pi}{2} \leq \theta \leq \frac{\pi}{2}$ (see Figure 2.8 for examples

taken from [40]). The name given to this particular problem is *limited-angle* CT. Figure 2.8 (a) is important here because it represents the situation with inspection of aircraft control surfaces.




Figure 2.8: Example cases for which only limited-angle raysums are available: (a) object is very long in one direction, and (b) other objects obstruct scanning.

2.3.1. Ill-Posed Nature of Limited Angle Computed Tomography

The Fourier slice theorem states that sampling a specimen with raysums is equivalent to sampling slices in Fourier space. It is obvious, therefore, that in order to sample a specimen completely one must measure raysums through a range of π radians, i.e. $-\frac{\pi}{2} \leq \theta \leq \frac{\pi}{2}$. As illustrated in Figure 2.8, circumstances arise where projections cannot be sampled throughout a full range, i.e., $-\Theta \leq \theta \leq \Theta$, and $0 < \Theta < \frac{\pi}{2}$, thus

leaving a portion of the Fourier space unsampled. When the full range of angles is available, then the sampled Fourier space covers a disc-shaped region. In the limited angle case, a portion of the Radon transform is missing and the sampled Fourier space covers only the butterfly-shaped region shown in Figure 2.9.




Figure 2.9: Effect of limited-angle Radon transform on Fourier domain sampling: (a) range of angles sampled (shaded region) and (b) effective sampling of Fourier domain (shaded region is sampled).

Fourier space sampling illustrates clearly the ill-posed nature of limited-angle CT. Ill-posedness occurs because there is no unique solution, but instead of a small null space containing only high-frequency oscillations, the null space also contains components corresponding to the unsampled Fourier space. Although both the full-range and limited-angle problems are ill-posed, we say the limited-angle problem is more ill-posed because it has a greater number of dimensions in its null space. Davison [10] and later Louis [23] use singular value decomposition to explore the ill-conditioned nature of limited-angle CT. Louis concludes “... *the limited angle problem is much more ill-posed than the full range problem. ... Clearly the algorithms applied in this situation have to be carefully designed and it is desirable to use as much a priori knowledge as possible.*” Louis also shows some images which are vectors in the null space of R . These null-space images show what types of artifacts may appear in limited-angle CT reconstructions.

2.3.2. Reconstruction from Limited-Angle Data

The Fourier slice theorem shows that limited-angle raysum data do not sample the Fourier space completely. Therefore, any reconstruction method based on limited angle data necessarily fills in the unsampled regions of Fourier space. Possibilities for filling in are unlimited which is precisely the source of the ill-posedness cited above. *A priori* assumptions can give constraints that lead to a unique solution. Current techniques for limited-angle CT fall into these three categories distinguished by the *a priori* constraints that guide the reconstruction to a unique solution and complete the Fourier space:

1. fill the unsampled Fourier space with zeros,
2. impose spatial support constraints on the solution, and
3. extrapolate the sampled Fourier space data.

The first category, fill with zeros, is the minimum-norm solution. Of all possible values for the unsampled region, all zeros contains the least amount of energy and therefore must be the minimum-norm solution. The second category, impose spatial support constraints is similar. Constraining spatial support does not necessarily determine the missing data on its own. One still seeks a minimum-norm solution as in category one, but with the additional constraint added. The last category, extrapolate the Fourier space, assumes that the available data establish all important trends in the Fourier space, an assumption that is not necessarily valid. Examples from all three categories follow.

Section 2.2.4 showed briefly the minimum-norm solution to a system of equations as computed by SVD. Refer to the example in Section 2.2.4. The discrete raysum operator for the limited angle case is:

$$R = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}.$$

with singular values of $\sqrt{2}$, $\sqrt{2}$, 0, and 0. Restricting the angle adds another singular value of zero to the problem and increases the number of dimensions in the null space

from one to two. Two orthogonal vectors form a basis for the null space, one identical to that for the full-range case, i.e.:

$$\mathbf{n}_1 = \begin{bmatrix} -1/2 \\ 1/2 \\ 1/2 \\ -1/2 \end{bmatrix} = \begin{bmatrix} -1/2 & 1/2 \\ 1/2 & -1/2 \end{bmatrix}$$

and the other, which is entirely attributable to the limited-angle sampling, is:

$$\mathbf{n}_2 = \begin{bmatrix} -1/2 \\ 1/2 \\ -1/2 \\ 1/2 \end{bmatrix} = \begin{bmatrix} -1/2 & 1/2 \\ -1/2 & 1/2 \end{bmatrix}$$

The new null space component is a step across a vertically oriented edge. If \mathbf{x} is:

$$\mathbf{x} = \begin{bmatrix} 1 & 2 \\ 1 & 2 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 1 \\ 2 \end{bmatrix},$$

then the raysum data, \mathbf{y} , is:

$$\mathbf{y} = \begin{bmatrix} 3 \\ 3 \end{bmatrix},$$

and the minimum-norm reconstruction, $\hat{\mathbf{x}}$, is:

$$\hat{\mathbf{x}} = \begin{bmatrix} 3/2 \\ 3/2 \\ 3/2 \\ 3/2 \end{bmatrix} = \begin{bmatrix} 3/2 & 3/2 \\ 3/2 & 3/2 \end{bmatrix}.$$

Note that the vertical edge present in the original image is missing in the reconstruction because the minimum-norm solution omits the portion of \mathbf{x} in the null space, i.e., the vertical edge. Suppose instead that \mathbf{x} is:

$$\mathbf{x} = \begin{bmatrix} 1 & 1 \\ 2 & 2 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 2 \\ 2 \end{bmatrix},$$

then the minimum-norm reconstruction is:

$$\hat{\mathbf{x}} = \begin{bmatrix} 1 \\ 1 \\ 2 \\ 2 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 2 & 2 \end{bmatrix}$$

The reconstruction is accurate in this case because \mathbf{x} is orthogonal to the null space of R . The salient point here is that minimum-norm reconstruction works so long as the function being reconstructed contains no component in the null space, and the success or failure of reconstruction is case-dependent.

As one might reasonably guess, any reconstruction method intended for full-range data can compute the minimum-norm solution by assuming that all the missing raysums are zero (Tuy [43]). Consequently, one does not need to use the unwieldy and slow singular value decomposition to compute the minimum-norm reconstruction from limited-angle data.

The second category of limited-angle CT uses the often reasonable assumption that an object lies within known spatial bounds. This is a common approach to improving the quality of limited-angle CT reconstruction. Grünbaum [16], for example, makes the *a priori* assumption that the specimen is contained within a rectangular region of support and linear attenuation values outside that region must be zero. The limited spatial support effectively interpolates the Fourier domain, and so there is an implicit assumption to justify interpolation that the Fourier space is smooth. Grünbaum gives results for synthetic data but there is no indication of the circumstances for which the effective interpolation is valid.

Although the Grünbaum rectangular region improves upon minimum-norm solutions, in some circumstances it is possible to further constrain spatial support. Sato *et al.* [33] employ a triangular region of support to reconstruct a synthetic image of the letter 'A' (see Figure 2.10). They perform reconstruction using an iterative method that alternates between the Fourier and spatial domains, applying constraints in each:

triangular support in the spatial domain and a butterfly-shaped region in the Fourier space. They provide a proof of convergence, but not a rate of convergence. Their method is, in fact, the method of projection onto convex sets (see Chapter 4) which guarantees convergence. As is the case for Grünbaum, examples are for synthetic data only. Such tight restrictions on spatial support are useful but not trivial to implement in practice.




Figure 2.10: Triangular region of spatial support from Sato *et al.* [33] method for limited-angle CT reconstruction of the letter 'A'.

Grünbaum and Sato *et al.* are just two examples of using spatial support restrictions to aid limited angle CT reconstruction. There are more examples in the literature, such as Dusaussoy and Abdou [11] and Tam *et al.* [40]. It is interesting to note that Tam's algorithm appears to be identical to that of Sato *et al.*

The third category of reconstruction method is interpolation of the Fourier domain. Methods in this category extrapolate data in the Fourier domain to fill in the unsampled regions. Soumekh [36] models the Fourier domain with a periodic function that can be represented by a Fourier series expansion. He then makes the *a priori* assumption that only a finite number of coefficients are necessary to represent the series, i.e., coefficients for high frequencies are zero or insignificant. The available Fourier space data are sufficient to determine the finite number of coefficients. The truncated Fourier series then interpolates to fill in the missing Fourier domain data. Although the interpolation by Soumekh does not translate into simple assumptions about spatial support, it does restrict the class of objects that may be accurately reconstructed. Reconstruction

is successful with synthetic examples, but the extrapolation requires that the available data establish all important trends. Heiskanen *et al.* [20] give another example of direct interpolation of raysum data.

Minimum norm solutions assume that there is no important information in the null-space of the raysum operator. Spatial support restrictions and direct extrapolation methods assume that the sampled data establish all trends in the specimen, i.e., everything you need lies in the range of the limited-angle raysum operator. The common feature among all categories is that none can properly reconstruct a function unless the null space is either trivial or unimportant. In fact, the following is true for limited-angle CT:

1. limited-angle CT requires filling in the unsampled Fourier space of a function,
2. there is no universally correct way to fill in the missing data, and
3. *a priori* assumptions should guide the reconstruction as much as possible.

This thesis proposes moving beyond *a priori* assumptions in limited-angle CT and using data from other sources to constrain the reconstruction, i.e., constraint-based data fusion for limited-angle CT. Data fusion fills in the missing data correctly, even when *a priori* assumptions for interpolation are not valid.

2.4. Ultrasound

The previous section showed that CT reconstruction from limited-angle data is ill-posed and has a large null space. To reliably reconstruct tomographic images of a specimen, limited-angle CT methods rely significantly on *a priori* assumptions to constrain the solution. Constraint-based data fusion avoids reliance on such assumptions by using other data sources to provide constraints. Ultrasound is a commonplace NDE technique that gives data complementary to x-rays and has the potential to provide constraints to limited-angle CT. This section presents some background on ultrasound as a prelude to its use in a limited-angle CT data fusion system.

2.4.1. Forward Models

In ultrasonic testing, a transducer transmits a high-frequency sound wave into the surface of a specimen. The sound travels through the specimen unchanged until it encounters an interface between materials of different acoustic impedance. At the interface some energy is reflected and some is transmitted. Reflections (echoes) at the surface contain information about the internal structure of the specimen. Take, for example, a single interface between two layers in a layered specimen; the upper and lower layers have acoustic impedances z_1 and z_2 respectively. The amplitude of a reflected signal is r , the reflection coefficient for the interface, times the amplitude of the incident signal, where r , is given by [24]:

$$r = \frac{z_2 - z_1}{z_2 + z_1}. \quad (2-7)$$

A negative value of r indicates a phase reversal. The following expression gives the echoed signal:

$$y(t) = s(t) * w(t) + r(t), \quad (2-8)$$

where $y(t)$ is the echoed signal, $s(t)$ is the impulse response of the material, $w(t)$ is the sound wavelet introduced at the surface, and $r(t)$ is noise (* is the convolution operator). $s(t)$ contains the desired structural information. It is a train of impulses, each impulse corresponding to an interface in the specimen. The magnitude and sign of the impulses depends on r for the interface and the attenuation in the path of the sound to and from that interface. Equation (2-8) gives a simple, non-parametric (i.e. material parameters are not specified), forward model of the ultrasonic process [25].

More sophisticated parametric models define the output of an ultrasonic system based on material properties. These models fit into two categories: (1) equal time slices and (2) variable time slices [25]. The models assume a lossless layered medium like that shown in Figure 2.11. $m(t)$ is the signal injected into the specimen and $y(t)$ is the output

as in Equation (2-8). $u_j(t)$ and $d_j(t)$ are the upward and downward traveling signals in the j^{th} layer.

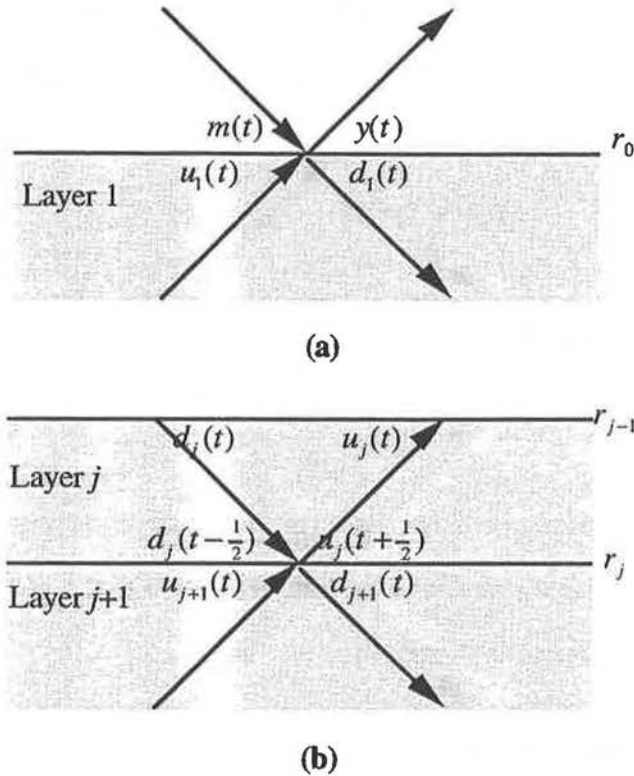


Figure 2.11: Schematic of a lossless layered ultrasonic medium: (a) top layer and (b) internal layers.

The equal time slice model divides the medium into layers of equal travel time. Each layer is a half sample period thick, which makes the two-way travel time equal to one sample period per layer. The set of reflection coefficients, $r_0, r_1, r_2, \dots, r_n$, where n is the number of samples, completely defines the model. The following equations describe the ultrasonic system:

$$\begin{aligned} d_{j+1}(t) &= (1 + r_j) d_j(t - \frac{1}{2}) - r_j u_{j+1}(t) \\ u_j(t + \frac{1}{2}) &= r_j d_j(t - \frac{1}{2}) + (1 - r_j) u_{j+1}(t), \quad j = 1, 2, \dots, n \end{aligned} \quad (2-9)$$

for internal layers, and

$$\begin{aligned} d_1(t) &= (1 + r_0) m(t) - r_0 u_1(t) \\ y(t) &= r_0 m(t) + (1 - r_0) u_1(t) \end{aligned} \quad (2-10)$$

for the top surface.

The variable time slice model divides the specimen into slices corresponding to the actual material layers. Reflection coefficients, $r_0, r_1, r_2, \dots, r_n$, and the travel times, $\tau_0, \tau_1, \tau_2, \dots, \tau_n$, where n is the number of layers in the specimen, completely define the model. The following equations describe the ultrasonic system:

$$\begin{aligned} d_j(t + \tau_j) &= (1 + r_{j-1})d_{j-1}(t) - r_{j-1}u_j(t) \\ u_j(t + \tau_j) &= r_j d_{j-1}(t) + (1 - r_j)u_{j+1}(t) \end{aligned} \quad (2-11)$$

for the internal layers, and

$$\begin{aligned} d_0(t) &= m(t) \\ y(t) &= u_1(t) \end{aligned} \quad (2-12)$$

define the input and output (in this case for a unidirectional sensor just beneath the upper surface). Variable time slices have the inherent advantage that the parameters correspond directly to the structure of the specimen.

2.4.2. Ultrasound Inversion

Both the non-parametric (Equation (2-8)) and the parametric (Equations (2-9) to (2-12)) models have corresponding inverse problems reviewed here.

Non-parametric modeling leads to the inverse problem of ultrasonic deconvolution. Equation (2-8) expressed with linear matrix operators is:

$$\mathbf{y} = \mathbf{W}\mathbf{s} + \mathbf{r}, \quad (2-13)$$

where \mathbf{y} is the vector of output data, \mathbf{s} is the vector of impulses, and \mathbf{r} is an error vector. \mathbf{W} is a matrix operator which performs a convolution with the ultrasound wavelet. The least-squares approach to solving Equation (2-13) gives the optimization problem of minimizing the quadratic norm of the residual, $\mathbf{W}\mathbf{s} - \mathbf{y}$, i.e., minimizing the objective function $E(\mathbf{s})$ where:

$$E(\mathbf{s}) = \|\mathbf{W}\mathbf{s} - \mathbf{y}\|^2. \quad (2-14)$$

Unfortunately W is not full column rank and the optimization is ill-posed; the solution is not unique. One way to deal with the ill-posedness is to modify the problem by regularization. Rather than solve for \mathbf{s} from Equation (2-13), solve for \mathbf{s} from:

$$\begin{bmatrix} W \\ \lambda D \end{bmatrix} \mathbf{s} = \begin{bmatrix} \mathbf{y} \\ \mathbf{0} \end{bmatrix}. \quad (2-15)$$

Solution of Equation (2-15) by least-squares minimizes the objective function:

$$E(\mathbf{s}) = \|\mathbf{W}\mathbf{s} - \mathbf{y}\|^2 + \lambda^2 \|\mathbf{D}\mathbf{s}\|^2. \quad (2-16)$$

D is chosen so that the matrix $\begin{bmatrix} W \\ \lambda D \end{bmatrix}$ is full column rank. Minimizing $E(\mathbf{s})$ in (2-16)

gives the following symmetric positive definite system:

$$(\mathbf{W}^T \mathbf{W} + \lambda^2 \mathbf{D}^T \mathbf{D}) \mathbf{s} = \mathbf{W}^T \mathbf{y}, \quad (2-17)$$

which is essentially the classic Wiener deconvolution [24]. In Wiener deconvolution the operator D corresponds to multiplication by the noise spectrum in the frequency domain.

A more recent method called l_1 -norm deconvolution uses l_1 norms rather than l_2 and assumes that $s(t)$ contains impulses only [45]. Based on experimentation, developers of the l_1 -norm method state that the l_1 norm gives superior results and is faster when compared to the l_2 -norm, and that the assumption that $s(t)$ contains only impulses leads to improved temporal resolution.

Non-parametric inversion only recovers the impulse response, $s(t)$, and no direct structural information. Inversion with a parametric model allows recovery of structural information in the form of an acoustic impedance profile. There are two basic methods for parametric inversion: direct inversion and optimization.

The first, direct inversion, solves for the uppermost interface from the first data point (Equation (2-10)). It then solves for subsequent layers from subsequent data points and solutions to upper layers in top-to-bottom order. There are two major problems with direct solution. First, it assumes that the data are deconvolved, i.e., the data are the series of impulses $s(t)$ rather than the convolved data $y(t)$. Second, the method is ill-

conditioned. Small errors in a layer propagate to lower layers and grow rapidly in amplitude.

In the second type of parametric inversion, optimization, one creates an objective function:

$$E(\mathbf{z}) = \|U(\mathbf{z}) - \mathbf{y}\|^2, \quad (2-18)$$

where \mathbf{z} is the impedance profile and $U(\mathbf{z})$ is the forward ultrasonic process, and then minimizes it among the possible \mathbf{z} . $U(\mathbf{z})$ is non-linear so $E(\mathbf{z})$ is non-convex and difficult to minimize. The minimum may not be unique and so the optimization is ill-posed. In spite of these problems, optimization is superior to direct inversion.

Examples of parametric inversion are given by Goupillaud [15], who uses an equal time slice model for ultrasonic inversion and more recently Habibi-Ashrafi and Mendel [18], Mendel and Goutsias [26], and Zala and McRae [45] who use a variable time slice model. Most current literature on ultrasonic inversion treats the material as lossless although in reality this is not the case. Usually the specimen absorbs some sound energy in a manner analogous to x-ray absorption. The absorption coefficient is difficult to measure precisely and ultrasonic inversion is extremely sensitive to it. So ultrasonic inversion for an absorbing medium is even more complex than for an imaginary lossless medium. Zala and Churchill [46] give an excellent review of ultrasonic inversion methods as well as the practical problems encountered when the medium is not lossless.

2.4.3. Minimal Ultrasound for Fusion with Computed Tomography

Ultrasound data complement x-ray data because they measure features that x-rays cannot. Whereas raysums are insensitive to discontinuities along the path of radiation, ultrasound explicitly detects discontinuities in its path. This complementary nature of x-ray and ultrasound portends their effective fusion in limited-angle CT.

Ultrasound inversion for a layered lossless specimen is difficult but possible. In the more realistic situation of an absorptive specimen ultrasonic inversion is tenuous; the solution is extremely sensitive to the absorption coefficient which is generally unknown and highly variable.

Fortunately, CT reconstruction of sandwich specimens by fusion with other data sources does not require the full potential of ultrasound. Instead, as shown in Chapter 3, only measurements of the face sheets are necessary. As long as the external layer is composed of a homogenous slab, thickness measurement only requires identification of two reflections in the ultrasound signal: one from the top surface and one from the bottom. The two reflections are distinct making their identification easy and thickness measurement simple.

There are several reasons why thickness measurement might not be simple. Delaminations in a composite face sheet can prevent detection of bottom surface reflections. Face sheets in some regions of a specimen may have multiple layers and so the reflection from the bottom must be separated from other reflections. In spite of these confounding factors, it is reasonable to assume that ultrasound can provide sufficient thickness measurements of the face sheet of a sandwich specimen to allow accurate interpolation of thickness data over the entire face sheet.

2.5 Chapter Summary

Section 2.1 of this chapter introduced the *constraint-based* data fusion model of Clark and Yuille [9]. Their elegant approach implements fusion by using the output of one data inversion module as constraints for another. Constraint-based data fusion provides a coherent framework within which to build data fusion systems.

The basic inverse problem of x-ray raysum inversion is known as CT (computed tomography). Conventional CT is ill-posed, but the ill-posedness is easily handled by suppressing high-frequency components in the reconstructed solution. Practical

considerations often dictate that raysum data cannot be collected over the full range of angles in the Radon transform. In these limited-angle cases the inverse problem is still ill-posed, but the null space is much larger and reconstruction requires additional constraints. Many methods exist to solve limited-angle CT. They all rely on *a priori* assumptions (which may not be valid) to constrain the solution. This reliance portends a more-accurate solution from constraint-based data fusion when *a priori* assumptions are not valid.

