



FULL WAVEFORM INVERSION OF
ULTRASONIC PHASED ARRAY DATA FOR
THE TOMOGRAPHIC RECONSTRUCTION OF
HETEROGENEOUS MEDIA FOR NDT

Russell Niven

Department of Mathematics and Statistics

University of Strathclyde

A thesis submitted for the degree of

Doctor of Philosophy

2020

Copyright Declaration

This thesis is submitted to the University of Strathclyde for the degree of Doctor of Philosophy in the Faculty of Science.

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

The copyright of this thesis belongs to the author under the terms of the United Kingdom Copyright Acts as qualified by the University of Strathclyde Regulation 3.51. Due acknowledgement must always be made of the use of any material contained in, or derived from, this thesis.

Date: 29th October 2020

Signed: Russell Niven

Acknowledgement

I would like to thank my supervisors, Anthony Mulholland and Katherine Tant, whose patience, guidance, and red pen has been essential for the completion of the PhD programme as well as my own personal and professional development. Also, Anthony Gachagan, who provided me with access and support in the field of ultrasonic engineering.

I would also like to thank my colleagues in the Department in Mathematics & Statistics for their technical and emotional support. In particular, the ladies in the office who were incredibly friendly and always great to chat to (hence the daily visit)! Lyndsay Kerr, who I shared an office and Taylor Swift Tuesday with, Kate Pyper and Johnny Love for their Statistical knowledge. And the boys (Joseph, Craig, Joe, and Feargus) that provided entertainment and emotional support. I would like to thank everyone in the Centre of Ultrasonic Engineering who helped me with all the engineering support. In particular, Richard O'Leary, David Lines, Marcus Ingram, and Matthew McInnes.

A special thank you to Ainsley Miller, whose friendship and support was essential for the completion of this thesis, in particular the support with the rj-MCMC and formatting of this thesis. For also introducing me to Cara Richardson whose kind words and support gave me the strength to complete this thesis.

Abstract

Within the ultrasonic non-destructive community the benchmark imaging algorithm assumes that the host material is homogeneous. This can lead to poor flaw detection, location, and characterisation. However, when the heterogeneous nature is accounted for within an imaging algorithm there is a significant improvement in the flaw images.

This work builds upon that work by creating a full waveform inversion framework to build a tomography method that reconstructs spatially heterogeneous wave speeds maps from simulated ultrasonic phased array data. This framework consists of: a Voronoi tessellation to spatially parametrise the wave speed map; a bespoke semi-analytical model that encompasses the data within an ultrasonic A-scan while being computationally efficient; a bespoke objective function to quantify how well the semi-analytical model compares to the observed data; and a Bayesian framework, namely the reversible jump Marko chain Monte Carlo method (rj-MCMC), to perform the tomographic reconstruction in the form of a posterior distribution. The reconstructed wave speed maps are then used in con-

junction with an imaging algorithm (TFM), to provide enhanced flaw imaging. The quality of the imaged flaws are then quantified by calculating the signal-to-noise ratio, the flaw location error, and probability of detection via ROC curves.

This framework is first applied to a synthetic randomly heterogeneous material with a side drilled hole present. It is then applied to a layered medium that contains three inclusions contained in different layers. The latter case being studied extensively in different scenarios and phased array set-ups. A new objective function is then proposed (a time windowed Hilbert transform) which is applied to the layered medium allowing for the comparison of different objective functions.

Contents

Contents	vi
1 Background	1
1.1 Background	1
1.2 Non-Destructive Testing	2
1.2.1 Ultrasonic NDT	3
1.2.2 Flaw Detection, Positioning, and Characterisation	5
1.2.3 Material Reconstruction	8
1.3 Inverse Problems	11
1.3.1 Optimisation Techniques	12
1.3.2 Bayesian Inverse Problems	13
1.3.3 Full Waveform Inversion	15
1.3.4 Material Parametrisation	18
1.4 Outline of Thesis	19
2 Full Waveform Inversion Tomography	22
2.1 Inversion Scheme	23
2.1.1 Observed Data	23

CONTENTS

2.1.2	Material Parametrisation	23
2.1.3	Objective Function	24
2.1.4	Data Uncertainty Parametrisation	26
2.1.5	A Bayesian Framework	26
2.1.6	The rj-MCMC Method	29
2.1.7	Sampling the Posterior Distribution	37
2.1.8	Imaging a Flaw	38
2.1.9	Quality Metrics	40
2.2	Waves in Layered Medium	42
2.3	Discussion	47
3	Development of the Forward Model	48
3.1	Forward Model	48
3.1.1	Deriving the amplitudes reflected wave and transmitted waves	49
3.1.2	Approximation of the Reflected & Transmitted Wave . . .	54
3.1.3	Pseudo-code for Approximating the Reflection & Transmitted Wave	59
3.2	Discussion	66
4	Tomographic Reconstruction of a Randomly Heterogeneous Material	68
4.1	Randomly Locally Isotropic Heterogeneous Material	69
4.2	Improving Flaw Detection using FWI rj-MCMC	73
4.3	Comparing Full Waveform Inversion to Time of Flight	78
4.4	Posterior Distribution vs a Randomly Generated Heterogeneous Map	81