In the field of non-destructive evaluation ultrasound is commonplace. It provides a source of sensory data that complements x-rays and may provide constraints to limited-angle CT. Many types of inspection are possible with ultrasound, but the limited-angle CT problem of this thesis requires only the ability to measure the thickness of a face sheet in a sandwich structure.

Chapter 3 shows that limited-angle CT fails to sample adequately the face sheet structure of a sandwich specimen causing limited-angle CT reconstruction methods constrained by generalized *a priori* assumptions to fail. The chapter then proposes a novel system for limited-angle CT that fuses ultrasound data with x-ray data to accurately reconstruct sandwich structures.

Chapter 3

Limited-Angle Computed Tomography for Sandwich Structures Using Data Fusion

Section 3.1 of this chapter defines the null space of the limited-angle Radon transform. The nature of the null space suggests that wide structures with correspondingly thin Fourier transforms will lie in the null space and be invisible in limited-angle raysum data. Face sheets in sandwich structures have this property. Therefore, accurate limited-angle CT reconstruction of sandwich structures is not possible because the face sheets are not properly measured and interpolation of the data is invalid.

The limitations of limited-angle CT present an opportunity to use constraint-based data fusion to improve reconstruction. Sections 3.2 and 3.3 introduce a novel strongly-coupled feed-forward fusion system for limited-angle CT. The system exploits the complementary nature of x-ray and ultrasound by using the latter to measure face sheet features that are otherwise invisible in raysum data. Fusing face sheet measurements into reconstruction removes the face sheets from the null space. Thus, the fusion system

produces an accurate reconstruction of a sandwich structure where it would otherwise be impossible.

Section 3.2 describes the data acquisition process for the three sources of data: range finder, ultrasound, and x-ray. Section 3.3 shows how the fusion system combines the multiple data sources, including the mathematical formulation of the applicable inverse problem.

3.1. Null Space of the Limited Angle Radon Transform

Define the limited-angle Radon transform of the cross section $\mu(x, y)$, $[\mathcal{R}_\Theta \mu](l, \theta)$, as:

$$[\mathcal{R}_\Theta \mu](l, \theta) = [\mathcal{R} \mu](l, \theta): -\Theta \leq \theta \leq +\Theta, 0 < \Theta < \pi/2. \quad (3-1)$$

The limited-angle transform is identical to the Radon transform except that it is defined over a limited range of angles only. By definition, a function μ_n is in the null space of the \mathcal{R}_Θ if $[\mathcal{R}_\Theta \mu_n](l, \theta) = 0$. Any structure in a specimen that lies in the null space does not contribute to the data measured by the transform. In this sense, such structures are invisible to data acquisition. It is important, therefore, to know the extent of the null space and which structures are contained within it. The remainder of this section defines the null space of the limited-angle Radon transform and an important class of structures that lie almost entirely within it.

As stated above, by definition $\mu_n \in \mathcal{N}(\mathcal{R}_\Theta)$ if and only if:

$$[\mathcal{R}_\Theta \mu_n](l, \theta) = 0. \quad (3-2)$$

Taking the Fourier transform of Equation (3-2) gives:

$$[\mathcal{F}_l \mathcal{R}_\Theta \mu_n](\omega_l, \theta) = 0. \quad (3-3)$$

By the Fourier slice theorem (see Chapter 2, Section 2.2.3):

$$[\mathcal{F}_l \mathcal{R}_\Theta \mu_n](\omega_l, \theta) = [\mathcal{F}_{x,y} \mu_n](\omega_l \cos \theta, -\omega_l \sin \theta), \quad (3-4)$$

and

$$[\mathcal{F}_{x,y} \mu_n](\omega_x, \omega_y) = 0, \quad -\Theta \leq \theta \leq \Theta, \quad \tan \theta = \frac{\omega_y}{\omega_x}. \quad (3-5)$$

The case where $\omega_x, \omega_y = 0$ requires special consideration. The limited-angle transform does sample this case, and, as for Equations (3-2) through (3-5):

$$[\mathcal{F}_{x,y}\mu_n](0,0) = 0. \quad (3-6)$$

Therefore the null-space is:

$$\mathcal{N}(\mathcal{R}_\Theta) = \left\{ \mu_n(x,y): \begin{array}{l} [\mathcal{F}_{x,y}\mu_n](\omega_x, \omega_y) = 0, -\Theta \leq \theta \leq \Theta, \\ \tan \theta = \frac{\omega_y}{\omega_x}, [\mathcal{F}_{x,y}\mu_n](0,0) = 0 \end{array} \right\}. \quad (3-7)$$

In words, the null space is the set of all functions whose Fourier transform is zero within the sampled regions of the Fourier domain.

Equation (3-7) illustrates the effect of limiting the range of angles. As Θ decreases the portion of a function (in Fourier space) that must be zero for the function to lie in the null space also decreases and the null space becomes larger. A practical consequence of a larger null space is that more of the structure of a specimen lies within it and is not sampled by the limited-angle Radon transform. Although CT reconstruction is ill-posed for both large and small null spaces, reconstruction is more difficult with the large null space because more structural information is missing.

Equation (3-7) also hints at what types of structures lie in the null space. As a rule, narrow structures have wide Fourier transforms and wide structures have narrow Fourier transforms. One expects, therefore, that a wide structure will lie in the null space, provided that it is oriented so that its narrow Fourier transform lies in the unsampled region. An example is the face sheets of sandwich structure. It is well known that the outer extremities of a structure carry stresses due to bending. Therefore, to achieve economy of weight and material, designers often create structures to carry loads in a thin outer shell separated by some core material, i.e., a sandwich structure. One example is the control surfaces of an airplane. The outer shell, i.e., the face sheets, carries the loads while the purpose of the core is to provide a rigid connection between the face sheets. If the structure is wide, as is the case for aircraft components, it does not allow acquisition

of full range raysum data, and furthermore, the face sheet structure lies almost entirely in the limited-angle Radon transform null space. We now show that this is the case.

The function $s_{r,\phi,a,d}(x,y)$, shown in Figure 3.1, represents a hypothetical face sheet and is given by:

$$s_{r,\phi,a,d}(x,y) = \begin{cases} a, & -\frac{d}{2} \leq (x \sin \phi + y \cos \phi) - r \leq \frac{d}{2}, \\ 0, & \text{otherwise} \end{cases} \quad (3-8)$$

i.e., a rectangular function in the y -axis direction with amplitude a and width d , shifted r from the origin in the y -axis direction and rotated by angle ϕ . The Fourier transform of $s_{r,\phi,a,d}(x,y)$, $[\mathcal{F}_{x,y} s_{r,\phi,a,d}](\omega_x, \omega_y)$, is:

$$[\mathcal{F}_{x,y} s_{r,\phi,a,d}](\omega_x, \omega_y) = \delta(\omega'_x) a d \operatorname{sinc}(\pi d \omega'_y) e^{-j\omega'_y r}, \quad (3-9)$$

where

$$\begin{bmatrix} \omega'_x \\ \omega'_y \end{bmatrix} = \begin{bmatrix} \cos \phi & -\sin \phi \\ \sin \phi & \cos \phi \end{bmatrix} \begin{bmatrix} \omega_x \\ \omega_y \end{bmatrix},$$

i.e., the Fourier transform is a sinc function along the y axis, phase shifted (for r) and rotated by angle ϕ .




Figure 3.1: Function $s_{r,\phi,a,d}(x,y)$; shaded region has value a and zero everywhere else.

Figure 3.2 shows a sandwich structure with two face sheets surrounding some core material. The width of the structure restricts the range of angles to $-\Theta \leq \theta \leq \Theta$

where $\Theta < \frac{\pi}{2} - \phi$. Thus, with the exception of the zero-frequency components, the face sheets lie within the null space of the limited-angle Radon transform. The presence of the face sheets confounds reconstruction. Their position in the null space eliminates the possibility of accurate reconstruction based on general *a priori* assumptions, so specific information about the face sheets is essential. Data fusion solves this problem by incorporating face sheet measurements gleaned from other data sources. The remainder of this chapter describes a novel data fusion system for limited-angle CT that acquires face sheet data and fuses them with raysum data in reconstruction.




Figure 3.2: Sandwich specimen with face sheets at angle ϕ . The width of the specimen prevents raysum acquisition outside the range $-\Theta \leq \theta \leq \Theta$. Note that $\Theta < (\pi/2) - \phi$, so the face sheets lie almost entirely within the limited-angle Radon transform null space.

3.2. Data Acquisition for Limited-Angle CT Data Fusion System

Clearly, limited-angle CT for sandwich structures based solely on raysum data is not feasible. The face sheets of the sandwich lie in the limited-angle Radon transform null space so that the assumptions necessary for interpolation are insufficient. This section proposes a novel system for limited-angle CT, starting with a description of the data acquisition. The system is based on raysum acquisition with modifications to allow

face sheet measurements from range finders and ultrasound. A description of a strongly-coupled feed-forward data fusion system that combines the data follows in the next section.

3.2.1 Raysum Data Acquisition

Figure 3.3 shows schematically the raysum data acquisition apparatus for thin objects. The apparatus sits on a large *C*-shaped manipulator (the *C*-arm) that allows the x-ray source (mounted on one end of the *C*) and the sensor (mounted on the other end of the *C*) to maneuver about a thin specimen.




Figure 3.3: Raysum data acquisition system.

To acquire an image, the *C*-arm sweeps the array across an area. A frame buffer assembles the image from the lines of data acquired by the array. Image acquisition is not important for CT, but a carefully arranged set of raysums is. Figure 3.4 shows, looking at the linear array end-on, the *C*-arm motion during data acquisition. The source-sensor array scans the specimen along a line as if acquiring an image, but only pixels from the centre of the array are used. The projection angle (the angle between vertical and the source-sensor plane) is initially zero. Subsequent scans acquire projection data for angles in the range $-\theta \leq \theta \leq \theta$.




Figure 3.4: Scan method for limited-angle raysum acquisition: (a) parallel with vertical and (b) at an angle to vertical.

3.2.2 Range Data

An important assumption often used in limited-angle CT is that of limited spatial support for the specimen, i.e., the specimen fits into a bounding region. Constraint-based data fusion improves solutions of inverse problems by substituting reliance on constraining assumptions with reliance on measured data. It is sensible, therefore, to measure the bounding shape of the specimen rather than rely on assumptions if possible.

Fortunately, devices exist that measure precisely the bounding shape of an object. An example is a laser range finder, an instrument that measures the distance from the instrument to an opaque surface in front of it. It projects a small dot of laser light onto the surface of the specimen. The optical system of the range finder detects the dot and triangulates its position. By sweeping the dot along a line and triangulating at each position, the range finder acquires range data along a line. Whereas the linear array x-ray detector collects raysum data along a line, a laser range finder collects range data along a line. To collect range data over an area, the range finder sweeps across the area in a manner analogous to the linear array x-ray detector when acquiring an image. One example of a commercially available laser range finder is the *Saturn-2000* built by *Servo-Robot* [35]. Two range finders mounted on the *C-arm* apparatus, one beside the x-ray

source and the other beside the detector, can measure the bounding region of the specimen with a single sweep of the *C*-arm. Figure 3.5 shows the modified apparatus for acquisition of raysum and range data.




Figure 3.5: Scanning method modified to include laser range finder for bounding region measurement.

Laser range finders are not the only method of determining the bounding region, but they are available and are accurate. An alternative is stereo vision or a modified stereo system. Two video cameras mounted in place of the range finders can acquire a sequence of images as the apparatus sweeps across the specimen. From these images it should be possible to compute range data, i.e., compute depth from motion where the motion is known. This idea is similar to stereo vision but whereas stereo uses only two cameras to compute depth, the sweeping motion of the *C*-arm provides many images.

3.2.3 Thickness Data

Measurement of the bounding region of the specimen by a range finder gives the position of the outer side of each face sheet. To completely measure the face sheet it is also necessary to know the thickness of the face sheet. Ultrasound is well suited for this measurement.

To acquire thickness measurements use a set of ultrasound pulse-echo traces taken along the face sheet in the scan direction. Ideally the traces have two reflections, one from the outer surface and one from the inner surface. From the position of the two

reflections calculate the time for the sound to travel between the two surfaces. Assuming that the speed of sound in the face sheet material is known compute the thickness of the face sheet from the speed and travel time.

Figure 3.6 shows the apparatus for ultrasound data acquisition. Conceptually, the ultrasound system sits on the C-arm with the range finders and x-ray system. Practical requirements for ultrasound (e.g. water squirters) and x-ray (high-voltage circuits) prohibit their coexistence. Therefore, the ultrasound apparatus is separate from the x-ray and range apparatus.




Figure 3.6: Apparatus for ultrasound data acquisition.

It is possible that thickness measurement may fail occasionally. This can occur for many reasons, e.g., scattering obstructions, defects within the face sheet, or internal structures which confound the reflections. Fortunately, the thickness of the face sheets is not rapidly varying and it is a straightforward process to interpolate missing thickness values.

3.3 The Fusion System

This section introduces a novel data fusion system for limited-angle CT of sandwich structures based on the apparatus shown in Section 3.2. Section 3.3.1 shows

the flow of data which is essentially a feed-forward strongly-coupled data fusion system, and Section 3.3.2 shows the mathematical formulation of the relevant inverse problem.

3.3.1 Data Flow

A strongly-coupled feed-forward data fusion system processes the data acquired by the apparatus of Section 3.2. Figure 3.7 shows a block diagram of the system.

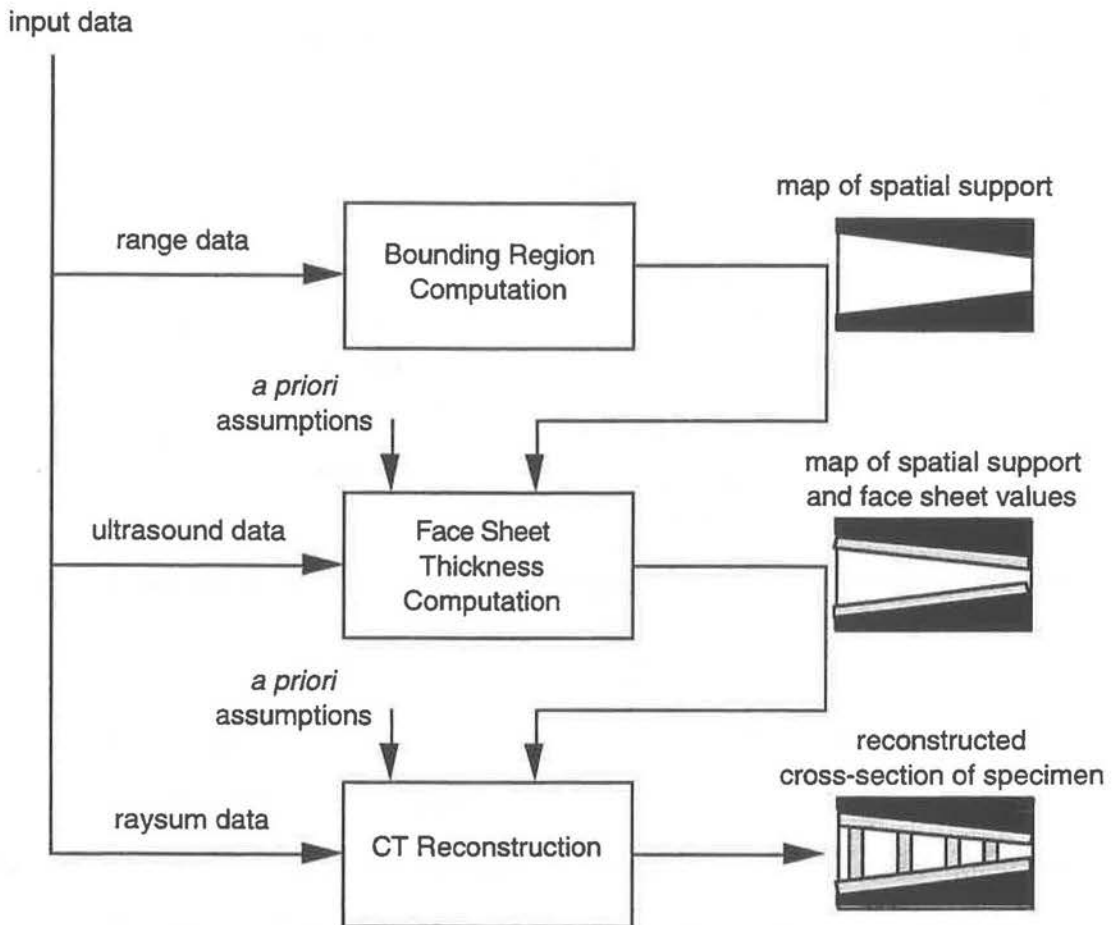


Figure 3.7: Block diagram of strongly-coupled feed-forward data fusion system for limited-angle CT of sandwich structures.

The first module of the system, bounding region computation, takes range data as input and produces a map of the spatial support for the specimen. The map segments the cross section into exterior regions and regions within the specimen.

The second module, face sheet thickness computation, takes ultrasound data as input and computes a map segmenting the specimen into exterior regions, face sheet

regions, and interior regions. The ultrasound data, constrained by the known speed of sound, give face sheet thickness, and spatial support constraints define the positions of the face sheets. Combined, the face sheet positions and face sheet thickness give a map of the face sheet regions. Regions that are not exterior nor in a face sheet must be, by default, in the specimen interior.

The map of the exterior and face sheet regions constrains the third module, CT reconstruction, whose input and output are the raysum data and the reconstructed cross-section respectively. As shown in Section 3.1, the face sheet structures are invisible to limited-angle raysums, and thus confound accurate reconstruction. The complementary nature of x-ray and ultrasound data allow the segmentation based on range and ultrasound to constrain the reconstruction to overcome this limitation. Whereas with the raysum data alone the reconstruction fills the unsampled Fourier space based on invalid *a priori* assumptions, the new constraints restrict the solutions to conform to the face sheets. These constraints do not necessarily make the problem well-posed so *a priori* assumptions are still necessary for reconstruction, but these assumptions no longer lead to erroneous results. Wide internal structures with edges parallel or near parallel to the face sheets are still a problem. In practice though, internal structures are much narrower than the face sheets so more of their Fourier transform lies in the sampled Fourier space, allowing better interpolation from *a priori* assumptions.

3.3.2. Mathematical Formulation

Acquisition of raysums is represented by the discrete linear system:

$$R\mathbf{x} = \mathbf{y}, \quad (3-10)$$

where \mathbf{x} is a vector of linear attenuation coefficients in a specimen and \mathbf{y} is a vector of raysums. R is a matrix representing the discrete Radon transform. Each raysum measurement is given by the inner product:

$$y_i = \langle \mathbf{r}_i, \mathbf{x} \rangle, \quad (3-11)$$

where y_i is the i^{th} raysum and \mathbf{r}_i is the i^{th} row of R . The j^{th} element of \mathbf{r}_i , $r_{i,j}$, is the length of the i^{th} ray that passes through the j^{th} element of \mathbf{x} . For limited-angle raysum acquisition, only rows corresponding to the raysums that are possible to measure are in R . The available raysums determine R and the number of singular values equal to zero. From Section 3.1 we know that R has a non-trivial null space containing the face sheet structures. Therefore, accurate CT reconstruction based on Equation (3-10) alone is ill-posed.

Range data and ultrasound data provide additional constraints so that reconstruction recovers the face sheet structure. The range finder gives a bounding region for the specimen. This leads to a map of regions inside and outside the specimen and the following system of equations:

$$W\mathbf{x} = W\mathbf{x}_F \quad (3-12)$$

where:

$$W = \text{diag}(w_1, w_2, \dots, w_j, \dots, w_N) \quad ,$$

$$w_j = \begin{cases} 1 & x_j \text{ outside specimen} \\ 0 & x_j \text{ inside specimen} \end{cases} \quad , \text{ and}$$

$$x_{F_j} = \begin{cases} \mu_{\text{air}} & x_j \text{ outside specimen} \\ 0 & x_j \text{ inside specimen} \end{cases}$$

The subscript F denotes fusion. μ_{air} is the linear attenuation coefficient for air and is zero. \mathbf{x}_F is a partial reference image based on the fusion data. The range finder data inversion gives the elements on the diagonal of W . Intermediate values between zero and one on the diagonal indicate that a particular element lies across a boundary between regions.

The ultrasound module gives the thickness of the face sheets. Face sheet thickness and the location of the outer side of each face sheet determine which elements of \mathbf{x} are exterior, which are in a face sheet, and which are interior. This gives a new matrix W and partial reference image \mathbf{x}_F where:

$$w_j = \begin{cases} 1 & x_j \text{ in face sheet or exterior} \\ 0 & x_j \text{ in specimen interior} \end{cases}, \text{ and}$$

$$x_{F_j} = \begin{cases} \mu_{\text{air}} & x_j \text{ outside specimen} \\ \mu_{\text{face sheet}} & x_j \text{ in face sheet} \\ 0 & x_j \text{ in specimen interior} \end{cases}.$$

$\mu_{\text{face sheet}}$ is the linear attenuation coefficient of the face sheet material.