CONTENTS

4.5	Discussion	83
5	Tomographic Reconstruction of a Layered Medium	85
5.1	TFM of Layered Medium	87
5.2	Setting up and Initiating the rj-MCMC Algorithm	93
5.3	Results using the PCC Objective Function	96
5.4	Reconstruction of a Layered Medium using 32 rays from the Full Matrix Capture	97
5.4.1	Reconstruction using a Pulse-Echo set up	99
5.4.2	Reconstruction using a Pitch-Catch set up - Case A	100
5.4.3	Reconstruction using a Pitch-Catch set up Without the Noise Variable - Case B	103
5.4.4	TFM+ from the Reconstructed Material Map Using 32 Rays	116
5.5	Reconstruction of a Layered Medium using 1024 rays from the Full Matrix Capture	119
5.5.1	TFM+ from the Reconstructed Material Map Using 1024 Rays	128
5.6	Reconstruction of a Layered media with the Addition of Noise	131
5.7	Discussion	135
6	Tomographic Reconstruction of a Layered Medium from Finite Element Simulated Data	137
6.1	Adjusting to Simulated Data	137
6.1.1	Adjusting to Simulated Data - Planar Wave	138
6.1.2	Adjusting to Simulated Data - Spherical Wave	143

CONTENTS

6.2	Reconstruction of a Layered Medium from Simulated Data using the Log Based Objective Function	152
6.2.1	TFM+ from the Reconstructed Material Map from Simulated Data	156
6.3	Discussion	159
7	Tomographic Reconstruction using a Time Windowed Objective Function	161
7.1	The Hilbert Transform	164
7.2	Implementing the Hilbert Transform Envelope	168
7.3	Tomographic Reconstruction of a Layered Medium using a Hilbert Transform	173
7.3.1	Reconstruction over 300,000 samples	173
7.3.2	TFM+ from the Reconstructed Map using the Hilbert Transform Time Windowed Objective Function	182
7.3.3	Reconstruction over 600,000 samples	187
7.4	Tomographic Reconstruction of a Layered Medium using the Hilbert Transform Time Windowed Objective Function on Finite Element Simulated Data	192
7.4.1	TFM+ from the Reconstructed Material Maps from Simulated data	196
7.5	Discussion	199
8	Conclusion and Future Work	201
8.1	Conclusion	201

CONTENTS

8.2 Future Work	204
References	206

Chapter 1

Background

1.1 Background

The work done in this thesis aims to build a framework, based on a Full Waveform Inversion (FWI) method which uses ultrasonic phased array data, to obtain a wave speed map that approximates the heterogeneous microstructure found in material of interest in non-destructive testing. With the approximation of the wave speed map being used in conjunction with an imaging algorithm, it is hoped that flaw detection and characterisation will be improved. This work is an inverse problem within a non-destructive testing (NDT) context. This chapter provides a review of current literature on NDT, flaw detection and characterisation methods, inverse problems, and full waveform inversion.

1.2 Non-Destructive Testing

Non-destructive testing (NDT) is a general term for a wide range of techniques industry use to obtain information of structural integrity of an object, without impairing its usefulness [1, 2]. It is also commonly known as non-destructive evaluation (NDE) or non-destructive inspection (NDI). The most common NDT techniques are visual inspection [3], radiography [4], liquid penetrant inspection [5], eddy current [6], and ultrasonic testing [7], all of which have their strengths and weaknesses. For example, while visual inspections are low cost they are limited to finding defects that break the surface, or those which create surface roughness and texture changes [1]. Visual inspection can be aided with liquid penetrants that work by penetrating the crack to increase its visibility. Eddy currents and other magnetic methods operate on the principle of detecting distortion of magnetic fields. In an NDE context, surface breaking and subsurface flaws generate local magnetic flux leakage which can be detected to infer their location and size in a component [1]. Radiography relies on the generation of electromagnetic radiating waves, commonly known as X-rays, to internally inspect a component [3]. Although highly successful, it is an expensive form of NDE due to required facilities and safety regulations to protect the operators [8]. Ultrasonic testing is one of the most popular NDT techniques [2] as it is relatively cheap, portable, offers potential for real-time and in-situ measurements, and is not restricted to surface defects [9]. This is the technique that forms the focus of this thesis.

1.2.1 Ultrasonic NDT

Ultrasound is the branch of acoustics which considers sound waves with frequencies which lie above the audible limit for humans (about 20kHz [10, 11]) and it has important applications within the fields of physics [12, 13], industrial technology [7], medicine [14, 15], and bioacoustics [16, 17]. The application of ultrasound to non-destructive testing was made possible in 1880 by the discovery of the piezoelectric effect by Pierre and Jacques Curie [18]. Then during World War II, after the discovery of radar, there was significant progress in imaging systems and led to an increase in ultrasonic research [1]. It was then that the pulse-echo ultrasonic flaw detector was developed, independently, by Firestone in the USA [19] and Sproule [20] in the UK. Ultrasonic NDT is now currently the most widely used method of NDT, excluding the simple tests of visual inspection and weighing [21].

Ultrasonic methods have their advantages in that they can accurately position, size and shape defects, using relatively inexpensive technology. They can also be used to estimate material properties. The main disadvantage of ultrasound is that complex shapes, such as elbows, nozzles, rough profile welds, create complex signals which are difficult to process as these shapes create beam profiles that have distortions [22]. Materials with highly scattering properties are also difficult to test, for example, within an NDT context, inspection of austenitic welds and composite laminates pose an ongoing challenge [23, 24].