Equations (3-10) and (3-12) combine to form the system:

$$\begin{bmatrix} R \\ W \end{bmatrix} \mathbf{x} = \begin{bmatrix} \mathbf{y} \\ W\mathbf{x}_F \end{bmatrix}. \quad (3-13)$$

Solving for \mathbf{x} from Equation (3-13) yields a solution that is constrained by Equation (3-12) and therefore constrained by the range and ultrasound data. Although Equation (3-12) further constrains the problem, the reconstruction is still ill-posed. The matrix $\begin{bmatrix} R \\ W \end{bmatrix}$ is not necessarily square and is not full column rank. Equation (3-13) eliminates the face sheet structures from the null space of $\begin{bmatrix} R \\ W \end{bmatrix}$, allowing a reconstruction that properly accounts for the face sheets. Interior structures with edges parallel or nearly parallel to the face sheets are still at least partly in the null space.

It is possible to weight the fusion data in Equation (3-13) to a greater or lesser degree by using other values on the diagonal of W , where a large weight indicates a greater degree of confidence in the constraint. For example, one can choose the weights to be the reciprocals of the standard deviations of constraint measurements [39]. The unitary weights are arbitrary here, but they happen to produce good results using singular value decomposition and regularization (see next chapter). Weights in W are irrelevant to the method of projection onto convex sets (also in the next chapter) where the parameters that control constraint set size determine the relative weighting of constraints.

In the context of the data fusion system proposed here, strongly-coupled feed-forward data fusion consists of:

1. forming the system of equations, and

2. solving for \mathbf{x} from Equation (3-13).

It remains to specify a method of solution. Chapter 4 presents and compares three possible methods.

3.4 Chapter Summary

The null space of the limited-angle Radon transform is the set of all functions whose Fourier transform is zero over the region sampled by the transform. This null space specifically includes the face sheets of sandwich structures, with the exception of the zero-frequency component. Therefore, an important part of a sandwich structure is invisible to limited-angle raysums and accurate reconstruction is not possible from raysum data alone.

A novel limited-angle CT system, proposed in this chapter, uses constraint-based data fusion to solve this problem. The system specifies data acquisition methods for range data, ultrasound data, and x-ray raysum data. A strongly-coupled feed-forward data fusion system combines the data from the different sources.

The range data define a bounding region for the specimen. Outside that region linear attenuation must be zero. Range data also locate one side of each face sheet. Ultrasound data give thickness measurements over the face sheet. Thickness measurements constrained by the face sheet locations determine which regions of the specimen are face sheet. Face sheet regions have a known linear attenuation. A part of the specimen that is neither inside the bounding region nor in the face sheets is, by default, in the interior of the specimen. No assumptions are made about linear attenuation in the interior.

The segmentation of the specimen into exterior, face sheet, and interior regions gives a set of constraints expressed by the linear system:

$$W\mathbf{x} = W\mathbf{x}_f$$

where W is a diagonal matrix whose elements indicate whether or not an element of \mathbf{x} is in a region where linear attenuation is known. \mathbf{x}_F is a partial reference image based on the fusion data.

Raysum acquisition is modeled by the linear system:

$$R\mathbf{x} = \mathbf{y}.$$

Range data and ultrasound data constrain the reconstruction from raysums in the combined linear system:

$$\begin{bmatrix} R \\ W \end{bmatrix} \mathbf{x} = \begin{bmatrix} \mathbf{y} \\ W\mathbf{x}_F \end{bmatrix}.$$

Data fusion in this context consists of building the above system of equations and solving for \mathbf{x} . The matrix $\begin{bmatrix} R \\ W \end{bmatrix}$ is not square, nor is it full column rank, so the inversion is still ill-posed. However, the new system reduces the size of the null space so that it no longer contains the face sheets. Only internal structures with edges parallel or nearly parallel to the face sheets are in the null space. In practice, internal structures are narrow enough that limited-angle data should recover them adequately.

The next chapter shows three numerical methods suitable for solving for \mathbf{x} from Equation (3-13).

Chapter 4

Numerical Methods for Limited-Angle Tomography System

Chapter 3 introduced a novel method for limited-angle computed tomography that exploits data fusion to properly reconstruct a cross section of a specimen with face sheets. The method ultimately requires solution of the linear system of equations:

$$\mathbf{Ax} = \mathbf{b}$$

where:

$$A = \begin{bmatrix} R \\ W \end{bmatrix} \text{ and } \mathbf{b} = \begin{bmatrix} \mathbf{y} \\ W\mathbf{x}_f \end{bmatrix}.$$

The nature of A dictates that solving for \mathbf{x} is an ill-posed problem; if a solution exists, it is not unique.

A plethora of methods exists for solving linear systems. This chapter reviews the following three methods suitable for the limited-angle CT fusion problem:

1. singular value decomposition (SVD),
2. regularization and the conjugate gradient method (R/CG), and
3. projection onto convex sets (POCS).

Note that these are not the only methods possible.

A brief description of each method follows, including details of application to the problem of interest, and its advantages and disadvantages. Section 4.4 summarizes the methods in tabular form. Results of application of each method to synthetic data, and application of R/CG and POCS to real data are in Chapter 5.

4.1 Singular Value Decomposition

Strang [38] suggests that singular value decomposition (SVD) "... is not nearly as famous as it should be." Whether or not this is true, SVD is an excellent tool for analysis of linear systems. This presentation of SVD, based on Golub and Vanloan [13], Stoer and Bulirshch [37], Forsythe, Malcolm and Moler [12], Strang [38], and Press *et al.* [31], focuses on the decomposition itself rather than the algorithm that computes it.

4.1.1 Properties of the Decomposition

Start with the generic linear system:

$$Ax = b. \quad (4-1)$$

We know nothing about A except that it is $m \times n$ and $m \geq n$. For $m < n$ fill A with rows of zeros to get $m = n$. SVD decomposes the matrix A into the product of three matrices:

$$A = U\Sigma V^T. \quad (4-2)$$

The following tableau elucidates the decomposition:

$$A = \left[\begin{array}{c|c|c} U & \Sigma & V^T \\ \hline m \times m & m \times n & n \times n \end{array} \right]. \quad (4-3)$$

U and V are both orthogonal, but the significance of SVD lies in the structure of Σ :

$$\Sigma = \begin{bmatrix} \sigma_1 & & & \mathbf{0} \\ & \sigma_2 & & \\ & & \ddots & \\ \mathbf{0} & & & \sigma_n \\ & & & & \mathbf{0} \end{bmatrix}. \quad (4-4)$$

The values σ_j relate to the eigenvalues of $A^T A$; $\sigma_j = \sqrt{\lambda_j}$, where λ_j is the j^{th} eigenvalue of $A^T A$. Since $A^T A$ is positive semi-definite, $\lambda_j \geq 0$ and $\sigma_j \geq 0$. Because the last $m - n$ rows of Σ are zeros, the last $m - n$ columns of U are unnecessary.

Ignoring these rows and columns gives the abbreviated decomposition:

$$A = \left[\begin{array}{c} U \\ m \times n \end{array} \right] \left[\begin{array}{c} \Sigma \\ n \times n \end{array} \right] \left[\begin{array}{c} V^T \\ n \times n \end{array} \right], \quad (4-5)$$

where:

$$\Sigma = \begin{bmatrix} \sigma_1 & & & \mathbf{0} \\ & \sigma_2 & & \\ & & \ddots & \\ \mathbf{0} & & & \sigma_n \end{bmatrix}. \quad (4-6)$$

Rearranging Equation (4-2) gives:

$$AV = U\Sigma, \quad (4-7)$$

from $V^T V = I$. Individual columns in Equation (4-7) are:

$$A\mathbf{v}_j = \sigma_j \mathbf{u}_j. \quad (4-8)$$

Consider Equation (4-8) for each of two categories for σ_j :

$$1) \quad \sigma_j = 0$$

$$A\mathbf{v}_j = \mathbf{0}, \text{ so}$$

$$\mathbf{v}_j \in \mathcal{N}(A)$$

$$2) \quad \sigma_j > 0$$

$$A \frac{1}{\sigma_j} \mathbf{v}_j = \mathbf{u}_j, \text{ so}$$

$$\mathbf{u}_j \in \text{Range}(A).$$

The following items summarize important aspects of the decomposition:

1. If $\sigma_j = 0$ for any j then the matrix $A^T A$ is singular and the linear system of Equation (4-1) is under determined,

2. The null space of A , $\mathcal{N}(A)$, and its orthogonal complement, $\mathcal{N}^\perp(A)$, are given by:

$$\mathcal{N}(A) = \text{span}(\mathbf{v}_j : \sigma_j = 0), \text{ and}$$

$$\mathcal{N}^\perp(A) = \text{span}(\mathbf{v}_j : \sigma_j \neq 0),$$

and

3. The range of A , $\text{Range}(A)$, and its orthogonal complement, $\text{Range}^\perp(A)$ are given by:

$$\text{Range}(A) = \text{span}(\mathbf{u}_j : \sigma_j \neq 0), \text{ and}$$

$$\text{Range}^\perp(A) = \text{span}(\mathbf{u}_j : \sigma_j = 0).$$

SVD indicates whether or not a system of equations is under-determined, and it computes orthogonal bases for the null space and range of a matrix and their orthogonal complements.

4.1.2 Least-Squares Minimum-Norm Solution

In addition to finding the null space and range of a matrix, SVD leads directly to the least-squares minimum-norm solution of an under-determined linear system. The least-squares minimum-norm solution for \mathbf{x} , $\hat{\mathbf{x}}$ is:

$$\hat{\mathbf{x}} = V\Sigma^+U^T\mathbf{b}, \quad (4-9)$$

i.e., $\hat{\mathbf{x}}$ is the shortest \mathbf{x} that minimizes $\|\mathbf{Ax} - \mathbf{b}\|^2$ and is unique. Equation (4-9) uses

$V^TV = I$ and $U^TU = I$, and the pseudo-inverse Σ^+ given by:

$$\Sigma^+ = \begin{bmatrix} \sigma_1^+ & & & \mathbf{0} \\ & \sigma_2^+ & & \\ & & \ddots & \\ \mathbf{0} & & & \sigma_n^+ \end{bmatrix}, \quad (4-10)$$

where:

$$\sigma_j^+ = \begin{cases} 0, & \sigma_j = 0 \\ \frac{1}{\sigma_j}, & \text{otherwise} \end{cases}. \quad (4-11)$$

Often in practice $\sigma_j \neq 0$, but σ_j is very small. The condition number of A , $\text{cond}(A)$,

where:

$$\text{cond}(A) = \frac{\max(\sigma_j)}{\min_{\sigma_j \neq 0}(\sigma_j)}. \quad (4-12)$$

describes this situation. If $\min_{\sigma_j \neq 0}(\sigma_j)$ is small then $\text{cond}(A)$ is large and the matrix is ill-conditioned. Either the machine precision or the accuracy of the data determine the maximum tolerable condition number. Let ε be the limiting precision, then the condition number should not exceed $1/\varepsilon$, i.e.:

$$\text{cond}(A) = \frac{\max(\sigma_j)}{\min_{\sigma_j \neq 0}(\sigma_j)} \leq 1/\varepsilon \quad (4-13)$$

yields a well-conditioned problem.

SVD provides the opportunity to condition the problem by throwing out any σ_j that is too small and replacing it with a zero. Consequently, the null space increases to include precisely the part of the problem that was ill-conditioned at the cost of a loss of range. Redefinition of the pseudo inverse to condition the matrix gives:

$$\sigma_j^+ = \begin{cases} 0, & \sigma_j < \varepsilon \max(\sigma_j) \\ \frac{1}{\sigma_j}, & \text{otherwise} \end{cases} \quad (4-14)$$

4.1.3 Application to Limited-Angle CT

Computation of the tomographic reconstructions proposed in Chapter 3 using SVD follows these steps:

1. build the matrix $A = \begin{bmatrix} R \\ W \end{bmatrix}$.
2. use a canned SVD algorithm, e.g. `svdcmp()` from [31], to compute U , V , and the diagonal of Σ .

3. set

$$\sigma_j = \begin{cases} 0, & \sigma_j < \varepsilon \max(\sigma_j) \\ \sigma_j, & \text{otherwise} \end{cases}$$

for $j = 1, 2, \dots, n$.

4. use a canned SVD back-substitution algorithm, e.g. `svbksb()` from [31], to compute the least-squares minimum-norm solution, $\hat{\mathbf{x}}$, from U , V and Σ , and data $\mathbf{b} = \begin{bmatrix} \mathbf{y} \\ W\mathbf{x}_F \end{bmatrix}$.

Repeat step 4 as needed for different data provided A does not change.

SVD has $O(mn^2 + n^3)$ complexity in time. Current computer technology is too slow to tackle CT problems of a practical size with SVD and third-order complexity staves off the day when it will. Storage requirements are also exorbitant. In the experiments reported in Chapter 5 the image size is 72×200 giving $n = 14,400$. The matrix A must contain at least $n^2 = 207,360,000$ elements. Using a 32-bit floating point representation for each element, A occupies about 790 Mbytes of memory. SVD requires two arrays of this size. Arrays of this size are not practical using current technology and this is only a moderately sized problem. Images of $n = 200 \times 200 = 40,000$ are routine in CT.

It may be possible to capitalize on sparsity in A and reduce both computation time and storage. Unfortunately, sparsity in A does not translate directly to sparsity in U and V . Currently no efficient methods exist for SVD of sparse systems. Should such techniques come to fruition, SVD has the advantage that once a decomposition is complete, solutions are easily computed for different data vectors \mathbf{b} , i.e., do the decomposition once to compute U , V and Σ , taking a long time if necessary, and repeatedly solve for any number of different \mathbf{b} 's quickly.

To summarize, SVD has the following advantages:

1. definition of range and null space,
2. controlled conditioning of the problem, and
3. well-posedness due to unique least-squares minimum-norm,

and the following disadvantages:

1. third-order computational complexity, and
2. exorbitant storage requirements.

If one is willing to sacrifice resolution of reconstruction SVD can perform CT reconstruction. The resolution will be too low for practical applications, but the decomposition gives a useful analysis of the problem.

4.2 Regularization and Conjugate Gradient Method

4.2.1 Regularization

As shown in the previous section, SVD gives a least-squares solution to

$$A\mathbf{x} = \mathbf{b},$$

i.e., it minimizes the quadratic objective function $E = \|A\mathbf{x} - \mathbf{b}\|^2$. The minimum is not unique because A has a non-trivial null space, so SVD finds the shortest \mathbf{x} that minimizes $E = \|A\mathbf{x} - \mathbf{b}\|^2$ (least-squares minimum-norm solution).

An alternative is to modify the system so that the least-squares solution is unique without the minimum-norm constraint. Tikhonov and Arsenin [41] describe one such technique, called regularization, to solve singular and ill-conditioned systems of linear algebraic equations. This section examines regularization in the context of the proposed limited-angle CT system.

Begin with the linear system:

$$A\mathbf{x} = \mathbf{b}. \quad (4-15)$$

To have a unique least squares solution it is necessary for A to be full column rank.

When A is not full column rank, regularization modifies A by adding more rows to give:

$$\begin{bmatrix} A \\ \alpha\Omega \end{bmatrix} \mathbf{x} = \begin{bmatrix} \mathbf{b} \\ \alpha\Omega\mathbf{x}_0 \end{bmatrix}, \quad (4-16)$$

where $\alpha > 0$, and Ω is some matrix. If Ω is full column rank, then the matrix $\begin{bmatrix} A \\ \alpha\Omega \end{bmatrix}$ is

also full column rank and the least squares problem of Equation (4-16) is well-posed. If Ω is nearly full column rank, then it may be sufficient to make $\begin{bmatrix} A \\ \alpha\Omega \end{bmatrix}$ full column rank.

To see that this is indeed the regularization proposed by Tikhonov and Arsenin, look at

the quadratic objective function minimized in the least-squares solution of Equation (4-16):

$$E = \left\| \begin{bmatrix} A \\ \alpha \Omega \end{bmatrix} \mathbf{x} - \begin{bmatrix} \mathbf{b} \\ \alpha \Omega \mathbf{x}_0 \end{bmatrix} \right\|^2, \quad (4-17)$$

which expands to

$$E = \|A\mathbf{x} - \mathbf{b}\|^2 + \alpha^2 \|\Omega(\mathbf{x} - \mathbf{x}_0)\|^2. \quad (4-18)$$

Compare Equation (4-18) to Tikhonov and Arsenin's functionals:

$$M^\alpha[z, \tilde{u}, A] = \|Az - \tilde{u}\|^2 + \alpha \|z - z^1\|^2, \quad \alpha > 0, \quad (4-19)$$

and

$$M^\alpha[z, \tilde{u}, A] = \|Az - \tilde{u}\|^2 + \alpha \Omega[z], \quad \alpha > 0. \quad (4-20)$$

The addition of the regularizing functional (the last terms in Equations (4-19) and (4-20)) is equivalent to adding more rows to Equation (4-15)*.

Now consider the elements of the regularizing functional, α , Ω and \mathbf{x}_0 . If $\alpha = 0$ Equation (4-16) reduces to the original singular system of equations, Equation (4-15). On the other hand, if $\alpha \rightarrow \infty$ Equation (4-16) reduces to:

$$\Omega \mathbf{x} = \Omega \mathbf{x}_0,$$

which is not the problem of interest. So α varies through a continuum. At one end a small α gives a problem that is close to the original but is ill-conditioned, while at the other end, a large α gives a different problem but one that is well-conditioned**.

Varying α makes a tradeoff between conditioning and fidelity of solution to the data.

We now examine two options for selecting Ω . The first option is $\Omega = I$, the identity matrix, which is full column rank, so $\begin{bmatrix} A \\ \alpha I \end{bmatrix}$ is also full column rank. This leads to

the following variations of the problem:

* Tikhonov and Arsenin use α where I use α^2 .

** In this case ill-posedness and ill-conditioning are different degrees of the same property. $\alpha = 0$ gives an ill-posed problem, whereas $\alpha > 0$ but very small gives a well-posed problem that is ill-conditioned.

- 1) If $\mathbf{x}_0 = \mathbf{0}$ then $\|\Omega(\mathbf{x} - \mathbf{x}_0)\|^2 = \|\mathbf{x}\|^2$. Regularizing in this case gives a well posed problem by biasing the solution towards the origin.
- 2) If $\mathbf{x}_0 = \bar{\mathbf{x}}$, where $\bar{\mathbf{x}}$ is a vector of \bar{x} 's and \bar{x} is the average value of \mathbf{x} , then $\|\Omega(\mathbf{x} - \mathbf{x}_0)\|^2 = \|\mathbf{x} - \bar{\mathbf{x}}\|^2$. Regularizing in this case gives a well posed problem by biasing toward the mean. This approach appears in the CT literature [1] [21] [22].

The second option is $\Omega = \begin{bmatrix} D_1 \\ D_2 \end{bmatrix}$, where D_1 is a discrete partial derivative in the x -axis direction in the reconstruction and D_2 is a discrete partial derivative in the y -axis direction. In this case we consider only $\mathbf{x}_0 = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}$. The matrix $\begin{bmatrix} D_1 \\ D_2 \end{bmatrix}$ is not full column rank but is nearly full column rank. To see this consider the corresponding system of partial differential equations:

$$\begin{aligned} \frac{\partial f(x, y)}{\partial x} &= 0 \\ \frac{\partial f(x, y)}{\partial y} &= 0 \end{aligned}$$

for which the solution is:

$$f(x, y) = c,$$

where c is a constant. There is only one degree of freedom in the solution, that is the selection of the constant. If A can resolve the single degree of freedom in $\begin{bmatrix} D_1 \\ D_2 \end{bmatrix}$ then

$\begin{bmatrix} A \\ D_1 \\ D_2 \end{bmatrix}$ is full column rank and least squares solution of Equation (4-16) is well-posed.

Note that in this case, if α is too large then Equation (4-16) reduces to $\begin{bmatrix} D_1 \\ D_2 \end{bmatrix} \mathbf{x} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}$ and

$A\mathbf{x} = \mathbf{b}$ diminishes from the problem. The single degree of freedom in the solution is not resolved and the least-squares problem is ill-conditioned.

The first option, $\Omega = I$, biases the solution towards \mathbf{x}_0 . Therefore, when $\mathbf{x}_0 = \mathbf{0}$ the solution is drawn towards a vector of zeros. We know *a priori* that such a bias is not

warranted. Likewise, when $\mathbf{x}_0 = \bar{\mathbf{x}}$ the solution is biased towards the mean. This bias, although more sensible than the previous one, is still not warranted by knowledge of the specimen. Values of \mathbf{x} are known to be of certain values only and the bias draws solutions to some value known not to exist in the specimen (except by chance). Although it is possible to use the fusion data to avoid some of this problem by making $\Omega = I - W$, the problem still exists for the interior regions of the reconstruction.

The second option, $\Omega = \begin{bmatrix} D_1 \\ D_2 \end{bmatrix}$ and $\mathbf{x}_0 = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}$, biases the solution to smooth images.

Although we know that the cross section is not perfectly smooth, it is smooth over most local regions, i.e., the specimen is piece-wise constant. Although there is a penalty incurred by smoothing local regions containing discontinuities, this option is more palatable than the first. Experiments reported in Chapter 5 use the second option for the limited-angle CT problem.

4.2.2 Parameter Selection

Selecting a value for the parameter α is an important consideration. If α is too small then the problem is ill-conditioned, but if α is too large then the solution deviates greatly from the data. In general, we want to have α as small as possible while still having a well-conditioned problem. Experiments presented in Chapter 5 use a trial-and-error approach to select α . Since the conjugate gradient method solves the problem quickly (see Section 4.2.3), it is reasonable to compute solutions for a range of α . Then simply use the smallest value of α for which CG converges to a reasonable solution. A reasonable solution is not so smooth that important detail is invisible, but does not have confounding artifacts due to ill-conditioning. Admittedly, the process is very subjective.