There are several ways that sound waves can propagate in solid media and these can be differentiated by the direction of particle vibration. These waves include longitudinal, transverse or shear, Rayleigh, Lamb, and Love waves [11]. In the case of longitudinal waves (also known as compressional or dilational waves)[11, 25], the particles move in the direction of the wave propagation. In the case of shear waves, the particles vibrate in the perpendicular direction to the wave propagation direction, and are therefore only present in solid media [11, 26]. Rayleigh, or surface waves, are waves that travel along the surface of a relatively thick medium propagating to about the depth of one wavelength [11, 27]. Lamb and Love waves are different forms of plate waves which propagate within materials that are a few wavelengths in thickness. In the case of Lamb waves the particles vibrate in the direction of wave propagation [11, 28], and in the case of Love waves the particles vibrate in the plane perpendicular to the direction of the wave [11, 29]. It is possible for a wave of one type to be transformed into another when it interacts with an interface or inhomogeneity and this is known as mode conversion [11]. One key example of this is when a longitudinal wave hits an interface at an angle, which is not normal to the interface, then some of the energy will be transformed into creating a shear wave.

The main way of producing ultrasonic waves for NDT is with an ultrasonic transducer, which is a device that converts electrical energy into mechanical energy and vice versa. Typically these will be piezoelectric or capacitive transducers however there is some recent developments in the area of laser induced ultrasonics [30, 31]. Piezoelectric transducers utilise the piezoelectric effect where an elec-

trical charge is generated by changing the shape of certain crystals by means of a mechanical force, and conversely applying an electrical charge to the crystal changes its shape. Capacitive transducers use electrostatic fields between a conductible diaphragm and a backing plate. The use of a driving AC current then causes the diaphragm to oscillate in transmission mode [10]. Laser transduction works by generating ultrasonic waves using a pulsed laser, which induces thermal stress and strain in the solid material which in turn generates elastic waves. These can then be detected by a laser interferometer [32].

The production and implementation of phased array transducers surged in the early 2000's [7]. These multi-element transducers allow for a range of different dynamic inspections from a single location, such as steering and/or focusing of ultrasonic beams. A phased array with N elements (typically between 64 and 256) can generate (from a single fixed position) N^2 time traces arising from each transmit/receive pair of elements; this data acquisition strategy is called Full Matrix Capture [9].

1.2.2 Flaw Detection, Positioning, and Characterisation

One of the main objectives of NDT is to detect interior flaws and remove affected components from service before their failure. These flaw detection methods can use time-domain, frequency, or time and frequency information and although ultrasonic NDT data is typically collected in the time domain, the Fourier transform facilitates examination across both domains [33, 34].

Within ultrasonics two of the primary visualisation techniques are A-scans, and B-scans [35]. A-scans (amplitude scans) display the amplitude (energy) of a one-dimensional ultrasonic signal against time. The B-scan (brightness scan) is a cross-section of the test material through which individual A-scans are collected. There are various types of B-Scans such as plane B-scans, focused B-scans, and sector B-scans [9].

With an ultrasonic array, it is possible to gather a complete set of A-scans from all combinations of transmit-receive pairs, an approach called Full Matrix Capture (FMC). This data can be utilised in an offline post-processing technique called the Total Focusing Method (TFM) (also known as delay-and-sum beamforming [36] or Kirchoff migration [37]). It is the gold standard in delay-and-sum beamforming imaging. It can provide much better images than other techniques such as SAFT, CSM, and time reversal MUSIC, which will be discussed later. Compared to these techniques TFM can image closely spaced scatterers and is less susceptible to interference [38, 39]. It works by discretising the region of interest into a grid and each point is then given an intensity value calculated by summing each signal from all of the elements in the array [9]. It is known as the gold standard of NDT imaging and, with modern computational architectures, can produce real time results [40]. One of the central assumptions of the standard TFM algorithm which is that the region of inspection is homogeneous. In cases where the component has a complex microstructure this can lead to errors in the calculation of delay laws on which the method is based. This has been improved in [41] which takes into account the variations in wave speed within the region of

inspection in an algorithm known as TFM+. The work presented in [24] modifies the standard TFM to take into account anisotropic materials, increasing its applications to composite materials; an increasingly used material in aerospace and marine structures. Another adaptation is the halfskip TFM (HSTFM) which was developed by [42, 43] and uses the ray paths that reflect off the back wall to enhance the reflectivity amplitude of any embedded flaws. This improved the measurement of the depth of cracks compared to TFM.

The Synthetic Aperture Focusing Technique (SAFT) is another method for detection and characterisation of defects embedded in the bulk of the component [44]. The method also combines pulse-echo measurements made at multiple transmitter/receiver locations [45] but creates this synthetically using a single transmit-receive pair of sensors. Early versions of SAFT used delay-and-sum beamforming which was then followed by frequency based SAFT (FD-SAFT) [46]. In recent years FD-SAFT has been expanded on research for applications in layered media [47]. Another modern ultrasonic imaging technique is Planewave Imaging (PWI) which was introduced at the beginning of the decade (2010) [48]. This transmits planewaves by exciting a number of array elements simultaneously and this can increase the frame-rate and decrease the scanning time in comparison to FMC TFM [49]. This technique is also used in medical ultrasound [50, 48, 51].