Cross validation [13] [14] is an alternative technique for selection of α not explored in this thesis. If, for example, $\Omega = I$ and $\mathbf{x}_0 = \mathbf{0}$ then cross validation selects α^2 by minimizing the cross-validation weighted square error [13]:

$$C(\alpha^2) = \frac{1}{m} \sum_{k=1}^m w_k [\mathbf{a}_k \mathbf{x}_k(\alpha^2) - \mathbf{b}_k]^2.$$

Paraphrasing Golub and Van Loan [13], minimization of $C(\alpha^2)$ is tantamount to choosing α^2 such that the final solution does not overly depend on any single measurement. Intuitively, cross validation minimizes the sensitivity to the data. Cross validation is possible for other Ω s and \mathbf{x}_0 s; see Bates and Wahba [3] for example. Although the trial-and-error method is satisfactory for the work presented in this thesis, cross validation may warrant future consideration should selection of the regularization parameter become critical, bearing in mind that cross validation has relatively high computational costs [3].

4.2.3 The Conjugate Gradient Method

Having regularized the problem, the resulting symmetric positive definite system requires solution. The conjugate gradient method (CG), recommended by Artzy, Elfving, and Herman [1] for fast convergence in CT, is one of many possible methods for solving the system. This section presents a summary of CG. For more details see Axelsson and Barker [2], Press *et al.* [31], or Pierre [29].

The following steps constitute the CG algorithm for finding the least-squares solution of $A\mathbf{x} = \mathbf{b}$, where A is positive-definite [1] [31]:

- (1) $\mathbf{g}_0 = \mathbf{h}_0 = \mathbf{b} - A\mathbf{x}_0$
- (2) $\mathbf{s}_i = A\mathbf{h}_i$

$$\lambda_i = \frac{\langle \mathbf{g}_i, \mathbf{h}_i \rangle}{\langle \mathbf{h}_i, \mathbf{s}_i \rangle}$$
- (3) $\mathbf{x}_{i+1} = \mathbf{x}_i + \lambda_i \mathbf{h}_i$
- (4) (direct method) $\mathbf{g}_{i+1} = \mathbf{b} - A\mathbf{x}_{i+1}$
 (indirect method) $\mathbf{g}_{i+1} = \mathbf{g}_i - \lambda_i \mathbf{s}_i$
- (5) $\gamma_i = \frac{\langle \mathbf{g}_{i+1}, \mathbf{g}_{i+1} \rangle}{\langle \mathbf{g}_i, \mathbf{g}_i \rangle}$ or $\gamma_i = \frac{\langle (\mathbf{g}_{i+1} - \mathbf{g}_i), \mathbf{g}_{i+1} \rangle}{\langle \mathbf{g}_i, \mathbf{g}_i \rangle}$
- (6) $\mathbf{h}_{i+1} = \mathbf{g}_{i+1} + \gamma_i \mathbf{h}_i$
- (7) go to step (2) for next iteration.

Note that in the context of CG, A is not $\begin{bmatrix} R \\ W \end{bmatrix}$, instead $A = R^T R + W^T W + \alpha^2 \Omega^T \Omega$ (see Section 4.2.4). Steps (1) and (4) initialize and compute the sequence $\mathbf{g}_0, \mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_{n-1}$, where \mathbf{g}_i is the negated gradient of E at \mathbf{x}_i , i.e., $\mathbf{g}_i = -\nabla E(\mathbf{x}_i)$. With initial conditions defined in step (1), steps (5) and (6) produce a sequence of mutually A -conjugate vectors $\mathbf{h}_0, \mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_{n-1}$, i.e.:

$$\mathbf{h}_i^T A \mathbf{h}_j = 0, \quad i \neq j,$$

where the \mathbf{h} 's are non-trivial. Because the \mathbf{h} 's are mutually A -conjugate they are necessarily linearly independent. Therefore, to minimize E it is sufficient to do n separate minimizations along the independent directions $\mathbf{h}_0, \mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_{n-1}$ (steps (2) and (3)).

Each iteration requires multiplication by the matrix A , which has $O(n^2)$ computational complexity. The complete minimization requires n iterations for a net complexity of $O(n^3)$. Two factors make $O(n^3)$ pessimistic though. First, often A is sparse and multiplication by A can be faster than $O(n^2)$. Second, requiring n iterations is a worst case scenario. Axelsson and Barker [2] give the more precise bound:

$$p(\varepsilon) \leq \frac{1}{2} \sqrt{\kappa(A)} \ln\left(\frac{2}{\varepsilon}\right) + 1,$$

where $p(\varepsilon)$ is the number of iterations required to converge to precision ε . $\kappa(A)$ is the spectral condition number of A defined as:

$$\kappa(A) = \frac{|\lambda_{\max}|}{|\lambda_{\min}|},$$

where λ_{\max} and λ_{\min} are the maximum and minimum eigenvalues of A . Clustered eigenvalues lead to further reductions in the number of iterations. We have no measure of the eigenvalues for A , but Davison [10] and Louis [23] show that singular values for R are clustered, which leads to clustered eigenvalues for $R^T R$. Results in Chapter 5 indicate that convergence is much better than the worst case, so it is likely that the

limited-angle CT problem here benefits from both the condition number and the clustering of eigenvalues.

4.2.4 Application to Limited-Angle CT

Application of regularization to the limited-angle CT problem (with data fusion) results in the modified system of equations:

$$\begin{bmatrix} A \\ D_1 \\ D_2 \end{bmatrix} \mathbf{x} = \begin{bmatrix} \mathbf{b} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \quad (4-21)$$

or more precisely:

$$\begin{bmatrix} R \\ W \\ D_1 \\ D_2 \end{bmatrix} \mathbf{x} = \begin{bmatrix} \mathbf{y} \\ W\mathbf{x}_F \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix}. \quad (4-22)$$

Minimizing the objective function:

$$E = \|R\mathbf{x} - \mathbf{y}\|^2 + \|W(\mathbf{x} - \mathbf{x}_F)\|^2 + \alpha^2 \|D_1\mathbf{x}\|^2 + \alpha^2 \|D_2\mathbf{x}\|^2, \quad (4-23)$$

or equivalently solving the following symmetric positive-definite linear system:

$$(R^T R + W^T W + \alpha^2 D_1^T D_1 + \alpha^2 D_2^T D_2) \mathbf{x} = R^T \mathbf{y} + W^T \mathbf{x}_F. \quad (4-24)$$

finds the least squares solution to Equation (4-22).

To solve the limited-angle data fusion problem by regularization:

1. Implement algorithms to compute the forward raysum operator, R , and the back projection operator, R^T .
2. Implement algorithms to compute $W = W^T$ and $D_1^T D_1 + D_2^T D_2$ (a discrete Laplacian operator).
3. Compute the vector $R^T \mathbf{y} + W^T \mathbf{x}_F$.
4. Use CG to solve Equation (4-24).

In summary, the advantages of regularization and CG are that it is fast for limited-angle CT, and it gives a well-posed and well-conditioned problem. The disadvantage is that regularization makes a tradeoff between fidelity of solution to the data and

conditioning of the problem. If one is willing to concede fidelity to achieve a well-posed and well-conditioned problem, regularization and CG offer a practical approach to limited-angle CT.

4.3 Projection onto Convex Sets

Projection onto convex sets (POCS) is a flexible method, able to incorporate a variety of constraints. As such it is particularly useful for constraint-based data fusion. The following discussion of POCS is based on Brègman [7], Youla and Webb [44], Gubin *et al.* [17], and Bauschke and Borwein [4]. POCS is used extensively for CT and subsumes the method known as ART (algebraic reconstruction technique) [8] [22] [42]. See Censor and Herman [8] for a review of projection methods, or Bauschke and Borwein [4] for a detailed survey.

4.3.1 General Description of POCS

POCS assumes a set of m constraints on the domain of A , all of the form:

$$\mathbf{x} \in C_i, \quad (4-25)$$

for $1 \leq i \leq m$, where C_i is a convex set. Combined, the constraints restrict the solution to lie in the convex set C_0 , i.e.:

$$\mathbf{x} \in C_0 = \bigcap_{i=1}^m C_i. \quad (4-26)$$

POCS makes the important assumption that C_0 is non-empty. Each convex set C_i has a corresponding projection operator P_i of the form:

$$P_i \mathbf{z} = \begin{cases} \text{point in } C_i \text{ closest to } \mathbf{z}, & \mathbf{z} \notin C_i \\ \mathbf{z}, & \text{otherwise} \end{cases} \quad (4-27)$$

Application of P_i to some point \mathbf{z} finds the point in C_i closest to \mathbf{z} . Points within C_i are fixed points of P_i . Let $T_i = 1 + \lambda_i(P_i - 1)$ be the relaxed operator corresponding to P_i , with relaxation coefficient, λ_i , in the range $0 \leq \lambda_i \leq 2$. The relaxed projection operators (and the projection operators for $\lambda = 1$) are non-expansive mappings, i.e.:

$$\|T_i \mathbf{x} - T_i \mathbf{y}\| \leq \|\mathbf{x} - \mathbf{y}\|.$$

The non-expansive property leads to the convergence properties of POCS.

Consider the parallel combination of a set of non-expansive mappings:

$$T_{\text{par}} = \sum_{i=1}^m \gamma_i T_i,$$

where the γ_i are weights, $0 \leq \gamma_i \leq 1$ and $\sum \gamma_i = 1$. The sequence $\{\mathbf{x}_n\} = \{T_{\text{par}}^n \mathbf{x}_0\}$ converges weakly to a point in C_0 . Furthermore, if \mathbf{x} is in Euclidean space then the convergence is strong. Note that if the γ_i vary between iterations, then T_{par} becomes the sequential combination of the T_i :

$$T_{\text{seq}} = T_m T_{m-1} \dots T_1,$$

by appropriate selection of the γ_i . This property gives flexibility, guaranteeing that the method will converge for both parallel and sequential variations, as well as for any other allowed values of γ_i . For arbitrary convex constraint sets there is no guarantee of rate of convergence. However, for constraints such as hyperslabs, or when the sets overlap (see Youla and Webb [44], Theorem 3) then convergence is at a linear rate*, although there is no way to know the value for β . For the purposes of this research we set $\lambda_i = 1$ always, and select the γ_i to give sequential application of the projection operators. Therefore, we compute the sequence $\{\mathbf{x}_n\} = \{P^n \mathbf{x}_0\}$ where $P = P_m P_{m-1} \dots P_1$.

Figure 4.1 illustrates POCS convergence. Two constraints require that the solution lie in the intersection of two half planes. Arrows in Figure 4.1 trace the progress of $\{P^n \mathbf{x}_0\}$ towards \mathbf{x} . Convergence is strong, because \mathbf{x} is in two-dimensional Euclidean space, and linear, because of the constraint set overlap.

* By linear convergence we mean $\|\mathbf{x}_n - \mathbf{x}\| \leq \alpha \beta^n$, where $0 \leq \beta < 1$ and $\alpha \geq 0$. In some references this is called geometric convergence.




Figure 4.1: Example of POCS convergence.

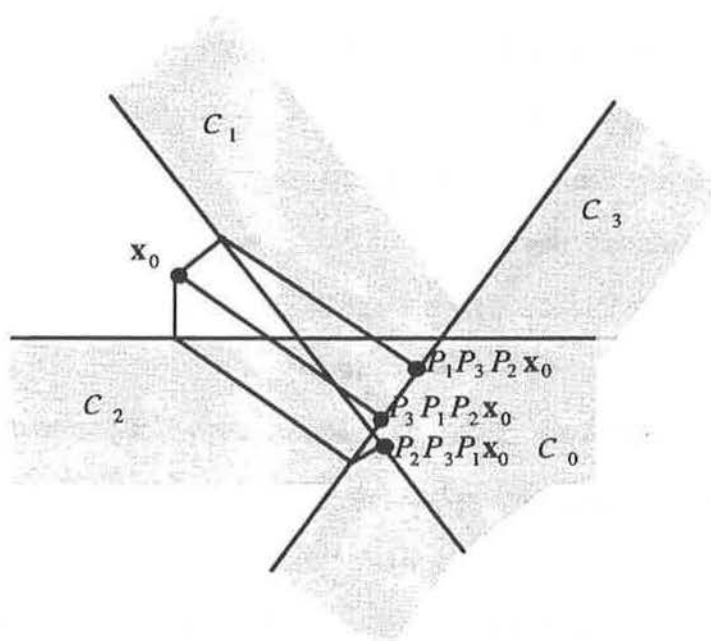


Figure 4.2: Demonstration of the variability of POCS convergence. POCS converges to three distinct but correct solutions depending on the order of projection operators.

It is tempting to relate the POCS solution to the minimum-norm solution of SVD. If $\mathbf{x}_0 = \mathbf{0}$ then $P_0 \mathbf{x}_0$, where P_0 is the projection onto the intersection C_0 , is the minimum-norm solution. Despite this temptation, POCS (as described here) does not compute $P_0 \mathbf{x}_0$ but only guarantees that it finds a solution that lies within C_0 . Figure 4.2 illustrates

this. Each of three paths converges to a point common to three convex sets. Only the order of application of the projection operators distinguishes between the paths. Each path leads to a different but correct solution, but only one solution is $P_0\mathbf{x}_0$, and this is by coincidence only. The ultimate solution varies with the order of projection operators. If the minimum-norm property is essential, then one may use Hildreth's method [8] or Brègman's method [7] [8].

4.3.2 Constraint Sets for Limited-Angle CT

Application of POCS to the limited-angle CT problem with data fusion requires specification of convex constraints and corresponding projection operators. The flexibility of POCS leads to diversity of constraint sets for CT. Three different constraint types are sufficient for the problem at hand. These are:

1. fit to raysum data,
2. fit to range and ultrasound data, and
3. amplitude constraints.

Equation (4-28), below, forms the basis of the first constraint type:

$$R\mathbf{x} = \mathbf{y}. \quad (4-28)$$

Consider each raysum measurement, y_i , independently. Each measurement constrains the solution to lie within the hyperplane defined by:

$$\langle \mathbf{r}_i, \mathbf{x} \rangle - y_i = 0, \quad (4-29)$$

where \mathbf{r}_i is the i^{th} row of R , and $\langle \cdot, \cdot \rangle$ denotes an inner product. For m raysum measurements there are m hyperplanes each of which is a convex set given by:

$$C_{R_i} = \{ \mathbf{x} : \langle \mathbf{r}_i, \mathbf{x} \rangle - y_i = 0 \}, \quad (4-30)$$

where the subscript R refers to raysum. ART is the POCS method that uses only this type of constraint set. The corresponding projection operator is found in many sources dating back to Brègman [7] [8] [27] [28]. The projection operator for C_{R_i} is:

$$P_{R_i} \mathbf{z} = \begin{cases} \mathbf{z} + \frac{y_i - \langle \mathbf{r}_i, \mathbf{z} \rangle}{\langle \mathbf{r}_i, \mathbf{r}_i \rangle} \mathbf{r}_i, & \mathbf{z} \notin C_{R_i} \\ \mathbf{z}, & \text{otherwise} \end{cases} \quad (4-31)$$

The constraints of Equation (4-30) require a precise fit to data. Because it is essential to POCS that C_0 be non-empty, such precise constraints are too restrictive. Consider instead this modified linear constraint, a hyperslab:

$$C_{R_i} = \{ \mathbf{x} : |\langle \mathbf{r}_i, \mathbf{x} \rangle - y_i| \leq \varepsilon_R \}, \quad (4-32)$$

where the allowed margin for error in the fit to raysum data, ε_R , determines the thickness of the slab. Equation (4-32) breaks down into the two intersecting half planes:

$$\begin{aligned} C_{R_{i,a}} &= \{ \mathbf{x} : \langle \mathbf{r}_i, \mathbf{x} \rangle - y_i \leq \varepsilon_R \} \\ C_{R_{i,b}} &= \{ \mathbf{x} : y_i - \langle \mathbf{r}_i, \mathbf{x} \rangle \leq \varepsilon_R \} \end{aligned} \quad (4-33)$$

The modified projection operators are:

$$P_{R_{i,a}} \mathbf{z} = \begin{cases} \mathbf{z} + \frac{y_i + \varepsilon_R - \langle \mathbf{r}_i, \mathbf{z} \rangle}{\langle \mathbf{r}_i, \mathbf{r}_i \rangle} \mathbf{r}_i, & \langle \mathbf{r}_i, \mathbf{z} \rangle - y_i > \varepsilon_R \\ \mathbf{z}, & \text{otherwise} \end{cases} \quad (4-34)$$

and

$$P_{R_{i,b}} \mathbf{z} = \begin{cases} \mathbf{z} + \frac{y_i - \varepsilon_R - \langle \mathbf{r}_i, \mathbf{z} \rangle}{\langle \mathbf{r}_i, \mathbf{r}_i \rangle} \mathbf{r}_i, & y_i - \langle \mathbf{r}_i, \mathbf{z} \rangle > \varepsilon_R \\ \mathbf{z}, & \text{otherwise} \end{cases} \quad (4-35)$$

Equations (4-34) and (4-35) are simple variations of Equation (4-31). The modified constraint sets allow control over the margin of error, ε_R , to vary the size of the sets and ensure that C_0 is non-empty.

There are two ways to incorporate the second constraint type, fit to fusion data. The first is analogous to the fit to raysum data. Fusion constraints originate from the linear equation:

$$W\mathbf{x} = W\mathbf{x}_F, \quad (4-36)$$

where W is a diagonal matrix with elements on the diagonal equal to one or zero. Ignore all the trivial rows of W because they provide no useful constraints. The remaining rows define the convex sets:

$$C_{F_i} = \left\{ \mathbf{x} : |w_i(x_i - x_{F_i})| \leq \varepsilon_F \right\} \quad (4-37)$$

where the subscript F refers to fusion. The corresponding projection operator is:

$$P_{F_i} \mathbf{z} = \begin{cases} \mathbf{z} : z_i = x_{F_i} + \frac{\varepsilon_F}{w_i}, & z_i - x_{F_i} > \frac{\varepsilon_F}{w_i} \\ \mathbf{z} : z_i = x_{F_i} - \frac{\varepsilon_F}{w_i}, & x_{F_i} - z_i > \frac{\varepsilon_F}{w_i} \\ \mathbf{z} & \text{otherwise} \end{cases} \quad (4-38)$$

The projection operator above clips pixel values to lie within a region about the partial reference image \mathbf{x}_F , with the size of the region determined by ε_F and w_i .

A second approach to fusion constraints originates with Oskoui-Fard and Stark [27]. They use a full reference image to constrain limited-angle CT. The concept is similar to the fusion problem here. Oskoui-Fard and Stark assume that they have a complete reference image, i.e., $W = I$, and define the convex set:

$$C_F = \left\{ \mathbf{x} : \|\mathbf{x} - \mathbf{x}_F\| \leq \varepsilon_F \right\}. \quad (4-39)$$

The corresponding projection operator is:

$$P_F \mathbf{z} = \begin{cases} \mathbf{x}_F + \varepsilon_F \frac{\mathbf{z} - \mathbf{x}_F}{\|\mathbf{z} - \mathbf{x}_F\|}, & \|\mathbf{z} - \mathbf{x}_F\| > \varepsilon_F \\ \mathbf{z} & \text{otherwise} \end{cases} \quad (4-40)$$

Unfortunately, the assumption of a full reference image is not practical and, in the limited-angle CT problem at hand, the reference image from the range and ultrasound data can only be partial. Redefinition of the constraint set for a partial reference gives:

$$C_F = \left\{ \mathbf{x} : \|W(\mathbf{x} - \mathbf{x}_F)\| \leq \varepsilon_F \right\}, \quad (4-41)$$

with the corresponding projection operator:

$$P_F \mathbf{z} = \begin{cases} (I - W)\mathbf{z} + W\mathbf{x}_F + \varepsilon_F \frac{W(\mathbf{z} - \mathbf{x}_F)}{\|W(\mathbf{z} - \mathbf{x}_F)\|}, & \|W(\mathbf{z} - \mathbf{x}_F)\| > \varepsilon_F \\ \mathbf{z} & \text{otherwise} \end{cases} \quad (4-42)$$

For reasons explained in Section 4.3.3 the latter form of constraint is best for the limited-angle CT fusion system.

Amplitude constraints are the third and last constraint type. It is physically impossible to have a negative linear attenuation coefficient. Also, one can usually

determine a maximum allowable linear attenuation based on knowledge of what materials are in the object. Therefore, the amplitude of a pixel in a reconstruction must lie between a lower and an upper bound, giving the following convex constraint sets:

$$C_{A_i} = \{ \mathbf{x} : a \leq x_i \leq b, \quad 0 \leq a < b \}, \quad (4-43)$$

where the subscript A refers to amplitude. The C_{A_i} are also hyperslabs. Their corresponding projection operator is:

$$P_{A_i} \mathbf{z} = \begin{cases} \mathbf{z} : z_i = a & z_i < a \\ \mathbf{z} : z_i = b & z_i > b \\ \mathbf{z} & \text{otherwise} \end{cases} . \quad (4-44)$$

Amplitude constraints amount to simply clipping the data so that pixel values lie between the prescribed upper and lower bounds.

4.3.3 Constraint and Parameter Selection

The previous section describes two versions of data fusion constraints. The difference between the two is not trivial. To illustrate this, consider the problem:

- Solve for \mathbf{x} from $A\mathbf{x} = \mathbf{b}$ by minimizing the objective function

$$E = \|A\mathbf{x} - \mathbf{b}\|_{\infty}$$

where $\|\cdot\|_{\infty}$ denotes an l_{∞} norm. This problem is equivalent to the linear programming problem:

- Minimize ε subject to the constraints

$$|\langle \mathbf{a}_i, \mathbf{x} \rangle - b_i| \leq \varepsilon, \quad i = 1, 2, \dots, m.$$

The constraints above are identical to the raysum constraints from the previous section.