Time reversal (TR) is another imaging method that was proposed in acoustics [52, 53]. A signal is recorded by a transducer and then re-transmitted back into the medium using a ‘last in first out’ strategy to focus on the signal source (which could be an embedded scatterer). It has applications in acoustic imaging,

nondestructive testing, medicine, and communication. The work done in [28] uses Time-Reversal of Lamb waves to classify defects. Another TR technique is TR imaging with multiple signal classification (TR-MUSIC) which has the potential to locate multiple defects with scatterers between them [54].

A final class of defect imaging algorithms are the Sampling methods which are numerical methods used to solve the inverse scattering problem to accurately locate and characterise defects [55]. These include the linear sampling method [55] which has the drawback of requiring a large amount of data. This issue can be solved with the Multi-frequency Sampling Method, the Time Domain Sampling Method (TDSM) [55], the Factorisation Method (FM), orthogonal sampling method [56], or the direct sampling method [57]. These qualitative inverse scattering methods are non-iterative and have been shown to achieve super resolution in ultrasonic phased array imaging applications (able to detect and characterise defects smaller than the wavelength) [43].

1.2.3 Material Reconstruction

In recent years, much effort has been devoted by the ultrasonic NDT community to develop capability in characterising the spatially varying material properties of a component [41]. This can be done destructively on a small scale using a method called electron backscatter diffraction (EBSD) which is a technique that allows for the characterisation of sub-micrometer grain structure [58]. This technique uses diffraction patterns acquired by a scanning electron microscope.

Current research in this area focuses on deployment of EBSD in situ to gain insight into characterisation of the micro-structure and the orientation of the crystalline structure [59, 60, 61]. It is also used within additive manufacturing (AM) frameworks to assess changes in the micro-structure of materials during the manufacturing stage [62, 63]

An alternative to EBSD is spatially resolved acoustic spectroscopy (SRAS), a non-contact laser ultrasonic technique, which utilises surface acoustic waves (SAW) generated by a laser pulse to map the grain structure of a material [64, 65, 66, 67, 68]. Similarly to EBSD, SRAS has been used to image material texture during additive manufacturing processes.

For a non-destructive method of material reconstruction there are various methods. In the weld specific case, MINA (modelling of anisotropy based on notebook of arcwelding) uses information from the welding procedure such as dimensions of the weld, number, inclination, order of weld passes, and the remelting rates to create a model of the welds anisotropy [69]. This model has facilitated numerical studies of how the ultrasound wave travels through the weld and this can be use to improve the ultrasonic testing of the weld [70, 71, 72].

In recent years stochastic methods have been employed to solve the inverse problem of reconstructing the spatially varying material properties of a component from non-destructive ultrasonic testing data. In the work presented in [73] a Markov chain Monte Carlo (MCMC) approach is used to extract an approximation of the locally anisotropic grain structure of a weld from time of flight data.

Using Dijkstra’s algorithm as the forward model, and initialising the algorithm from an approximation of the weld structure, an approximation of the weld’s grain map is obtained. This information is then used in conjunction with the TFM imaging algorithm and an improvement of 2mm in flaw location is achieved. An alternative stochastic approach is taken in [41, 74] where the reversible-jump Markov chain Monte Carlo (rj-MCMC) method is used to estimate the spatially varying material properties in a component. Unlike [73] this work uses no prior information on the distribution of material properties throughout the domain. However, these time of flight tomography approaches are based on ray theory which only estimates wave propagation in the high frequency regime, and so the resolution of the reconstructions is limited to approximately the wavelength. Furthermore, in the case where full aperture inspection is not possible (as is typical in NDT applications), these approaches suffer from non-uniqueness, even when the underlying geometry is simple (for example, layered isotropic media).

One way to overcome these limitations is to exploit more of the acquired data using a full waveform inversion approach. Full waveform Inversion (FWI) is an approach which aims to match the full reflected waveform observed in the phased array inspection with those generated using a model through some approximation of the material map. It forms the basis for the work presented in this thesis and is discussed further in Subsection 1.3.3.

1.3 Inverse Problems

Inverse problems describe the challenge of working backwards from some observed data to determine what gave rise to it. Since inverse problems are data by nature, they often require the minimisation of the difference between the observed data and the mathematically modelled data over the parameter space. Unlike in mathematical modelling (forward problems), where for a given set of inputs one can predict the behaviour of some system, inverse problems are usually ill posed; that is their solution is very sensitive to small changes in the observed data and may not be unique [75]. Strategies for coping with these difficulties include regularisation [76] and use of global optimisation schemes to avoid being trapped in local minima [77]. Inverse problems appear frequently in engineering, science and data analysis applications. For example they are fundamental in medical imaging [14, 20], seismology [78, 79], astronomy [80], non-destructive testing [32] and security [81].

In the context of this thesis, we are concerned with the inverse problem of reconstructing spatially varying material properties of solid objects from observed scattered wave data. [78, 82, 83, 84, 85, 86, 87].

1.3.1 Optimisation Techniques

A key component of many inversion frameworks is optimisation. The optimisation attempts to find the global minimum (or maxima) within an objective function landscape often contains many local minima and these can be numerous; particularly in the situation where only partial data is available as found in the limited aperture case. It can then arise that some of these minima generate identical values for the objective function and hence a number of material configurations arise that satisfy the problem constraints. This non-uniqueness of the problem can be alleviated somewhat by the use of suitable priors which constrain the search space to increase the probability of the correct material map reconstruction being achieved. The purpose of the objective function (also known as a misfit or cost function) is to quantify the difference between the modelled predicted and experimental data being minimised [88, 89, 90]. Typically, the inversion scheme will be initialised with some set of parameters selected subject to any a priori information and constraints. This objective function will be evaluated for this initial model and, dependent on the nature of the optimisation method (whether it be stochastic or deterministic), the model will be perturbed. The optimisation method will dictate how the perturbations are made and how the model space is explored, and will continue until some stopping criteria is met [89, 90].

Optimisation methods can be separated into two different categories: deterministic or stochastic [91]. Deterministic methods use the analytical properties of a problem to generate a sequence of points which converges to a global solu-

tion. However, these methods can fail to converge when the objective function is non-convex [92]. Stochastic methods, also called, optimisation under uncertainty, when applied in the context of a deterministic problem, describe optimisation frameworks which allow random perturbations to the model which allows more efficient exploration of the parameter space and escape from local minima [90]. Examples of stochastic methods include simulated annealing and genetic algorithms [93]. Simulated annealing is an iterative process that is analogous to the process of cooling metal or glass [89], where it is possible to accept a model that is objectively worse according to the objective function so that the algorithm does not get stuck inside a local minima/maxima [94]. Genetic algorithms are based on natural selection and survival of the fittest [95]. They work by selecting the best individuals (models) to reproduce and generate new models. These are evaluated and old models are replaced with the new ones. This is repeated until the population converges to an optimal solution [94].

1.3.2 Bayesian Inverse Problems

Formulating an inverse problem in a Bayesian framework (one where the optimisation is based on Bayes rule i.e. information is written in probabilistic terms) facilitates the study of the posterior distribution on the solution rather than a point estimate. It allows regularisation of the problem using prior information and can cope with non-linear inverse problems. Note that by a non-linear inverse problem I one whereby the objective function landscape is a non-linear function

of the material map parameters (the parameters to be reconstructed) and so often this surface contains many local minima as a result. Thus it is a convenient choice for the high dimensional, non-linear tomographic problem that is the focus of this thesis [96]. The posterior probability density function is given by Bayes' rule

$$p(A|B) = \frac{p(B|A)p(A)}{p(B)}, \quad (1.1)$$

where A is the current model parameters, B is the collected (observed) data, $p(A|B)$ is therefore the probability that the current model parameters can account for the observed data. This is equal to the likelihood $p(B|A)$ (how likely it is that the data B came from the model A), multiplied by the prior $p(A)$ (the probability distribution of the parameters A in a given iteration), over the evidence term $p(B)$ (probability of the observed data B). The evidence term $p(B)$ is an integral over all possible parameters and this is too computationally expensive to compute. The work around is to set up a Markov chain of models and then examine the ratio of subsequent versions of Equation (1.1) and in this way $p(B)$ is cancelled out and doesn't need to be calculated. The MCMC allows us to sample from a known probability distribution and allows us to numerically approximate the posterior distribution. The samples are generated with the Markov chain, a stochastic process where the current model is only dependent on the previous iteration (it has the memory-less property) [97]. The model is either accepted or rejected using the Metropolis-Hastings criteria which allows for the acceptance of inferior models similar to simulated annealing.

The MCMC method is used in inverse problems [98, 99] and in NDT to approximate the material map of weld [73]. There are other variations such as the Hamiltonian MCMC [88], and reversible jump MCMC (rj-MCMC). The rj-MCMC method allows for dimensional jumps and is used in Seismology [100], Geophysics [101], and NDT [41, 74].

1.3.3 Full Waveform Inversion

Full waveform inversion (FWI) is a method to find a high resolution model of a medium, be it a body part, subsurface of the earth, or a component. This is achieved by matching waveforms generated by a mathematical model of the underlying physics to an experimental data set [102]. The aim is to exploit much more information from the data, than say time of flight methods, and has become standard practice in seismology [79, 103, 104].

For example, the work done in [102] reconstructs acoustic 3D anisotropic field data and in a computationally efficient manner. In [105] they used plane wave FWI (a wave travelling along one spatial dimension) to reconstruct 3D spatially varying velocity map. The use of a plane wave allows for a decrease in the computational cost however this leads to cross-term artefacts (false artefacts) in the reconstructions. In [93] they inverted a 1D elastic wave using both synthetic and field data scenarios to derive the elastic characterisations (p-velocity, s-velocity, and density) of subsurface seismograms. However, their work cannot be applied in 2D or 3D due to the high computation cost and this highlights the relationship

between the number of model parameters to invert and computational time. The work done in [106] increased the convergence rate of FWI by taking into account the attenuation of the wave. The work in [107] used FWI on elastic subsurface waves however they assume a constant density due to the difficulty of reconstructing the density maps. In [88] a variational FWI approach was applied to a 2D synthetic dataset which provided accurate uncertainty estimates. They used the Stein variational gradient descent which was not as efficient as the Hamiltonian Monte Carlo method for the 2D case however they claim it can be more easily applied to 3D datasets where it will be more efficient.

Another application of FWI in the geophysics field is Ground Penetrating Radar FWI [108] which works on a much smaller spatial scale. The work done in [109, 110] used a frequency based FWI to reconstruct a horizontally layered medium. The model was very sensitive to the initial wavelet. Similarly, [111] used GPR FWI however the number of layers must be prescribed in this method.

Ultrasonic FWI has also been used in the medical field for the inspection of the breast as it is a safer alternative to X-rays which, although very successful, use ionising radiation [112]. The work done in [112] successfully inverted 3D ultrasonic data to generate images that were more accurate than standard 2D *in vivo* images. Similarly, ultrasonic FWI has been applied to the brain [113] to generate images with sub-millimetre resolution. However this required an almost full 360° array and hence was computationally expensive.