This leads to the following observations:

1. with the raysum constraints, POCS is solving something like the above linear programming problem for $R\mathbf{x} = \mathbf{y}$, and
2. the optimum value for ε_R (i.e., ε_R as small as possible with C_0 non-empty) gives the l_{∞} solution to $R\mathbf{x} = \mathbf{y}$ subject, of course, to the other constraints in the problem.

It is possible to solve the linear programming problem and find an optimal value for ε_R (e.g., Brègman [7]) but it is simpler instead to estimate what ε_R should be based on the accuracy of the data. If the estimate of ε_R is too large, then POCS may converge to a poor solution. On the other hand, if ε_R is too small, POCS converges to a region of good solutions and then oscillates within it. The following summarizes selection of ε_R :

1. find the optimal value of ε_R by solving the linear programming problem,
2. instead estimate ε_R based on known accuracy of the data, and
3. if estimating ε_R , err on the side of ε_R too small to keep the ultimate solution within a region of good solutions.

Now consider the fusion data constraints. While the raysum constraints find something like an l_∞ solution to $R\mathbf{x} = \mathbf{y}$, the fusion data constraints of the form:

$$C_F = \{\mathbf{x} : \|W(\mathbf{x} - \mathbf{x}_F)\| \leq \varepsilon_F\},$$

find an l_2 solution to $W\mathbf{x} = W\mathbf{x}_F$. As with ε_R , it is easiest to estimate a value for ε_F based on the accuracy of the data, while keeping in mind that it is best to err on the side of ε_F too small.

The difference between the l_∞ and l_2 constraints is significant. While the l_∞ norm is sensitive to any deviation in the data, the l_2 norm can effectively ignore a large local deviation by global averaging. The significance emerges when the different data sources are not consistent. An l_∞ fit to the raysum data forces consideration of all raysum measurements without averaging out local deviations. Should the raysum data not agree with the partial reference image, an l_2 fit allows the reconstruction to deviate locally from the reference so long as the global fit remains. This fusion task demonstrates the flexibility of POCS in facilitating the combination of l_∞ and l_2 constraints in the same problem.

4.3.4 Application to Limited-Angle CT

Limited-angle CT reconstruction using POCS follows these steps:

1. set $\mathbf{x}_0 = \mathbf{0}$.
2. perform an iteration by computing:

$$\mathbf{x}_1 = P_{A_n} P_{A_{n-1}} \dots P_{A_1} P_F P_{R_m} P_{R_{m-1}} \dots P_{R_1} \mathbf{x}_0.$$
3. repeat step 2 to compute the series $\{\mathbf{x}_n\} = \{P^n \mathbf{x}_0\}$ until convergence.

Advantages of the POCS method are:

1. it is fast for the problem at hand,
2. it incorporates varied constraints easily and blends well with constraint-based data fusion, and
3. it has linear convergence under many conditions.

Disadvantages of the POCS method are:

1. its solution depends arbitrarily on the order of application of projection operators, and
2. it requires care in the selection of constraint sets.

In practice the POCS proves to be fast and versatile making it an excellent choice for CT with data fusion.

4.4 Chapter Summary

This chapter presents three numerical algorithms to solve the proposed limited-angle CT problem with data fusion. Table 4.1 below summarizes each of the three methods.

The next chapter shows results from each of the three methods explored here. SVD, because of its computational complexity, is restricted to small sample problems. It does, however, show exactly what happens to the limited-angle CT problem with the proposed data fusion. Regularization with CG and POCS are both adequate for practical sized problem. Results in Chapter 5 show that sacrifices in accuracy made by regularization to achieve a well-conditioned problem are too great. The ill-posedness of the limited-angle problem is not easily handled by regularization when the data are less

than ideal. POCS proves to be the winner in spite of its ambiguity about the solution. It is important to remember that because each of the numerical methods incorporates different constraints, they are, in essence, solving different problems.

Property	Method		
	SVD	CG	POCS
Uniqueness	minimum norm	regularization	not unique - solution depends on order of operators and initial conditions
Conditioning	eliminate small singular values	use regularizing functional	enlarge constraint sets to get C_0 non-empty
Advantages	<ul style="list-style-type: none"> • computes null space and range • can condition the problem • gives well-posed problem 	<ul style="list-style-type: none"> • fast (for the task at hand) • control over conditioning 	<ul style="list-style-type: none"> • fast (for the task at hand) • incorporates a wide variety of convex constraints • converges at a linear rate
Disadvantages	<ul style="list-style-type: none"> • slow • exorbitant storage 	<ul style="list-style-type: none"> • makes a tradeoff between fidelity to data and conditioning 	<ul style="list-style-type: none"> • solution depends on order of operators • must ensure C_0 is non-empty
Constraints	<ul style="list-style-type: none"> • minimize $\ Ax - b\ ^2$ and $\ x\ ^2$ 	<ul style="list-style-type: none"> • $\Omega x = \Omega x_0$ • minimize $\left\ \begin{bmatrix} A \\ \alpha\Omega \end{bmatrix} x - \begin{bmatrix} b \\ \alpha x_0 \end{bmatrix} \right\ ^2$ 	<ul style="list-style-type: none"> • amplitude limits $a \leq x_i \leq b$
Other	<ul style="list-style-type: none"> • projects data onto range of A 	<ul style="list-style-type: none"> • projects data onto range of $\begin{bmatrix} A \\ \Omega \end{bmatrix}$ 	<ul style="list-style-type: none"> • projects onto constraint sets in the domain of A

Table 4.1: Summary of numerical methods considered for limited-angle computed tomography with data fusion.

Chapter 5

Experimentation and Results

This chapter presents experimental results that establish the validity of the proposed limited-angle CT data fusion system. In addition, the experiments explore the capacity of each of the three numerical methods described in Chapter 4 to deal with confounding errors in the data.

The experiments consist of a series of trials. Each trial applies one of the three reconstruction algorithms to some data to produce a reconstructed cross section for analysis. The algorithm, its parameters, and the data distinguish trials. Section 5.1 describes these distinguishing features, including salient features of the algorithm implementations and the applicable parameters. The section then describes the sources of data, both synthetic and real. Synthetic data trials verify the validity of the proposed limited-angle CT system, while the real data trials show practical application of the system. Errors in the real data highlight the capacity of each algorithm to deal with confounding errors.

Section 5.2 presents the trial results organized by algorithm. The first three trials for each algorithm compute reconstructions based on varying fusion constraints. These trials show a significant improvement in accuracy due to the fusion constraints. SVD shows, in addition to improved accuracy, the changing nature of the problem in terms of the number of singular values equal to zero.

Whereas SVD is limited to small problems, regularization combined with the conjugate gradient method (R/CG) and POCS can handle large problems and so are applied to real data. Trials based on R/CG and POCS, and real data, show the success of the proposed limited-angle CT system with real data. They also show the capacity of the reconstruction algorithms to cope with erroneous data, and the response to varying reconstruction parameters.

The salient conclusion of the experimentation is that, when fused with limited-angle CT, constraints based on ultrasound face sheet data improve the accuracy of reconstruction. Although measurement errors confound reconstruction of accurate cross sections from real data, the limited-angle CT system does, nevertheless, work in practice. Of the algorithms used, POCS proves to be superior because it produces accurate reconstructions, is less sensitive to its parameters, and does not force a compromise between image smoothness and conformity to data.

5.1 Description of Experiments

Experimentation presented in this thesis consists of a set of trials where each trial involves:

1. acquisition of some data,
2. application of a reconstruction algorithm (SVD, R/CG, or POCS) to the data, and
3. analysis of the algorithm output.

These items uniquely describe any trial:

1. the algorithm,

2. the parameters, and
3. the data.

This section describes these items by showing key features of the implementation of each algorithm and the sources of data, both synthetic and real.

5.1.1 Implementation of Algorithms

The raysum operator is central to the implementation of each reconstruction algorithm. This section starts by describing a compact representation of the raysum operator, and then describes the manner in which each implementation uses it.

Parameters, initial conditions, and stopping conditions for each algorithm are also given.

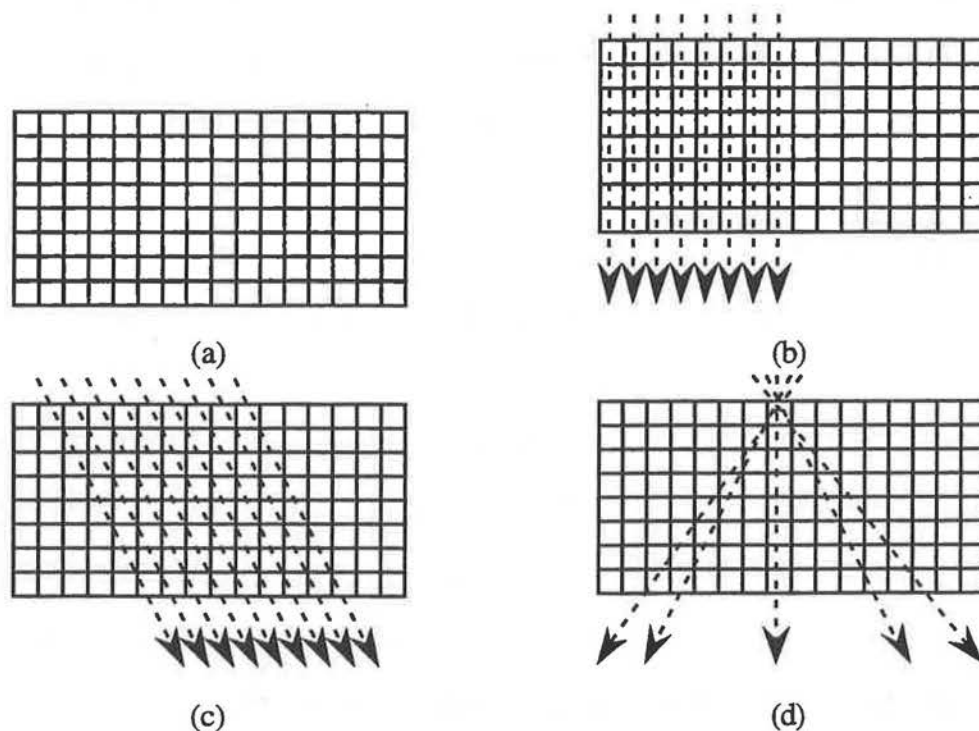


Figure 5.1: Steps to a compact raysum operator: (a) hypothetical grid for reconstruction, (b) parallel rays perpendicular to upper surface of reconstruction, (c) parallel rays at an angle to reconstruction, and (d) set of rays, one at each scan angle.

Raysum Operator

All three numerical methods for reconstruction require a representation of the raysum operator, R . Each row of the raysum operator is a vector in which each element is

the distance an individual ray passes through an element of the reconstruction. The matrix required to represent all rows of the raysum operator completely is huge (207,360,000 elements for a 200 *pixel* by 72 *pixel* reconstruction). However, most elements of each row are zero, i.e., R is sparse, and there are patterns in the operator that lead to a compact representation.

Figure 5.1 (a) shows a hypothetical grid for reconstruction. Elements of the raysum operator rows are the distances a particular ray travels through each box in the grid. A clipping algorithm (from computer graphics) computes these distances. Figure 5.1(b) shows the set of rays corresponding to a scan across the specimen parallel to vertical. A scan at an angle to vertical measures raysums for the set of rays shown in Figure 5.1 (c). Rays for a scan at a constant angle are parallel, and parallel raysums are identical to each other except that the columns (of the reconstruction grid) are offset corresponding to the motion of the scan. With the exception of the column offsets, the set of rays in Figure 5.1 (d), where there is one ray for each scan angle, completely describes all rays used in sampling. Therefore, to get a compact representation of the raysum operator:

1. record only one row of the operator for each scan angle, and
2. record only non-zero elements.

The compact raysum operator is a list of vectors with one element in the list for each scan angle. Each vector is a list of non-zero distances as well as the row and column coordinates for the corresponding grid elements. To reconstruct a row of the raysum operator quickly from the compact representation:

1. assume the row consists of all zeros,
2. select the vector for the appropriate angle from the compact representation,
3. apply the appropriate column offset to the column coordinates of the vector, and

4. put the offset elements of the compact representation into the reconstructed row.

R has one row for each combination of scan angle and column offset, but excludes those combinations that produce rays passing beyond the width of the grid. Only SVD requires the full raysum operator representation. For other algorithms, the compact representation allows quick recovery of the elements needed for computation.

Implementation of SVD

SVD requires an explicit representation of the matrix $\begin{bmatrix} R \\ W \end{bmatrix}$. Implementation of

the proposed limited-angle CT reconstruction follows these steps:

1. declare matrices for A , V , and Σ (note that since Σ is diagonal only the elements on the diagonal are represented),
2. put rows of R into rows of A from the compact raysum operator as described above,
3. put rows of W into rows of A ,
4. if A has fewer rows than columns add sufficient trivial rows to make A square,
5. call SVD subroutine, `svdcmp()` [31],
6. set non-zero singular values that are too small to zero, i.e., condition the matrix,
7. create the vector $\mathbf{b} = \begin{bmatrix} \mathbf{y} \\ W\mathbf{x}_F \end{bmatrix}$, and
8. compute the reconstruction from the decomposition and \mathbf{b} with `svbksb()` [31].

Because `svdcmp()` specifically deals with singular matrices it does not require any special precautions.

Implementation of R/CG

Section 4.2.3 gives the conjugate gradient method for finding the least squares solution to the positive definite system:

$$A\mathbf{x} = \mathbf{b}.$$

Two things are necessary to implement the method:

1. an algorithm to compute $A\mathbf{x}$, and
2. an algorithm to compute \mathbf{b} .

From Equation (4-24) we have:

$$A\mathbf{x} = (R^T R + W^T W + \alpha^2 D_1^T D_1 + \alpha^2 D_2^T D_2)\mathbf{x}, \text{ and}$$

$$\mathbf{b} = R^T \mathbf{y} + W^T \mathbf{x}_F.$$

Break the computation of $A\mathbf{x}$ down into three separate computations: $R^T R\mathbf{x}$, $W^T W\mathbf{x}$, and $(\alpha^2 D_1^T D_1 + \alpha^2 D_2^T D_2)\mathbf{x}$.

To compute $R^T R\mathbf{x}$, first compute $R\mathbf{x}$ using the compact raysum operator. $R\mathbf{x}$ is a column vector, each element of which is the inner product of a row of R and \mathbf{x} . Computation of the inner products in $R\mathbf{x}$ is straight forward from the compact representation of R .

R^T is the back summation operator. The i^{th} element of the vector $R^T \mathbf{b}$ is:

$$[R^T \mathbf{b}]_i = \sum_j r_{j,i} \mathbf{b}_j,$$

i.e., $[R^T \mathbf{b}]_i$, is the weighted sum of all the raysums for rays that pass through the corresponding grid element. Computation on this basis is inefficient, but a more efficient approach arises from observing that each raysum contributes to a set of grid elements in the back summation. That set of elements is the same as the set that contributed to the raysum. The weight applied to a raysum for its contribution to an element of the back summation is the same as the weight for that element in the original raysum. Thus the back summation can be computed in a manner analogous to the raysum. For the experimentation here, the vector $R^T R\mathbf{x}$ is computed by a two step process: first $R\mathbf{x}$ and then $R^T(R\mathbf{x})$. The compact raysum operator allows fast computation of both steps.

The matrix W is diagonal so $W^T = W$ and $W^T W = \text{diag}(w_1^2, w_2^2, \dots, w_N^2)$. From this it is easy to compute the vector $W^T W \mathbf{x}$.

$\alpha^2 D_1^T D_1 + \alpha^2 D_2^T D_2$ is a discrete Laplacian operator. Convolution with the kernel shown in Figure 5.2 best summarizes its implementation. Should different α^2 for the x and y directions be necessary, the more general kernel of Figure 5.3 applies.

$$\begin{array}{|c|c|c|} \hline & -\alpha^2 & \\ \hline -\alpha^2 & 4\alpha^2 & -\alpha^2 \\ \hline & -\alpha^2 & \\ \hline \end{array}$$

Figure 5.2: Convolution kernel for Laplacian operator.

$$\begin{array}{|c|c|c|} \hline & -\alpha_2^2 & \\ \hline -\alpha_1^2 & 2\alpha_1^2 + 2\alpha_2^2 & -\alpha_1^2 \\ \hline & -\alpha_2^2 & \\ \hline \end{array}$$

Figure 5.3: Convolution kernel for Laplacian operator with different coefficients for x and y directions.

Using $R^T R \mathbf{x}$, $W^T W \mathbf{x}$ and $(\alpha^2 D_1^T D_1 + \alpha^2 D_2^T D_2) \mathbf{x}$, compute $A \mathbf{x}$ from:

$$\begin{aligned} A \mathbf{x} &= (R^T R + W^T W + \alpha^2 D_1^T D_1 + \alpha^2 D_2^T D_2) \mathbf{x} \\ &= R^T R \mathbf{x} + W^T W \mathbf{x} + (\alpha^2 D_1^T D_1 + \alpha^2 D_2^T D_2) \mathbf{x} \end{aligned}$$

Computation of \mathbf{b} is similar to $A \mathbf{x}$; the back summation operator gives $R^T \mathbf{y}$ and the identity $W^T = W$ gives $W^T \mathbf{x}$. Add these two vectors to get $\mathbf{b} = R^T \mathbf{y} + W^T \mathbf{x}_F$.

Initial and stopping conditions are important to an iterative procedure such as CG. Although any arbitrary starting condition, \mathbf{x}_0 , is acceptable, these experiments use $\mathbf{x}_0 = W \mathbf{x}_F$ (or $\mathbf{x}_0 = \mathbf{0}$ for trials omitting fusion data) to make maximum use of the fusion data to avoid unnecessary iterations. Iteration stops when the norm of the residual, $\|A \mathbf{x} - \mathbf{b}\|$, falls below a specified threshold.

Implementation of POCS

POCS computes the sequence $\{\mathbf{x}_n\} = \{T^n \mathbf{x}_0\}$ where $T = T_m, T_{m-1}, \dots, T_1$ and $T_i = 1 + \lambda_i(P_i - 1)$. There is no known method for determining optimum values for the relaxation coefficients λ_i , so experiments here use $\lambda_i = 1$ for all i . This reduces to the sequence $\{\mathbf{x}_n\} = \{P^n \mathbf{x}_0\}$ where $P = P_m, P_{m-1}, \dots, P_1$ and P_i is the projection operator onto the i^{th} constraint set.

Three types of projection operators are necessary for the three different types of constraints. The first is the raysum constraints. Equations (4-34) and (4-35) give the two projection operators required for each row of the raysum operator R . The salient feature of these equations is the inner products $\langle \mathbf{r}_i, \mathbf{z} \rangle$ and $\langle \mathbf{r}_i, \mathbf{r}_i \rangle$. The compact raysum operator allows quick computation of these inner products.

The second constraint type is the data fusion constraint. Experiments presented here use Equation (4-42) to implement the data fusion constraints. The computation of the projection is straightforward; it involves only computing the inner products of vectors that are directly available.

The third, and last, constraint is the amplitude constraint. Its implementation is trivial. Simply clip all the elements of the solution to lie within the specified upper and lower bounds.

As with conjugate gradient, it is important to specify starting and stopping conditions. Unlike the conjugate gradient method, the starting point for POCS changes the ultimate solution. Although POCS does not generally compute the minimum-norm solution, these experiments use $\mathbf{x}_0 = \mathbf{0}$ in order to arrive at a solution near the minimum-norm solution. This starting condition does not guarantee such an outcome, but works well in practice. POCS iteration does not introduce spurious signals into the solution so $\mathbf{x}_0 = \mathbf{0}$ is a sensible starting point. Iteration terminates when the norm $\|\mathbf{x}_i - \mathbf{x}_{i-1}\|$ falls below a specified level, i.e., the solution stops moving (or slows sufficiently). Note that

unlike the conjugate gradient, this stopping condition does not guarantee that the solution is good, just that the solution is not going anywhere.

5.1.2 Synthetic Data

This section describes data synthesis for synthetic data trials. Data synthesis starts with an arbitrary cross section image as the subject of an experiment. Application of the raysum operator to the image yields the raysum data for the trial. The cross section also defines the spatial support data and the face sheet data directly.

Application to aircraft parts provides the initial motivation for this research. Accordingly, aircraft parts, in particular sandwich structures with graphite/epoxy composite face sheets around aluminum honeycomb core, are the basis for synthetic data. Thin aluminum in the honeycomb core pushes the resolution limits of the real data acquisition apparatus (described in the next section). To avoid problems with resolution, the real-data experiments use a plexiglass phantom. In the phantom, two plexiglass face sheets surround a set of vertical plexiglass members simulating honeycomb. All plexiglass parts are 3 *mm* thick which, with a reconstruction resolution of 0.5 *mm* by 0.5 *mm*, translates to 6 *pixels* thick in the reconstruction. Real honeycomb is thinner than 1 *pixel*. Synthetic data in these experiments simulate the conditions of the plexiglass phantom.




Figure 5.4: Synthetic cross section image for SVD trials.

Both time and numerical precision limit the SVD problem size to 10 *pixels* by 30 *pixels*. Figure 5.4 shows the synthetic image for SVD trials. Two parallel face sheets and

honeycomb are 2 *pixels* thick. Within the honeycomb cells, two blocks mimic the plexiglass inserts of the real phantom.

R/CG and POCS trials use 72 *pixels* by 200 *pixels*. At the specified resolution of 0.5 *mm* by 0.5 *mm* this corresponds to a cross section 36 *mm* thick by 100 *mm* wide. Low resolution restricts SVD trials to parallel face sheets, but R/CG and POCS trials have no such restriction. Non-parallel face sheets make the trials more interesting and they also quell possible doubts that results may rely on parallel face sheets. Figure 5.5 shows the synthetic cross section for R/CG and POCS synthetic data trials. Face sheets and vertical members are all 6 *pixels* thick.