The potential of FWI in the NDT industry has only recently come to the forefront of research as it is only now with the ever increasing computational resources that it has started to become feasible. The work done in [114] uses FWI inversion in conjunction with Reverse time migration to reconstruct a composite layered medium to be able to detect a flaw on the layer interface. Another application of ultrasonic FWI in NDT concerns guided wave tomography; the work carried out in [115, 116] uses a finite difference (FD) of wave propagation forward model in the frequency domain to construct a thickness map of a plate like structure.

Another FWI work, that needs to be noted, is the work in [117, 118, 119]. This series of papers successfully reconstructs an arbitrary layered media in a computationally efficient manner. The work is limited to layered media and is an iterative procedure which produced one layer at a time. However, the image reconstructions are not robust to noise as the method is ill conditioned.

One important aspect of FWI is the choice of objective function. The simplest one is the L_2 misfit (the L_2 norm of the difference between the observed and modelled data) which is easy to implement [88]. In [120, 121] they overcome the issue of cycle skipping (where the phases match are out by more than half a wavelength) that happens with using L_2 as a misfit. Another misfit is the Wasserstein distance, a metric that captures the time-shift and amplitude variations by calculating how much the modelled data needs to distort to resemble to the experimental data. This misfit is computationally intensive and more complex than the standard L_2 . The work in [122] proposed a new misfit based

on the instantaneous phase (propagation phase of a wave front independent of amplitude) difference and envelope (amplitude) ratios by using a Hilbert transform. Their misfit reduces the non-linear behaviour of waveforms and is robust to noise. However, the instantaneous phase is prone to cycle skipping.

1.3.4 Material Parametrisation

Another important aspect of FWI and inverse problems is the method employed to parametrise the material (method to describe the geometry and physical properties). A simple and common method is to use a uniform grid as this is easy to implement [102, 123, 124, 125, 126]. However, models with uniform parametrisation have sharp discontinuities and spatial smoothing may be required. There is also a compromise between the resolution and the uncertainty [127]. An alternative grid approach is used in [73] in the context of welds, which uses a uniform grid of both isotropic and anisotropic sub regions. An alternative approach to parametrise the weld geometry is used in [71] whereby just four key parameters represent the weld geometry. This low degree of freedom approach does not allow a local description of the material properties.

An alternative approach (one which strikes a balance in terms of degrees of freedom) is to use an irregular parametrisation such as Delaunay triangulation or Voronoi tessellations [123]. A Delaunay tessellation (or triangulation in 2D) is a series of random points (nodes) which are connected to their neighbour to create an irregular triangulation [128]. The size of the triangles is dependent on

density of points; a useful property when the application you are parametrising has significant variations in length scales [129]. The dual of Delaunay tessellation is the Voronoi tessellation. This partitions the spatial domain into irregular convex polygons [129]. For the purposes of parametrising some heterogeneous material, each polygon can be assigned a material property. The Voronoi diagram is a fundamental and important geometrical construct, its construction has been optimised [130], allowing it to be used in seismic tomography [100, 127, 131, 101, 132], NDT [41, 133, 134, 135], biological structures [136], epidemiology [137], and a plethora of other computational sciences. An alternative way to partition the spatial domain is using the Johnson-Mehl tessellation [138, 139, 140] where the cells are not always convex and their boundaries can be hyperbolic [140].

1.4 Outline of Thesis

In this chapter, the concept of Non-Destructive Testing (NDT) is introduced and different methods involved within NDT. The different types of ultrasonic waves that are used within ultrasonic NDT are discussed. This is followed by a review of techniques for flaw detection and characterisation, and NDT techniques for material reconstruction tomography. A general introduction into inverse problems outside of NDT is given alongside details on optimisation techniques, Bayesian inverse problems, full waveform inversion and on material parameterisation.

Chapter 2 details the specific techniques used for the inverse problem in this thesis. This includes material parametrisation via Voronoi tessellation, definition of the objective function, and an introduction to the reversible jump Markov chain Monte Carlo (rj-MCMC) method. The mathematical modelling of waves travelling through layered medium is detailed.

Chapter 3 until the end of Chapter 7 presents the **original work** in this thesis. It provides the development of the forward model. First a detailed forward model that is too computationally expensive and then a more computationally efficient semi-analytical model used within the inversion algorithm.

Chapter 4 is the first test case for the inversion algorithm; a locally isotropic randomly heterogeneous material and the use of simulated data to drive the algorithm. The full waveform inversion tomography is then compared to time of flight (ToF) tomography.

Chapter 5 focuses on a 2D layered medium. The tomographic inversion using a partial array aperture for different scenarios is explored followed by the results of the tomographic inversion using the complete FMC dataset. The robustness of the inversion method is tested by adding artificial noise to the observed data.

Chapter 6 starts by comparing and adjusting the forward model to simulated data generated in a finite element (FE) simulation. The results of the inversion for a layered medium using a simulated data are then presented.

Chapter 7 develops a new objective function based on a time windowed Hilbert transform. The results of the tomographic inversion are then presented for observed data created from both the forward model and from finite element simulated data.

The thesis concludes with a summary of the results presented and an overview of future work.

Chapter 2

Full Waveform Inversion Tomography

In this thesis, a full waveform inversion framework is created to build a tomographic method that will reconstruct spatially heterogeneous wave speed maps from simulated ultrasonic phased array data. To solve these inverse problems, the rj-MCMC method is selected as the optimisation technique. Voronoi tessellations are used to parametrise the material. The Pearson correlation coefficient was used as an objective function (this is then adapted in Chapter 5 and a new objective function is proposed in Chapter 7). In this chapter, the components of the inversion method are defined as the observed data, the material parameterisation, the objective function, and the optimisation technique. The methods used to image a flaw and metrics to quantify the quality of the flaw image is then given. Finally, the mathematics of wave propagation is outlined.

2.1 Inversion Scheme

2.1.1 Observed Data

In this thesis, the observed data is generated in two ways: using a finite element (FE) model (in the commercial package PZFlex and later its successor OnScale [141]) and using the semi-analytical model developed in Chapter 3. In the case of the FE generated data, physics neglected in the forward model of the inversion framework (such as refraction and mode conversion) are included and so there is no need to introduce synthetic noise. However, where the forward model is used to generate the observed data, proportional and additive noise is added to check the robustness of the inversion.

2.1.2 Material Parametrisation

To limit the degrees of freedom in our inverse problem, a Voronoi tessellation is used to create a low dimensional parametrisation of the material. This starts by creating a set of seeds S where each $s_i \in S$ is a two dimensional Cartesian coordinate. Non-overlapping convex polygons then tessellate the domain where the domain of the i^{th} polygon is defined as the space that is closest to seed s_i . Each cell is assigned a material property, in this thesis this will be the longitudinal material wave speed. This gives $3L + 1$ degrees of freedom (the (x, y) -coordinates of each cell s_i , the set of wave speeds V , and the number of seeds L , which is also an unknown). This method reduces the degrees of freedom required to describe

the material map and provides a means of affecting large regions with a single perturbation. This method has been used to parametrise tomographic imaging problems [41, 100, 101].

2.1.3 Objective Function

Within the inversion algorithm a method of comparing the modelled and observed data is needed to form an objective function. The objective function used is the Pearson correlation coefficient (PCC), which measures the similarity between two waves, with the intent of looking at the phase coherence between the modelled and experimental wave signatures rather than the agreement in their amplitudes. This is an advantage when considering attenuation (which is not included in

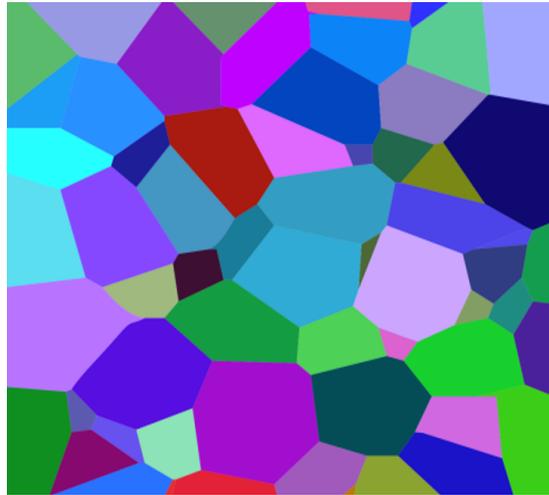


Figure 2.1: Example of a 2D partitioned using Voronoi tessellation

the semi-analytic model introduced in Chapter 3) as it compares the relative difference in height of the difference amplitudes rather than the absolute values of the amplitudes within an A-scan.

Each A-scan is stored as a vector of length n given by the number of time steps. For each transmit (i) and receive (j) ultrasonic transducer element pair the observed A-scan is denoted by $\mathcal{A}_{i,j}^E$ and the modelled A-scan by $\mathcal{A}_{i,j}^{M_p}$ (for the p^{th} material map). The PCC of the two A-scans is given by

$$\Psi(\mathcal{A}_{i,j}^E, \mathcal{A}_{i,j}^{M_p}) = \frac{\sum_{t=1}^n (x_t - \bar{x})(y_t - \bar{y})}{\sqrt{\sum_{t=1}^n (x_t - \bar{x})^2} \sqrt{\sum_{t=1}^n (y_t - \bar{y})^2}}, \quad x_t \in \mathcal{A}_{i,j}^E, \quad y_t \in \mathcal{A}_{i,j}^{M_p},$$

where \bar{x} and \bar{y} are the mean of $\mathcal{A}_{i,j}^E$ and $\mathcal{A}_{i,j}^{M_p}$, respectively [142]. The PCC gives a value between -1 and 1, where $\Psi = 1$ means there is a perfect agreement between the two signals and $\Psi = -1$ means the modelled signal is the phase cancelling counterpart to the observed signal (perfect negative correlation). Rescaling Ψ using $\hat{\Psi} = (1 - \Psi)/2$ gives $\hat{\Psi} \in [0, 1]$, with 0 being perfect agreement; this can then be used as an objective function in a minimisation problem. Comparing the A-scans for each transmit/receive pair the objective function can be written

$$\Phi(M_p) = \sum_{i,j=1}^{P^2} \hat{\Psi}(\mathcal{A}_{i,j}^E, \mathcal{A}_{i,j}^{M_p}), \quad (2.1)$$

for material model M_p as parametrised by a given Voronoi and where P is the number of transmitting elements.

This objective function is then adapted in Chapter 5 and then a new objective function is proposed and used in Chapter 7.