Figure 5.5: Synthetic cross section image for R/CG and POCS trials.


To conduct synthetic data trials it is necessary to synthesize data corresponding to the three data sources. Application of the forward raysum operator gives the synthetic raysum data. All synthetic data trials compute raysum data this way using single precision floating point arithmetic. CT reconstruction is tolerant of unbiased noise on the data, so trials with unbiased noise added are uninformative. In the absence of a better model for measurement errors, synthetic data trials use error-free data.

Range and ultrasound data give a map of regions that have known linear attenuation coefficients and the known values. Trials using raysum and spatial support data only assume regions external to the specimen are known and have linear attenuation coefficients of zero. With the addition of face sheet data, trials assume external and face sheet regions have known linear attenuation values: zero for external and linear attenuation of plexiglass for face sheets. Figure 5.6 shows the synthetic fusion data as

images. Values from the image of known regions form the diagonal of W while the values of known linear attenuation coefficients form x_F . For synthetic trials, the assumed linear attenuation of plexiglass is 0.40 cm^{-1} .



(a)



(b)



(c)



(d)

Figure 5.6: Synthetic fusion data vectors: for spatial support data only (a) known region and (b) known linear attenuation, and for spatial support and face sheet data (c) known region and (d) known linear attenuation.

5.1.3 Real Data

Real data come from measurements of a plexiglass phantom simulating a sandwich structure with graphite/epoxy composite face sheets and aluminum honeycomb core. As mentioned in the previous section, the thin honeycomb core pushes beyond the resolution of the experimental apparatus. So for the sake of experimentation, components of the phantom are thicker than what is actually found in real aluminum structures.

Figure 5.7 shows schematically the cross section of the plexiglass phantom.




Figure 5.7: Cross section of plexiglass phantom simulating a sandwich structure with graphite/epoxy composite face sheets and aluminum honeycomb core.

Figure 5.8 shows the apparatus for raysum data acquisition. The apparatus originates from a real-time radiography system custom-built by Philips for inspection of aircraft control surfaces. The x-ray tube and linear array are mounted on a C-arm manipulator system (see section 3.2.1). The x-ray tube on one end of the C-arm projects a cone of radiation through the specimen to a sensor on the other end of the arm. The C-shape allows the source and sensor to maneuver about and inspect various aircraft parts.

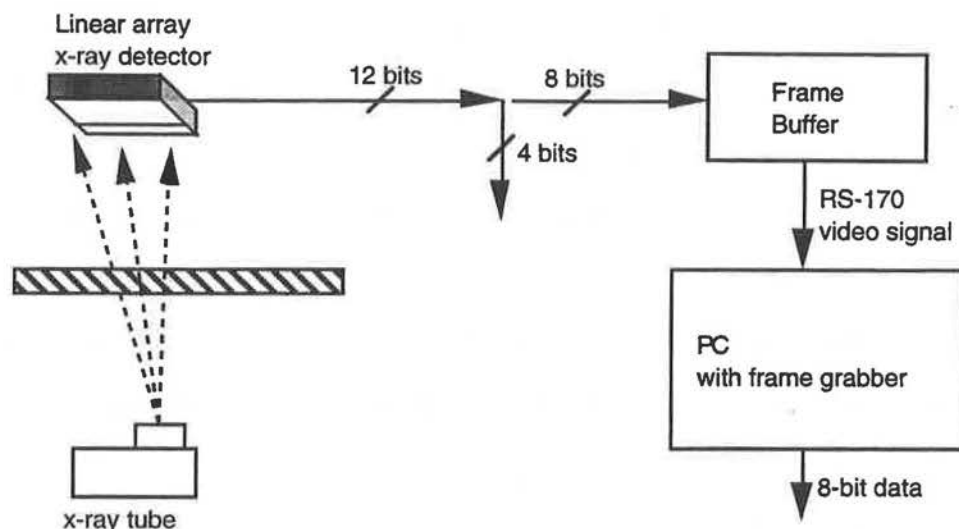


Figure 5.8: Apparatus for limited-angle CT raysum data acquisition system.

In the original configuration, the sensor was a real-time x-ray image intensifier and a video monitor displayed the radiograph in a shielded control room. The image intensifier has since been replaced by a Thompson linear array detector. Whereas the intensifier acquires images of an area, the linear array acquires a single line of data. To acquire an image the *C*-arm sweeps the array across an area. A frame buffer assembles the acquired lines of data into an image. The array is 1024 elements across with an inter-element spacing of 0.2 *mm* and a width of 0.5 *mm*. Analog to digital converters digitize the intensity data to 12 bits but at present the frame buffer uses only the eight most significant bits. The frame buffer produces an RS-170 video signal to display images on a video monitor. A PC-based video frame grabber acquires digital data by sampling the frame buffer RS-170 output.

The experimental trials use the following procedure to acquire raysum data:

1. sweep the *C*-arm across the specimen at a constant angle to the specimen and acquire data at 0.5 *mm* interval spacing,
2. with the frame buffer, assemble the data into an image,
3. digitize the image with the PC-based frame grabber,

4. compute the average of the centre column of the digitized image and the two columns adjacent to it (to compensate for anomalies created by the linear array), and
5. use the data from the averaged column as a set of parallel raysums.

In summary, the x-ray tube projects x-rays onto the linear array. The array creates an electrical signal from the x-rays and digitizes the signal to twelve bits. Eight bits of the twelve go to frame buffer which in turn produces an RS-170 video image. A PC-based frame grabber digitizes the image. Three adjacent columns of the digitized image form a set of parallel raysums.

There is a variety of error sources inherent in the apparatus. There is no filtration of the x-ray radiation other than that inherent in the manufacture of the tube, so beam hardening is a problem. The linear array requires only a fan of radiation, so any excess radiation contributes to scatter.

Raysum measurements are not uniformly sensitive to x-ray intensity measurements. Equation (2-3) describes the raysum as a function of intensity. For a homogenous specimen Equation (2-3) becomes:

$$\mu t = -\ln\left(\frac{I}{I_0}\right),$$

where I is the measured intensity, I_0 is the initial intensity, μ is the linear attenuation coefficient, and t is the material thickness. The sensitivity of μt to I is:

$$\frac{d(\mu t)}{dI} = -\frac{1}{I},$$

which is a function of I . At $I = 225$ the sensitivity is $-1/225 = -0.0044$. A one quantum change in I gives a change in μt of 0.0044. For plexiglass ($\mu = 0.45 \text{ cm}^{-1}$ at 60 kV) this corresponds to a change in material thickness of 0.01 cm. However, at an intensity of 50, the sensitivity is -0.02, corresponding to a 0.044 cm change in plexiglass thickness, about the resolution of the reconstructions. When I is low, i.e., the specimen is thick, the data are more sensitive to intensity errors. The sensitivity of the measurements to a quantum

change in intensity can be of the order of the reconstruction resolution. Certainly this data acquisition system needs improvement, but it is suitable for verifying the proposed limited-angle CT system.

The region reconstructed in experiments is 72 rows by 200 columns which, at a resolution of 0.5 mm by 0.5 mm, corresponds to a specimen 36 mm thick. The plexiglass phantom is 36 mm thick so no pixels in the reconstruction are external to the specimen. The region of reconstruction matches precisely the spatial support of the specimen and, for this experiment only, it is not necessary to acquire range data.

The design of the phantom allows face sheet thickness measurement directly with a micrometer. Therefore, ultrasound measurement is not necessary.* The 3 mm face sheet thickness (as measured by the micrometer) translates to 6 pixels. So for real data trials, the region of known linear attenuation consists of the top and bottom six rows of the reconstruction. The linear attenuation for these rows is that of plexiglass. At the x-ray tube setting of 60 kV used for these experiments, the linear attenuation coefficient of plexiglass is 0.45 cm^{-1} .

5.2 Results

This section presents experimental results. The experiments consist of a series of trials wherein each trial involves application of one numerical method to one data set to produce a reconstruction. The following three items distinguish trials:

1. the algorithm,
2. the parameters for the algorithm, and
3. the data.

For each numerical method, there are synthetic data trials followed by real data trials. The first three synthetic data trials for each algorithm produce reconstructions from:

* In fact, if one were to use ultrasound to measure face sheet thickness, the micrometer measurements would be used to calibrate the ultrasound measurements.

1. raysum data only,
2. raysum and spatial support data, and
3. raysum, spatial support and face sheet data.

Results of these three trials show the improvement in reconstruction accuracy with the incorporation of constraints.

Real data trials show (for R/CG and POCS):

1. the success of the proposed limited-angle CT with real data,
2. the ability of the algorithms to cope with errors in the data, and
3. the effect of parameters.

Real data are more effective than synthetic data for items two and three above. Synthetic data lead too easily to good reconstructions for a wide range of parameter values, even with Gaussian noise added. However, errors in the real data make the effects of the parameters obvious.

An error measure is essential for comparing the accuracy of the different methods. Experimental results presented here use the measure, e , defined by [28]:

$$e = \frac{\|\mathbf{x} - \hat{\mathbf{x}}\|}{\|\mathbf{x}\|} \times 100\%, \quad (5-1)$$

where \mathbf{x} is the true cross section and $\hat{\mathbf{x}}$ is the reconstructed cross section.

5.2.1 SVD Trials

Limitations of the SVD algorithm restrict SVD experimentation to trials based on the low resolution cross section of Figure 5.4. Raysum data for these trials cover seven scan angles from -60° to 60° in steps of 20° . Spatial support constraints assume the top and bottom rows of the reconstruction are known to be zero. Face sheet constraints assume the second and third rows from the top and bottom are known to be plexiglass.

As for all the synthetic data trials, SVD synthetic data trials give results for varying levels of constraints. SVD gives a compelling illustration of the change in the reconstruction problem with added constraints by recording:

1. the number of raysum rows in A ,
2. the number of spatial support rows in A ,
3. the number of face sheet rows in A ,
4. the number of singular values equal to zero, and
5. the error measure.


Figure 5.9 shows the reconstructions for each of the three trials and the corresponding absolute error vectors. Grey scales in the images are set to allow direct comparison of the reconstructions and the errors, but the maximum and minimum values (listed in the caption) differ for each trial. Figure 5.9(g) shows the absolute error vector for trial 3 rescaled to show more detail. Table 5.1 summarizes the SVD trial results.

Trial Number	Description	Rows of A	Number of Singularities	Error Measure e (%)
1	raysums only	<ul style="list-style-type: none"> • 152 raysum • 148 trivial 	148	50.7
2	raysums and spatial support	<ul style="list-style-type: none"> • 152 raysum • 60 spatial support • 88 trivial 	88	31.5
3	raysums, spatial support, and face sheets	<ul style="list-style-type: none"> • 152 raysum • 60 spatial support • 120 face sheet • 0 trivial 	13	4.8


Table 5.1: SVD trial results.

Results in Table 5.1 show how spatial-support and face-sheet constraints add rows to the matrix A . The number of singularities (singular values equal to zero) in A reflects the degree to which the problem is ill-posed. Reliance on raysum data alone is the poorest scenario with 148 singularities out of 300 singular values. Spatial support constraints improve upon the raysum data, reducing the number of singularities to 88, but still leaving the problem ill-posed. Incorporation of face-sheet constraints further reduces


the number of singularities to 13. Although the use of data fusion still results in an ill-posed problem, it markedly reduces the number of singularities and size of the null space.




(a)



(b)




(c)




(d)

Figure 5.9: Reconstructed images for SVD trials: (a) trial 1 reconstruction (max = 0.48, min = -0.01), (b) trial 1 error vector (max = 0.30, min = 0.0), (c) trial 2 reconstruction (max = 0.55, min = 0.0), and (d) trial 2 error vector (max = 0.23, min = 0.0).


Continued on next page.



(e)



(f)



(g)

Figure 5.9 (continued): Reconstructed images for SVD trials: (e) trial 3 reconstruction (max = 0.42, min = -0.04), (f) trial 3 error vector (max = 0.04, min = 0.0), and (g) rescaled trial 3 error vector (max = 0.04, min = 0.0).

Not only does the data fusion reduce the size of the null space, it improves accuracy because the reduced null space does not contain the face sheet structure. However, even after addition of the fusion constraints the null space is not trivial and it is possible to have a specimen with components in the remaining null space. The minimum-norm solution would not be accurate because it omits the null-space components. Such components would be wide horizontal edges in the interior of the specimen.

From data that are readily available in trials 1 through 3 it is possible to compute a Class I weakly-coupled fusion reconstruction, \mathbf{x}_{w-c} , from:

$$\mathbf{x}_{w-c} = W\mathbf{x}_F + (I - W)\hat{\mathbf{x}}$$

where $\hat{\mathbf{x}}$ is the solution to $R\mathbf{x} = \mathbf{y}$, i.e., the reconstruction in Figure 5.9(a). Intuitively, the weakly-coupled reconstruction is the partial reference image pasted onto of the unconstrained reconstruction. Figure 5.10 shows this weakly-coupled reconstruction and its absolute error vector. The error measure for the reconstruction is 13.9% (versus 4.8% for the strongly-coupled reconstruction). Clearly, the strongly-coupled constrained-based system produces superior results.




Figure 5.10: (a) Weakly-coupled reconstruction using SVD (max = 0.48, min = -0.01), and (b) the corresponding absolute error vector (max = 0.15, min = 0.0).

5.2.2 R/CG Trials

All R/CG trials compute 72 pixel by 200 pixel reconstructions. The first three trials use synthetic data based on Figure 5.5, while the remaining trials use real data from measurements of the plexiglass phantom of Figure 5.7. Both synthetic and real data trials use raysums for 13 scan angles from -60° to 60° in steps of 10° .

Figure 5.11 shows the reconstructions for the first three trials and the corresponding absolute error vectors. As for all synthetic data trials, the first three use varying levels of constraints. The parameter α^2 is set to a constant value, $\alpha^2 = 0.001$. The stopping condition is $\|Ax - b\| < 0.1$. Table 5.2 summarizes the R/CG synthetic data trials.

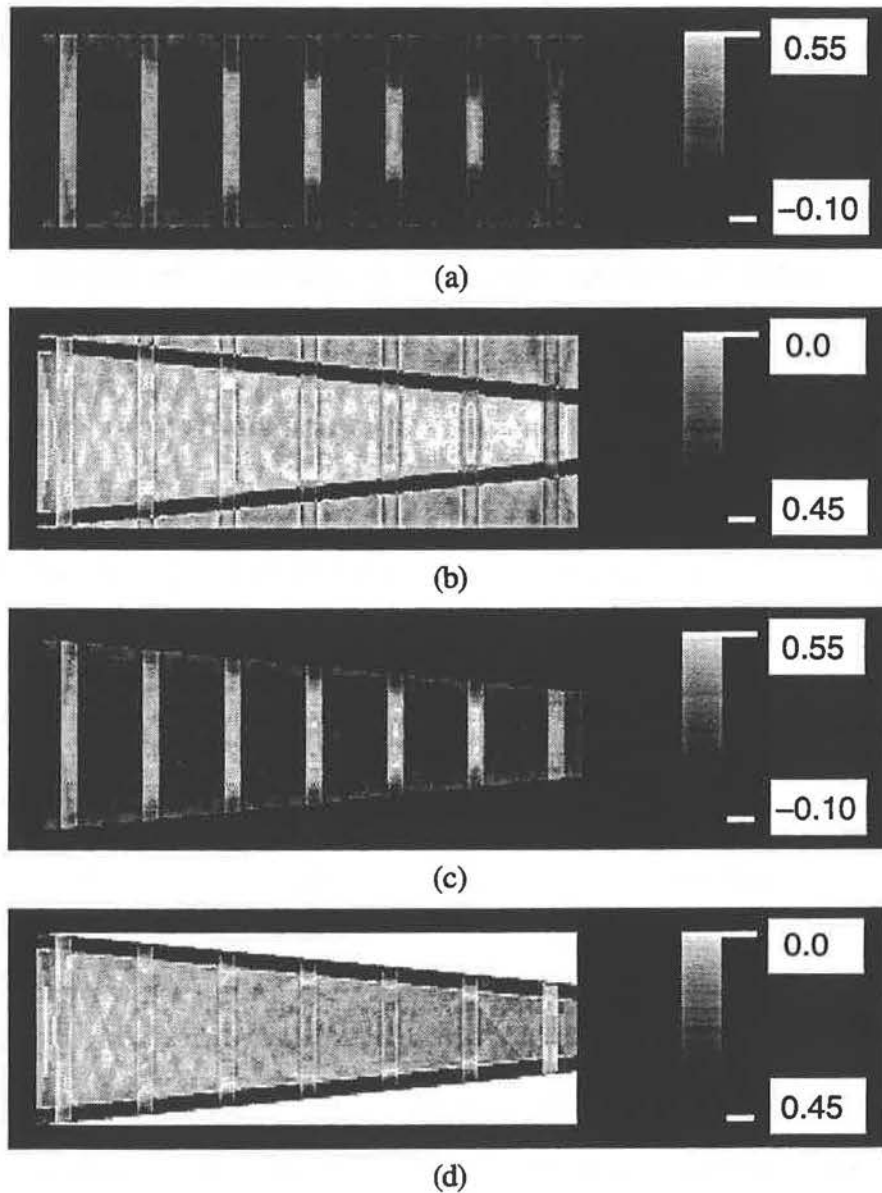




Figure 5.11: Reconstructed images for R/CG synthetic data trials: (a) trial 1 reconstruction (max = 0.49, min = -0.06), (b) trial 1 error vector (max = 0.40, min = 0.0), (c) trial 2 reconstruction (max = 0.55, min = -0.05), and (d) trial 2 error vector (max = 0.43, min = 0.0). Continued on next page.



(e)



(f)

Figure 5.11 (continued): Reconstructed images for R/CG synthetic data trials: (e) trial 3 reconstruction (max = 0.45, min = -0.07), and (f) trial 3 error vector.

Trial Number	Description	Regularization Parameter α^2	Iterations	Error Measure e (%)
1	raysums only	0.001	10	64.0
2	raysums and spatial support	0.001	12	55.3
3	raysums, spatial support, and face sheets	0.001	13	6.7

Table 5.2: R/CG synthetic data trial results.

The images of Figure 5.11 show subjectively the improvement in accuracy attained by incorporating fusion constraints. Error measures in Table 5.2 show this improvement quantitatively. As was the case for SVD results, with R/CG too there is a great improvement in reconstruction accuracy for trials using data fusion.

In trial three, the error measure for R/CG is comparable to that for SVD (6.7% versus 4.8%). The small difference is not significant and may be due to the difference in the data and problem size. Figure 5.11 exhibits an oscillatory pattern in the interior of the reconstruction. The reason for its existence is well known. Finite precision means that the synthetic data do not precisely match the forward model. Consequently, R/CG uses part of the null space of $\begin{bmatrix} R \\ W \end{bmatrix}$ (an oscillatory pattern) to match the data. A larger value for α^2 prevents this, but at the expense of smoothing the solution. The reconstructions of Figure 5.11 are about optimum, based on trial-and-error selection of α^2 .

Figure 5.12 shows the reconstruction of the synthetic cross section using weakly-coupled fusion. The error measure for the weakly-coupled reconstruction is 19.6% (versus 6.7% for the strongly coupled reconstruction), which indicates again the superiority of the strongly-coupled method.




Figure 5.12: (a) Weakly-coupled reconstruction using R/CG (max = 0.49, min = -0.05), and (b) the corresponding absolute error vector (max = 0.22, min = 0.0).

The remaining R/CG trials use real data to show:

1. the success of the method with real data,

2. the capacity of R/CG to handle errors in the data, and
3. effects of α^2 on the quality of reconstruction.


There are seven real data trials, designated 4(a) through 4(g). Trials 4(a) through 4(f) are identical except that the parameter α^2 varies from 0.001 to 100.0 in decade steps. Trial 4(g) reconstructs the image using a different α^2 for the x and y directions. All trials use a stopping condition of $\|Ax - b\| < 0.1$. A smaller tolerance does not significantly affect the reconstructions. Figure 5.13 shows the reconstructed images for these trials on the same grey scale to allow for easy comparison. Table 5.3 summarizes the results.

Trial Number	Description	Regularization Parameter α^2	Iterations	Maximum (cm^{-1})	Minimum (cm^{-1})
4 (a)	<ul style="list-style-type: none"> • real data • all constraints 	0.001	20	0.71	-0.30
4 (b)	<ul style="list-style-type: none"> • real data • all constraints 	0.01	20	0.71	-0.27
4 (c)	<ul style="list-style-type: none"> • real data • all constraints 	0.1	25	0.68	-0.20
4 (d)	<ul style="list-style-type: none"> • real data • all constraints 	1.0	40	0.62	-0.17
4 (e)	<ul style="list-style-type: none"> • real data • all constraints 	10.0	55	0.50	-0.05
4 (f)	<ul style="list-style-type: none"> • real data • all constraints 	100.0	102	0.38	0.12
4 (g)	<ul style="list-style-type: none"> • real data • all constraints 	$\alpha_1^2 = 0.001$ $\alpha_2^2 = 0.1$	24	0.68	-0.22


Table 5.3: R/CG real data trial results.

Reconstructions in Figure 5.13, in particular 5.13(c) and 5.13(g), resemble the cross section of the phantom in Figure 5.7. It is fair to conclude that the proposed fusion method for limited-angle computed tomography is successful with real data. One is


forced to admit, however, that at least for R/CG, the quality of reconstruction from real data is poorer than from synthetic data.




(a)



(b)




(c)



(d)


Figure 5.13: Reconstructed images for R/CG real data trials:
(a) $\alpha^2 = 0.001$, (b) $\alpha^2 = 0.01$, (c) $\alpha^2 = 0.1$, and (d) $\alpha^2 = 1.0$.
See Table 5.3 for maxima and minima. Continued on next page.



(e)



(f)



(g)

Figure 5.13 (continued): Reconstructed images for R/CG real data trials: (e) $\alpha^2 = 10.0$, (f) $\alpha^2 = 100.0$, and (g) $\alpha_1^2 = 0.001$ and $\alpha_2^2 = 0.1$. See Table 5.3 for maxima and minima.

R/CG does not appear to handle errors in the data well. In the synthetic case, the vector $\mathbf{b} = \begin{bmatrix} \mathbf{y} \\ W\mathbf{x}_F \end{bmatrix}$ is, within the limits of numerical precision, in the range of $\begin{bmatrix} R \\ W \end{bmatrix}$. Such is not the case with the real data. Consequently, the errors exaggerate the effects of ill-conditioning and larger null space components appear in the reconstruction.

The parameter α^2 does not help deal with errors in this case. It merely allows one to select the degree of compromise in the solution. Figure 5.13(a), for example, shows one extreme with $\alpha^2 = 0.001$. The image is sharp, but oscillations are severe as indicated by the maximum and minimum image values (they should be 0.45 cm^{-1} and

0.00 cm^{-1} respectively). At the other extreme, Figure 5.13(f) shows the reconstruction for $\alpha^2 = 100.0$. In this case the solution is so heavily biased towards a flat image that almost all structural detail is lost. Certainly this would not work with real honeycomb which is much thinner.

The ill-posed nature of this limited-angle CT problem requires more smoothing in the y direction than for x , so one can use a different α^2 for each direction. Figure 5.13(g) shows the reconstruction for $\alpha_1^2 = 0.001$ and $\alpha_2^2 = 0.1$. It supports the conclusion that the proposed method works, but none of the images of Figure 5.13 is particularly satisfying.

Aside from selecting the degree of compromise in the solution, α^2 affects the number of iterations required. Smoother reconstructions require more iterations. Reasonable solutions like those for trials 4(c) and 4(g) take about 25 iterations and do not impose an outrageous computational burden.

5.2.3 POCS Trials

There is a total of seven POCS trials. Three synthetic data trials show the improvement in reconstruction accuracy due to fusion constraints, and a fourth shows the effect of changing the stopping condition. A real data trial shows the success of the proposed limited angle CT method with real data. Two more real data trials illustrate the effect of altering the parameter ϵ_R .

Figure 5.14 shows the reconstructions for the four POCS synthetic data trials and the corresponding absolute error vectors. Because POCS employs amplitude constraints, the reconstructions all have a maximum value of 0.40 and a minimum value of 0.0. Table 5.4 summarizes the results. The first three synthetic data trials do reconstructions based on varying degrees of fusion constraints (as done for SVD and R/CG). The fourth trial shows the effect of changing the stopping condition for the algorithm.









Figure 5.14: Reconstructed images for POCS synthetic data trials: (a) trial 1 reconstruction, (b) trial 1 error vector (max = 0.4, min = 0.0), (c) trial 2 reconstruction, and (d) trial 2 error vector (max = 0.4, min = 0.0). Continued on next page.




(e)



(f)



(g)



(h)

Figure 5.14 (continued): Reconstructed images for POCS synthetic data trials: (e) trial 3 reconstruction, (f) trial 3 error vector (max = 0.22, min = 0.0), (g) trial 4 reconstruction, and (h) trial 4 error vector (max = 0.25, min = 0.0).

Trial Number	Description	Parameters	Iterations	Error Measure e (%)
1	raysums only	$\epsilon_R=0.001$ $\epsilon_F=0.1$ stop at $\ \mathbf{x}_i - \mathbf{x}_{i-1}\ < 0.1$	13	62.6
2	raysums and spatial support	$\epsilon_R=0.001$ $\epsilon_F=0.1$ stop at $\ \mathbf{x}_i - \mathbf{x}_{i-1}\ < 0.1$	33	38.9
3	raysums, spatial support, and face sheets	$\epsilon_R=0.001$ $\epsilon_F=0.1$ stop at $\ \mathbf{x}_i - \mathbf{x}_{i-1}\ < 0.1$	9	6.0
4	raysums, spatial support, and face sheets	$\epsilon_R=0.001$ $\epsilon_F=0.1$ stop at $\ \mathbf{x}_i - \mathbf{x}_{i-1}\ < 0.001$	146	5.5

Table 5.4: POCS synthetic data trial results.

Trials 1 through 3 exhibit the same marked improvement in accuracy of reconstruction observed for SVD and R/CG. Trial 3 shows that POCS computes a reconstruction comparable in accuracy to SVD and R/CG (6.0% for POCS versus 4.8% for SVD and 6.7% for R/CG). POCS, however, does not make a compromise towards a flat solution, while it is on par with R/CG computationally (i.e., it converges in about the same number of iterations with about the same amount of work per iteration). Of the three methods, only POCS manages to provide an accurate solution quickly without smoothing.

Only nine iterations are necessary for trial 3. This is not a lot of iterations, which compels one to ask whether or not stricter stopping conditions will improve the accuracy further. Trial 4 computes a reconstruction based on the stricter stopping condition

$\|\mathbf{x}_i - \mathbf{x}_{i-1}\| < 0.001$. The result is a large increase in the number of iterations (from 9 to 146) as expected, but only a slight improvement in accuracy (6.0% down to 5.5%). Obviously the stricter stopping condition has not paid off in a significant improvement in accuracy.

Figure 5.15 shows the reconstruction of the synthetic cross section using weakly-coupled fusion. As is the case for both SVD and R/CG, the weakly-coupled reconstruction is poorer than the strongly-coupled one (error measures of 18.0% versus 6.0%), showing the superiority of the strongly-coupled method.




Figure 5.15: (a) Weakly-coupled reconstruction using POCS, and (b) the corresponding absolute error vector (max = 0.24, min = 0.0).

Figure 5.16 shows the reconstruction for trial 5, based on real data. Parameters for the trial are $\varepsilon_R = 0.01$ and $\varepsilon_F = 0.1$. With a stopping condition of $\|\mathbf{x}_i - \mathbf{x}_{i-1}\| < 0.1$ the trial requires 30 iterations. The reconstruction is qualitatively superior to those for R/CG. Structural detail is clearly visible and the image is not smoothed. However, there are some artifacts visible in the interior of the reconstruction. These artifacts are an inevitable consequence of measurement errors. Although practical application requires

better reconstructions, the POCS method shows that if the errors are reduced, a sharp accurate reconstruction is possible.






Figure 5.16 Reconstructed cross section for POCS real data trial 5.

Two further trials with POCS and real data show the effect of varying ϵ_R , the extent to which POCS forces the solution to fit the raysum data. Large ϵ_R relaxes the fit to data (appropriate for unreliable raysums) while a small ϵ_R enforces a tight fit to the raysum data. Figure 5.17 shows the reconstructed cross sections from the real data with $\epsilon_R = 0.1$ and $\epsilon_R = 0.001$. Table 5.5 summarizes the results.



(a)



(b)

Figure 5.17: Reconstructed images from POCS real data trial 6: (a) $\epsilon_R = 0.1$, and (b) $\epsilon_R = 0.001$.

Trial Number	Description	Parameters	Iterations
6 (a)	raysums, spatial support, and face sheets	$\varepsilon_R=0.1$ $\varepsilon_F=0.1$ stop at $\ \mathbf{x}_i - \mathbf{x}_{i-1}\ < 0.1$	7
6 (b)	raysums, spatial support, and face sheets	$\varepsilon_R=0.001$ $\varepsilon_F=0.1$ stop at $\ \mathbf{x}_i - \mathbf{x}_{i-1}\ < 0.1$	32

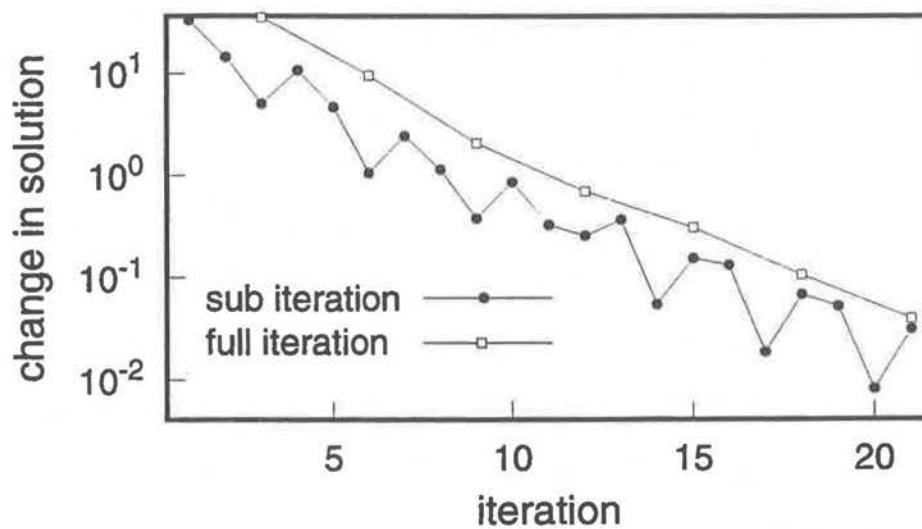
Table 5.5: POCS trial 6 (real data) results.

A large ε_R enlarges the constraint sets to ensure a non-empty intersection, and leads to faster convergence. Table 5.5 indicates that, as expected, POCS converges faster towards a solution with ε_R large. Figure 5.17 show that the reconstruction is of a higher quality with the small ε_R .

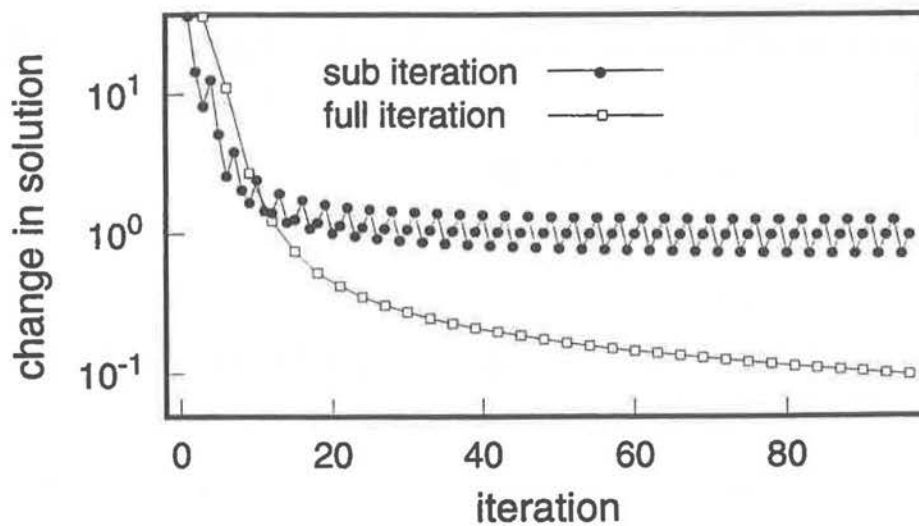
Plots of convergence for trial 6 in Figure 5.18 are instructive in understanding the effects of ε_R . There is one plot each for trials 6(a) and 6(b) with two lines in each plot. One line in each plot has a value for each integer on the x axis, 1, 2, 3, ..., with corresponding y -axis values showing:

1. change in \mathbf{x} from beginning of iteration 1 to after projection onto raysum constraints,
 2. change in \mathbf{x} from after projection onto raysum constraints to after projection onto fusion constraints in iteration 1,
 3. change in \mathbf{x} from after projection onto fusion constraints to after projection onto amplitude constraints in iteration 1,
 4. same as (1) but for iteration 2,
 5. same as (2) but for iteration 2,
 6. same as (3) but for iteration 2,
- etc.

The other line in each plot has data points for x -axis values 3, 6, 9, 12, ..., with corresponding y -axis values indicating the change in x between full iterations. So one line shows the movement in x within iterations (sub-iterations) while the other shows changes in x for full iterations.



(a)



(b)

Figure 5.18: Plots of convergence for POCS: (a) trial 6(a) - $\epsilon_R = 0.1$ and (b) trial 6(b) - $\epsilon_R = 0.001$.

The plot for trial 6(a) shows oscillations in sub-iteration movement that quickly subside along with movement between full iterations. Trial 6(b), on the other hand,

shows sub-iteration movement that continues while the solution slows between iterations. Plot 6(a) suggest a non-empty intersection of constraints while 6(b) appears to have an empty intersection. In 6(b) the solution jumps around in a regular pattern within an iteration while the solution does not move much between iterations. These plots do not provide conclusive evidence though that the intersections are non-empty and empty respectively.

In spite of the apparent empty constraint intersection in trial 6(b), the reconstruction is good. Other than wasted iteration time, there is no penalty for making ϵ_R too small; the solution just bounces around in a region of good solutions. Iterations are fast so it seems wise to select ϵ_R to be small. Only when the solution does not converge does it make sense to set ϵ_R larger.

5.3 Chapter Summary

The salient conclusion of this chapter is that data fusion constraints markedly improve the accuracy of reconstruction produced by the limited-angle CT system. All three reconstruction methods exhibit improved accuracy by including spatial support constraints and further improvement by including face sheet constraints. In particular, the SVD trials show a decrease in the number of singularities due to the added constraints, and a corresponding reduction in the size of the null space. The fusion constraints do not arbitrarily reduce the null space, but they reduce the null space in a manner consistent with measurements of the specimen. This consistency in the reduction of the null space results in the improved accuracy of reconstruction.

Although the source of real data used here is far from ideal, real data trials serve to show that the method is not restricted to synthetic data, but works with real data. In particular, POCS gives good reconstructions with the structure and anomalies clearly visible.

R/CG real data trials illustrate the compromises associated with the regularization parameter α^2 . A large value of α^2 leads to a smooth reconstruction. A small α^2 does

not smooth as much, but conditioning becomes a problem. Selecting α^2 proves to be a compromise between smoothing errors and conditioning errors.

POCS gives better reconstructions than R/CG. POCS reconstructions are sharp and free of the oscillatory patterns exhibited by R/CG. There is, however, some anomalous clutter in the interior of the image due to errors in the data. POCS is also relatively insensitive to its parameters. One can make ϵ_R too small without degrading the reconstruction and suffering only an increase in the number of iterations required.

Chapter 6

Conclusions and Discussion

Section 6.1 of this chapter summarizes the important conclusions of this thesis, i.e.: identification of the inability of limited-angle CT to correctly reconstruct sandwich structures, and the efficacy of the novel system using data fusion with ultrasound to overcome the inability. The section then generalizes the conclusions to cover a broader field of application. Section 6.2 extrapolates from the conclusions by discussing possibilities for further development with respect to practical application of the limited-angle CT system, use of CAD models for constraints, other types of constraints, and application of fusion to Compton back scatter imaging.

6.1 Conclusions

6.1.1 Novel Contributions

Current limited-angle CT techniques rely on general *a priori* assumptions to constrain reconstruction. This thesis contributes the novel observation that the success or failure of reconstruction depends on the validity of these assumptions for the specific specimen. It is not sufficient to claim that some limited-angle CT method is successful without defining the types of specimens that yield accurate reconstruction. In particular,

current methods cannot accurately reconstruct a sandwich structure, as shown in Section 3.1. The methods fail because the face sheets of sandwich structures lie almost entirely in the limited-angle Radon transform null space. Consequently, there is no valid basis for interpolation of the data to fill in missing data. This failure of limited-angle CT is significant because sandwich structures are commonplace; they carry loads due to bending moments economically with respect to weight and material.

Ultrasound measurements complement x-ray data well. Whereas x-rays are insensitive to discontinuities along the direction of the raysum, ultrasound specifically locates discontinuities along its path. This thesis describes a novel method for limited-angle CT that exploits the complementary nature of x-ray and ultrasound by fusing ultrasound measurements into the reconstruction process. The method yields accurate reconstructions of sandwich structures where conventional limited-angle CT cannot.

Ultimately, the fusion system requires formation and solution of the system of linear equations:

$$\begin{bmatrix} R \\ W \end{bmatrix} \mathbf{x} = \begin{bmatrix} \mathbf{y} \\ W\mathbf{x}_F \end{bmatrix}.$$

Raysum data give the equations $R\mathbf{x} = \mathbf{y}$ and range and ultrasound data give $W\mathbf{x} = W\mathbf{x}_F$, where \mathbf{x}_F is a partial reference image and W is a diagonal matrix. The partial reference image contains the linear attenuation coefficients for the face sheet and exterior regions. Weights on the diagonal of W indicate whether or not a pixel in \mathbf{x}_F is known (i.e., in a face sheet or exterior) or unknown (i.e., in the interior). The range and ultrasound equations remove the face sheet structures from the null space so that face sheets do not confound reconstruction, thus allowing accurate reconstruction where it would otherwise be impossible.

Previous limited-angle CT work has used a full reference image [28], but the use of a partial reference image as a constraint set for POCS is novel. Therefore, the mapping of the fusion problem here to the constraint set of Equation (4-41) is also novel.

6.1.2 Generalization of Results

This thesis focuses on a specific problem and a specific type of specimen. Despite the narrow focus, the conclusions have a broader implication. The novel CT method succeeds because the data fusion focuses directly on the limitations of the original problem. One cannot expect that the *ad hoc* addition of data to an inverse problem will improve the solution. Additional constraints that do not focus on uncertainty in the original problem are not beneficial.

For example, suppose that rather than use ultrasound data, the proposed method used neutron radiation. Although neutrons have certain properties that complement x-rays, they are not more sensitive to discontinuities. Since it is the failure of the x-rays to locate discontinuities along the rays that confounds reconstruction, one should expect no improvement due to fusion with neutron data (other than the cancellation of unbiased errors by redundant measurements). Another example is CT, which is itself a fusion problem that fuses a set of two-dimensional radiographs to produce a three-dimensional reconstruction. However, as Chapter 2 points out, just fusing views from a few angles is not sufficient. A full range of angles is necessary for reconstruction. Data acquired for one angle complement well data acquired at a perpendicular angle. This observation has implications for computational vision methods that derive scene descriptions from multiple images, the opaque version of CT *. It is important to select camera views to complement each other, i.e., an image should contain information that the others do not have.

Therefore, a key to success with data fusion is to ensure that additional data sources address uncertainty in the original problem. Ultrasound addresses well the weaknesses of x-rays, so the fusion method succeeds.

* The opaque problem differs from CT because opacity provides additional constraints, but leads to a correspondence problem instead.

6.2 Discussion

This section presents four areas for further development of work in this thesis. First, it considers topics vital to practical application of the proposed limited-angle CT method, including improved apparatus and ability to detect anomalies. Second, Section 6.2.2 discusses the use of CAD models (available for many NDE specimens) as constraints. Third is an examination of other possibilities for constraints, specifically non-convex ones. Fourth, and last, is a discussion speculating about fusion of Compton back scatter data and ultrasound data.

6.2.1 Practical Application

The apparatus used in experiments is a patchwork of devices designed for digital radiograph acquisition and is not well suited to the more rigorous requirements of CT. Although it is sufficient to perform the experiments presented in this thesis, practical application of the proposed limited-angle CT method requires improved apparatus.

Improvements to the apparatus must address sources of error in the data. CT tolerates small amounts of unbiased noise by canceling it out. Therefore, the small amount of photon noise in the data is not an issue in the experimentation. Should photon noise be a problem with a modified apparatus, longer integration times can compensate. Scattered radiation is a source of errors related to the specimen and is difficult to model. The best approach to elimination of scatter is collimation of the x-ray beam. Collimation prevents unnecessary radiation from reaching the specimen and reduces the subsequent scattered radiation. Whereas medical systems rely on collimation to reduce patient exposure to radiation, conventional radiography does not normally use collimation. The apparatus used for experiments in this thesis does not have the collimation that is essential for practical application. Chapter 5, Section 5.1.3, shows the sensitivity of raysum measurements to radiation intensity. The quantization interval in eight-bit intensity data corresponds to a large change in the raysum values. Twelve-bit data can improve this sensitivity, provided that the additional bits are significant. A saner

acquisition process that avoids sending intensity signals through multiple amplifiers, digital-to-analogue and analogue-to-digital converters will help to reduce errors in intensity data and increase the number of significant bits. The linear array scanner is excellent at its intended job, acquisition of digital radiographs, but it is not the best device for CT. For practical application of the proposed method, a more appropriate x-ray sensor is essential.

This thesis uses the l_2 norm of the error, e , as a measure of accuracy where:

$$e = \frac{\|\mathbf{x} - \hat{\mathbf{x}}\|}{\|\mathbf{x}\|} \times 100\%.$$

\mathbf{x} is the true cross section while $\hat{\mathbf{x}}$ is the reconstruction. $\|\cdot\|$ is an l_2 norm. Because the l_2 norm averages the square of the errors over the entire solution, it gives a global measure of accuracy. For example, if $\hat{\mathbf{x}}$ is identical to \mathbf{x} everywhere except for a large error at one sample, the l_2 norm will average the error at that sample over the whole image.

Consequently, the l_2 norm indicates a good global match between \mathbf{x} and $\hat{\mathbf{x}}$ while ignoring the large local error. A greater number of samples averages out large local errors to a greater degree.

A global error measure is well suited to some tasks, e.g., pattern recognition. However, for inspection tasks, where small local anomalies are important, an error measure sensitive to local errors is better. An l_∞ norm of the error has such local properties. For the previous example, the large local error determines the maximum error yielding a large l_∞ norm. The l_∞ norm of the error is important in identification of anomalies because it determines the size of local deviation that is significant. For example, if one expects a pixel in a reconstruction to be 0.5 cm^{-1} but observes instead 0.51 cm^{-1} , the difference can only be interpreted to be significant if the reconstruction has better than 0.01 cm^{-1} accuracy in the l_∞ sense. In the l_2 sense there is no way to know if a local deviation is significant.

Section 4.3.3 of Chapter 4 points out that the POCS constraints lead to a pseudo l_∞ norm solution. However, minimizing the residual, as POCS does, is not necessarily

the same as minimizing the error. The question of how to minimize the local error merits consideration.

The proposed limited-angle CT method has potential for application beyond inspection of sandwich structures. Often an outer shell obscures an internal structure just as face sheets obscure the interior of a sandwich structure. If such an obscured structure also requires limited-angle techniques, fusion with ultrasound will be helpful. The ultrasound allows one to mathematically peel away the outer shell so that it does not confound reconstruction of the interior. Thus, the method is not limited to control surfaces, but also allows inspection of some internal structures *in situ* that would otherwise require disassembly.

6.2.2 CAD Model as a Source of Constraints

With the advent of computer-aided design (CAD) it has become commonplace to design structures with computers. Consequently, many structures targeted for inspection have CAD models available. The question then arises: can CAD models provide a substitute for ultrasound data in the proposed limited-angle CT method? In short, the answer is yes, depending on two factors.

First, it is essential that face sheet data from the CAD model register precisely with the raysum data. This is not a trivial concern, but such registration is within the realm of current technology.

Second, one must consider the effects of errors in the CAD model, i.e., discrepancies between the CAD model and the true structure. Such errors may arise from manufacture or from changes in the structure during service. The success of the method then depends upon the ability of the range and raysum data to correct the model. The null space of the limited-angle Radon transform provides a guide to determining what errors are correctable. Raysum data cannot correct errors that are wide and thin, like the face sheets, while they can correct narrow ones. Therefore, small local variations in the face

sheet data should not be a problem. A global error in face sheet thickness or a wide anomaly in the face sheet will confound reconstruction.

POCS provides an advantage here. The mixture of l_2 and l_∞ fits allowed by POCS means that raysum data can make local corrections to the CAD data. While the l_∞ fit forces a degree of local conformity to the raysum data, some local errors in the CAD model are tolerable, provided that they are not so large that they affect the global fit.

Substitution of ultrasound data with a CAD model offers important benefits. Elimination of ultrasound allows simultaneous acquisition of all the data, and thus contributes a major time savings. Also, some structures may be difficult to test with ultrasound. In such cases the CAD model makes it easier to carry out a limited-angle CT inspection.

6.2.3 More Sophisticated Constraints

The proposed limited-angle CT method relies primarily on the data for solution but also uses some *a priori* assumptions. Solution by R/CG assumes a smooth solution (with dubious success) and POCS assumes fixed amplitude limits for the solution. The numerical methods examined require either a convex objective function (for CG) or convex constraint sets (for POCS). However, there are other legitimate assumptions that merit consideration and do not fit directly into CG or POCS. Two assumptions about graphite/epoxy and aluminum honeycomb sandwiches not considered in the previous chapters are:

1. a finite set of materials in the specimen, and
2. vertical continuity in the honeycomb.

Assumption one is reasonable because one usually knows exactly what materials are possible in the specimen. In the case of the plexiglass phantom, only plexiglass and air are possible while in a graphite/epoxy and aluminum honeycomb sandwich, only graphite/epoxy, aluminum, and air are possible. One can also allow for the existence of materials in typical anomalies such as water or oxidization products. A finite set of

possible materials leads to non-convex constraints in both CG and POCS. Assumption two is reasonable because, with the exception of anomalies, a material in the interior has identical material above and below it. Vertical continuity leads to convex constraints, and is not special, but it has the potential to improve reconstruction for thin honeycomb by reflecting known honeycomb properties in the constraints. Here we focus on the more difficult problem of implementation of the non-convex constraints.

In the context of R/CG, constraining the solution to consist of a finite set of materials requires that the objective function contain local minima at positions corresponding to allowed materials. With many local minima, the objective function is non-convex and CG is not guaranteed to find the global minimum, or even a close approximation. One method for minimizing such non-convex objective functions is the graduated non-convexity algorithm (GNC) [6] described as follows:

- (1) create a new objective function with a parameter, say α , such that as $\alpha \rightarrow \infty$ the function is convex, and as $\alpha \rightarrow 0$ the function becomes the original non-convex function,
- (2) set α large and solve the effectively convex minimization,
- (3) reduce α and solve the locally convex minimization at the previous minimum,
- (4) repeat step (3) until α is sufficiently small.

GNC performs a non-convex minimization by successively minimizing the locally convex objective function for decreasing values of α .

Realizing that GNC only ever needs to minimize a locally convex objective function suggests a simplification. Restrict the problem to minimizing the following quadratic objective function:

$$E = \left\| \begin{bmatrix} R \\ \frac{1}{\alpha} W \\ \alpha D_1 \\ \alpha D_2 \end{bmatrix} \mathbf{x} - \begin{bmatrix} \mathbf{y} \\ \frac{1}{\alpha} W \mathbf{x}_F \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \right\|^2$$

where all variables and parameters are as described in Chapter 4. The following procedure performs a minimization akin to GNC but only solves linear systems:

- (1) minimize E with α large, i.e., biased heavily towards a smooth solution and ignoring \mathbf{x}_F ,
- (2) if a pixel in the solution is near one of the allowed material values, assume that the pixel is that material and modify W and \mathbf{x}_F accordingly,
- (3) reduce α and repeat the minimization,
- (4) repeat steps (2) and (3) until α is sufficiently small.

The procedure starts by finding a very smooth solution. As the smoothness constraint relaxes, the solution migrates towards one that contains only allowed values.

Consideration of only one material value for each pixel avoids non-convexity, leads to a fast linear algorithm, and arrives at essentially the same solution as GNC. Figure 6.1 shows a reconstruction of the plexiglass phantom using the above method assuming that the solution contains only air ($\mu = 0\text{cm}^{-1}$) and plexiglass ($\mu = 0.45\text{cm}^{-1}$). The reconstruction is the result of the eleven quadratic optimizations summarized in Table 6.1. No attempt is made to find an optimal schedule for reducing α .



Figure 6.1: Reconstruction from successive R/CG iterations incorporating the assumption that pixels are either air ($\mu = 0\text{cm}^{-1}$) or plexiglass ($\mu = 0.45\text{cm}^{-1}$).

Optimization Number	α^2	$\frac{1}{\alpha^2}$	Number of Iterations
1	8	1/8	57
2	4	1/4	42
3	2	1/2	27
4	1	1	12
5	1/2	2	10
6	1/4	4	13
7	1/8	8	29
8	1/16	16	65
9	1/32	32	111
10	1/64	64	74
11	1/128	128	5

Table 6.1: Summary of successive minimizations used to produce reconstruction in Figure 6.1.

POCS can also solve the non-convex problem. Define the constraint set C_{M_i} as:

$$C_{M_i} = \{ \mathbf{x} : |x_i - x_m| \leq \varepsilon_M, x_m \in \mathcal{M} \},$$

where \mathcal{M} is the set of allowed values in the solution. C_{M_i} is a potentially non-convex set consisting of regions about solutions containing allowed values, with the parameter ε_M determining the size of the regions. If ε_M is large the regions overlap and C_{M_i} is convex, but if ε_M is small the regions do not overlap and C_{M_i} is non-convex. The corresponding projection operator is:

$$P_{M_i} \mathbf{z} = \begin{cases} \mathbf{z} : z_i = x_m + \varepsilon_M & z_i > x_m + \varepsilon_M, x_m \text{ minimizes } |z_i - x|_{x \in \mathcal{M}} \\ \mathbf{z} : z_i = x_m - \varepsilon_M & z_i < x_m - \varepsilon_M, x_m \text{ minimizes } |z_i - x|_{x \in \mathcal{M}} \\ \mathbf{z} & \text{otherwise} \end{cases}$$

POCS finds a solution in a manner analogous to GNC by varying ε_M as follows:

- (1) start with ε_M large and perform the reconstruction starting from $\mathbf{x} = \mathbf{0}$,
- (2) reduce ε_M and repeat the reconstruction starting from the previous solution,
- (3) repeat step (2) until ε_M is sufficiently small.

The method works because the constraint set C_{M_i} is always locally convex, i.e., once the projection operator has determined which material a pixel matches best, it considers only that material. As ε_M diminishes it may be necessary to increase the parameter ε_R to ensure that the intersection of constraints is non-empty. Figure 6.2 shows a

reconstruction based on the above method using $\varepsilon_R = 0.1$, $\varepsilon_F = 0.1$, and

$\varepsilon_M = 1, \frac{1}{2}, \frac{1}{4}, \dots, \frac{1}{512}$, stopping when $\|x_i - x_{i-1}\| < 0.1$. Note that with the assumption of a

finite set of materials the amplitude constraint is redundant and therefore it is omitted.

Table 6.2 summarizes the iterations in the reconstruction.



Figure 6.2: Reconstruction from successive POCS iterations incorporating the assumption that pixels are either air ($\mu = 0\text{cm}^{-1}$) or plexiglass ($\mu = 0.45\text{cm}^{-1}$).

POCS Solution Number	ε_M	Number of Iterations
1	1	1
2	1/2	1
3	1/4	1
4	1/8	4
5	1/16	9
6	1/32	9
7	1/64	4
8	1/128	2
9	1/256	2
10	1/512	2

Table 6.2: Summary of successive POCS applications used to produce reconstruction in Figure 6.2.

In summary, it is possible to incorporate non-convex constraints in the reconstruction process. Figures 6.1 and 6.2 show the results of using such constraints with R/CG and POCS respectively. It is not clear from this evidence that assuming a finite set of possible materials improves the accuracy of reconstruction. More accurate measurements using superior apparatus may be necessary to gain a real benefit from such constraints.

6.2.4 Compton Scatter and Ultrasound Fusion

This section speculates about the potential for application of data fusion to Compton back scatter imaging and ultrasound. Figure 6.3 shows schematically a Compton back scatter imaging system. An x-ray source projects a collimated beam into the specimen. As the specimen absorbs the x-rays, it emits scattered radiation along the beam. Compton effect scattering radiates in all directions, including back towards the source. An x-ray pin-hole camera directs radiation from varying depths in the specimen to an array of sensors producing intensity signals related to the amount of scatter originating at each depth. Philips COMSCAN, a commercial version of the apparatus, scans the specimen over two dimensions to produce a volumetric image. The method potentially offers high contrast images of anomalies not visible to conventional radiography*, and requires access to only one side of the specimen. See Chapter 2 for a description of ultrasound.

At present, there are no algorithms for solving the inverse problem posed by back scatter imaging. Instead, commercial devices, such as COMSCAN, simply scale the raw intensity data to produce an image. The inverse problem is difficult for two reasons:

1. linear attenuation varies greatly depending on position in the path from source to sensor because:
 - a. the spectral content of the incident radiation is different from the scattered radiation, and
 - b. beam hardening changes the spectral content of both incident and scattered radiation,and
2. radiation intensity diminishes with depth into the specimen, so noise and sensitivity increase.

* Theoretically, the contrast for a void in a homogenous specimen is infinite, but in practice spatial resolution limits contrast.

A possible approach to solving the inverse problem is to reconstruct the specimen one layer at a time, starting from the top and working down. The top layer is easily reconstructed because there is only air above it. Subsequent layers depend only on attenuation values for layers above so a downward step-wise procedure can work.




Figure 5.3: Apparatus for Compton back scatter imaging.

For ideal lossless media, ultrasound inversion is tractable, but in reality, materials are not lossless and ultrasound inversion is difficult. Losses confound ultrasound inversion in two ways:

1. losses are inconsistent and unknown, and
2. losses cause signal strength to diminish with depth and reduce accuracy.

For example, the effects of losses can make it difficult to distinguish between the normal reflection from an interface deep in the specimen and the reflection off of an air gap in the interface (an anomaly).

Both back scatter and ultrasound inversion can operate independently. However, they are both prone to errors with increasing depth in the specimen. The two methods do

not complement each other as in the proposed CT method, but their fusion may still be beneficial. For example, a strongly-coupled recurrent fusion system can combine the layer-by-layer back scatter reconstruction with a similar ultrasound algorithm. As the reconstruction proceeds downwards, the algorithm must decide what material (or what interface in the ultrasound model) to put in the reconstruction. When both back scatter and ultrasound data concur the algorithm proceeds in a straight-forward manner. When one has ambiguous data the other can suggest the correct interpretation. For example, when ultrasound cannot determine if an interface has an air gap in it, it can assume there is no air gap if the back scatter data do not indicate a drop in scattered radiation. Likewise, a drop in scattered radiation deep in the specimen may be due to a small void or just a random variation in the data. If it is a void, however, there should be a corresponding reflection in the ultrasound data.

It is not certain that such a fusion system will lead to a superior inspection technique. Nevertheless, if the system works, it will constitute an important contribution to NDE.

Bibliography

- [1] Artzy, E, Elfving, T, and Herman, G.T., "Quadratic Optimization for Image Reconstruction, II," *Computer Graphics and Image Processing*, vol. 11, pp 242-261, 1979.
- [2] Axelsson, O. and Barker, V.A., **Finite Element Solution of Boundary Value Problems Theory and Computation**, Academic Press, Orlando, 1984.
- [3] Bates, D.M. and Wahba, G., "Computational Methods for Generalized Cross-Validation with Large Data Sets," **Treatment of Integral Equations by Numerical Methods**, C.T.H. Baker and G.F. Miller editors, Academic Press, London, 1982.
- [4] Bauschke, H.H. and Borwein, J.M., "On Projection Algorithms for Solving Convex Feasibility Problems," Simon Fraser University, Department of Mathematics and Statics Research Report No. 93-12, June, 1993.
- [5] Bertero, M., Poggio, T.A. and Torre, V., "Ill-Posed Problems in Early Vision," *Proceedings of the IEEE*, Vol. 76, No. 8, pp 869-889, August 1988.
- [6] Blake, A., and Zisserman, A., **Visual Reconstruction**, MIT Press, Cambridge, MA, 1987.
- [7] Brègman, L.M., "The Method of Successive Projection for Finding a Common Point of Convex Sets," translated by J.G. Ceder, *Doklady*, Vol. 162, No. 3, pp 688-692, 1965.
- [8] Censor, Y. and Herman, G.T., "Row-Generation Methods for Feasibility and Optimization Problems Involving Sparse Matrices and Their Applications," *Sparse Matrix Proceedings 1978*, SIAM, Philadelphia, pp 197-219, 1978.
- [9] Clark, J.J. and Yuille, A.L., **Data Fusion for Sensory Information Processing Systems**, Kluwer Academic Publishers, Boston, 1990.

- [10] Davison, M.E., "The Ill-Conditioned Nature of the Limited Angle Tomography Problem," *SIAM Journal of Applied Mathematics*, Vol. 43, No. 2, pp 428-448, 1983.
- [11] Dusaussoy, N.J. and Abdou, I.E., "The Extended MENT Algorithm: A Maximum Entropy Type Algorithm Using Prior Knowledge for Computerized Tomography," *IEEE Transactions on Signal Processing*, Vol. 39, No. 5, pp 1164-1180, 1991.
- [12] Forsythe, G.E., Malcolm, M.A. and Moler, C.B., **Computer Methods for Mathematical Computations**, Prentice-Hall, Englewood Cliffs, N.J., 1977.
- [13] Golub, G.H. and Van Loan, C.F., **Matrix Computations**, The Johns Hopkins University Press, Baltimore, Maryland, 1983.
- [14] Golub, G.H., Heath, M., and Wahba, G., "Generalized Cross Validation as a Method of Choosing a Good Ridge Parameter," *Technometrics*, Vol. 21, pp 215-213, 1979.
- [15] Goupillaud, P.L., "An Approach to Inverse Filtering of Near-Surface Layer Effects from Seismic Records," *Geophysics*, Vol. 26, No. 6, pp 754-760, 1961.
- [16] Grünbaum, F.A., "A Study of Fourier Space Methods for Limited Angle Image Reconstruction," *Numer. Funct. Anal. and Optimiz.*, Vol. 2, No. 1, pp 31-42, 1980.
- [17] Gubin, L.G., Polyak, B.T., and Raik, E.V., "The Method of Projections for Finding the Common Point of Convex Sets." *U.S.S.R Computational Mathematics and Mathematical Physics*, Vol. 7, No. 6, pp 1-24, 1967.
- [18] Habibi-Ashrafi, F. and Mendel, J.M., "Estimation of Parameters in Lossless Layered Media Systems," *IEEE Transactions on Automatic Control*, Vol. AC-27, No. 1, pp 31-48, 1982.
- [19] Halmshaw, R., **Industrial Radiology: Theory and Practice**, Applied Science Publishers, London, 1982.
- [20] Heiskanen, K.A., Rhim, H.C. and Monteiro, J.M., "Computer Simulations of Limited Angle Tomography of Reinforced Concrete," *Cement and Concrete Research*, Vol. 21, pp 625-634, 1991.
- [21] Herman, G.T. and Lent, A., "Quadratic Optimization for Image Reconstruction I," *Computer Graphics and Image Processing*, Vol 5, pp 319-332, 1976.
- [22] Herman, G.T., **Image Reconstruction from Projections: The Fundamentals of Computerized Tomography**, Academic Press, New York, 1980.
- [23] Louis, A.K., "Incomplete Data Problems in X-Ray Computerized Tomography," *Numerische Mathematik*, Vol. 48, pp 251-262, 1986.
- [24] McRae, K.I., "Deconvolution Techniques for Ultrasonic Imaging of Adhesive Joints", *Materials Evaluation*, pp 1380-1384, 1990.

- [25] Mendel, J.M. and Habibi-Ashrafi, F., "A Survey of Approaches to Solving Inverse Problems for Lossless Layered Media Systems", *IEEE Transaction on Geoscience and Remote Sensing*, Vol. GE-18, No. 4, pp 320-330, 1980.
- [26] Mendel, J.M. and Goutsias, J., "One-Dimensional Normal-Incidence Inversion: A Solution Procedure for Band-Limited and Noisy Data," *Proceedings of the IEEE*, Vol. 74, No. 3, pp 401-414, 1986.
- [27] Oskoui-Fard, P, and Stark, H., "Tomographic Image Reconstruction Using the Theory of Convex Projections," *IEEE Transactions on Medical Imaging*, Vol. 7, No. 1, March 1988.
- [28] Oskoui-Fard, P, and Stark, H., "A Comparative Study of Three Reconstruction Methods for a Limited-View Computer Tomography Problem," *IEEE Transactions on Medical Imaging*, Vol. 8, No. 1, March 1989.
- [29] Pierre, D.A., **Optimization Theory with Applications**, Dover Publications, New York, 1986.
- [30] Poggio, T., Torre, V. and Koch, C., "Computational Vision and Regularization Theory," *Nature*, Vol. 26, pp 314-319, September 1985.
- [31] Press, W.H., Flannery, B.P., Teukolsky, S.A. and Vetterling, W.T., **Numerical Recipes in C the Art of Scientific Computing**, Cambridge University Press, Cambridge, 1990.
- [32] Radon, J., "Uber die Bestimmung von Funktionen durch ihre Integralwerte langs gewisser Mannigfaltigkeiten," *Ber. Verb. Saechs. Akad. Wiss., Leipzig, Math. Phys. Kl.*, Vol. 69, pp 262-277, 1917.
- [33] Sato, T., Norton, S.J., Linzer, M., Ikeda, O. and Hiram, M., "Tomographic Image Reconstruction from Limited Projections using Iterative Revisions in Image and Transform Spaces," *Applied Optics*, Vol. 20, No. 3, pp 395-399, 1981.
- [34] Scudder, H. J., "Introduction to Computer Aided Tomography", *Proceedings of the IEEE*, Vol. 66, No. 6, pp 628-637, 1978.
- [35] Servo-Robot, Saturn-2000 Integrated 3-D Laser Vision System Product Description, Montreal, Canada, circa 1989.
- [36] Soumekh, M., "Image Reconstruction from Limited Projections using the Angular Periodicity of the Fourier Transform of the Radon Transforms of an Object," *Acoustical Imaging*, Vol. 13, Proceedings of the Thirteenth International Symposium, Minneapolis, MN, USA, pp 31-42, October 1983.
- [37] Stoer, J. and Bulirsch, R., **Introduction to Numerical Analysis**, Springer-Verlag, New York, 1980.
- [38] Strang, G., **Linear Algebra and Its Applications**, 2nd ed., Accademic Press, 1980.
- [39] Strang, G., **Introduction to Applied Mathematics**, Wellesley-Cambridge Press, Wellesley, MA, 1986.

- [40] Tam, K.C., Eberhard, J.W. and Mitchell, K.W., "Incomplete-Data CT Image Reconstructions in Industrial Applications," *IEEE Transactions on Nuclear Science*, Vol. 37, No. 3, pp 1490-1499, 1990.
- [41] Tikhonov, A.N. and Arsenin, V.Y., **Solutions of Ill-Posed Problems**, translation editor F. John, V.H. Winston and Sons, Wahsington, D.C., 1977.
- [42] Trummer, M.R., "A Note on the ART of Relaxation," *Computing*, Vol. 33, pp 349-352, 1984.
- [43] Tuy, H., "Reconstruction of a Three-Dimensional Object from a Limited Range of Views," *Journal of Mathematical Analysis and Applications*, Vol. 80, pp 598-616, 1981.
- [44] Youla, D.C. and Webb, H., "Image Restoration by the Method of Convex Projections: Part 1 - Theory," *IEEE Transactions on Medical Imaging*, Vol. MI-1, No. 2, October 1982.
- [45] Zala, C.A. and McRae, K.I., "An Optimization Method for Acoustic Impedance Estimation of Layered Structures using Prior Knowledge", *Acoustical Imaging*, Vol. 18, pp 363-372, 1991.
- [46] Zala, C.A. and Churchill, L., "Inversion of Ultrasonic Traces," DREP Contractor's Report Series 90-18, Victoria, B.C, Canada, 1990.