2.1.4 Data Uncertainty Parametrisation

Data uncertainty comes in multiple forms. First there is a numerical discrepancy; using the same forward model in two different software packages gives slightly varying results. Then there is the uncertainty arising from simplifications in the forward model (for example, neglecting to include mode conversion). Lastly there is system noise; in an experimental set up this could be white noise generated from the equipment. In this work, the data uncertainty is parameterised by the noise parameter, σ_n . This parameter will be estimated during the inversion process. This parameter should prevent the data from being overfitted or underfitted. Since σ_n is an additional parameter to be determined, the dimension of the model is $3L + 2$.

2.1.5 A Bayesian Framework

From related investigations [41, 100] it is clear that the objective function in these Bayesian inverse problems will not be convex and will contain many local minima. It is also important in the context of this problem to be able to quantify the uncertainty in any map reconstruction and subsequent flaw imaging. It is sensible therefore to use a global optimisation approach which is suited to such

complex objective functions and one which will naturally give rise to a population of material map solutions; from which statements can readily be made about the uncertainty. For these reasons a Bayesian methodology was adopted.

The reversible-jump Markov Chain Monte Carlo (rj-MCMC) method produces a posterior distribution for trans-dimensional spaces. It is an ensemble approach based in a Bayesian framework in which all information is written in terms of probabilities. The posterior probability density function is given by Bayes' rule

$$p(M_p|\mathcal{A}_{i,j}^E) \propto p(\mathcal{A}_{i,j}^E|M_p)p(M_p)$$

where $p(M_p)$ is the prior probability density function of the model M_p and $p(\mathcal{A}_{i,j}^E|M_p)$ is the likelihood that the observed data $\mathcal{A}_{i,j}^E$ arises from that model.

The difference between the model and the observed data is measured by the objective function given in equation (2.1). The likelihood is then given by

$$p(\mathcal{A}_{i,j}^E|M_p) \propto \exp(-\phi(M_p)/2) \tag{2.2}$$

where

$$\phi(M_p) = \frac{\Phi(M_p)}{\sigma_n} \tag{2.3}$$

where σ_n is the estimated variance of the data noise.

To calculate the posterior probability density function we also require information on the prior, $p(M_p)$. In this work the prior probability density functions for each model parameter are chosen to be a uniform distribution as used in

[41, 100, 101]. The model parameters are assumed to be independent of each other. Therefore, the priori probability density function can be written as a product of the probability density functions of the individual model parameters

$$p(M_p) = p(L)p(V|L)p(S|L)p(\sigma_n) \quad (2.4)$$

where $p(L)$ is the prior on the number of Voronoi cells to parametrise the underlying structure. A discrete uniform distribution defined as

$$p(L) = \begin{cases} 1/\Delta L, & \text{if } L \in \mathcal{L} \\ 0, & \text{otherwise,} \end{cases} \quad (2.5)$$

is used. Where $\mathcal{L} = [L_{\min}, L_{\min} + 1, \dots, L_{\max} - 1, L_{\max}]$, $\Delta L = L_{\max} - L_{\min} + 1$ and integer bounds $L_{\min}, L_{\max} \in \mathbb{N}$. Each Voronoi seed $s_i \in S$ will have an associated wave speed, $v_i \in V$ which have a uniform distribution given by

$$p(v_i) = \begin{cases} 1/\Delta v, & \text{if } v_{\min} \leq v_i \leq v_{\max} \\ 0, & \text{otherwise} \end{cases} \quad (2.6)$$

where v_i is measured in ms^{-1} , $\Delta v = v_{\max} - v_{\min} + 1$ and v_{\min} and v_{\max} are integer bounds. It is assumed that the wave speed of one cell is independent of the others, therefore,

$$p(V|L) = \prod_{i=1}^S p(v_i). \quad (2.7)$$

The seed positions have a uniform distribution within the bounds of the spatial domain, the L seeds will have L distinct locations, therefore,

$$p(S|L) = \left[\frac{|I_*|!}{L!(|I_*| - L)!} \right]^{-1} \quad (2.8)$$

where $|I_*|$ is the cardinality of the computational domain with which the seed co-ordinates are assigned. Lastly, the prior on the level of uncertainty present in the system is given by

$$p(\sigma_n) = \begin{cases} 1/\Delta\sigma_n, & \text{if } \sigma_n^{\min} \leq \sigma_n \leq \sigma_n^{\max} \\ 0, & \text{otherwise,} \end{cases} \quad (2.9)$$

where $\Delta\sigma_n = \sigma_n^{\max} - \sigma_n^{\min} + 1$, and σ_n^{\max} and σ_n^{\min} are the bounds. Thus, the probability of a given model $p(M_p)$ is

$$p(M_p) = p(L)p(V|L)p(S|L)p(\sigma_n) = \frac{L!(|I_*| - L)!}{\Delta L(\Delta v)^M |I_*|! \Delta\sigma_n} \quad (2.10)$$

assuming the model parameters are within their predefined ranges, and is equal to 0 otherwise.

2.1.6 The rj-MCMC Method

The Markov chain Monte Carlo (MCMC) method provides an iterative stochastic model which generates samples from the Bayesian posterior probability density function. However, the dimensionality of the model is fixed in a standard MCMC,

unlike the rj-MCMC, which allows *jumps* in the dimensionality meaning it can add and delete model parameters, in this case the birth or death of a cell. The important step is the initial model. This could be a random model generated from the prior distributions or it could be a chosen model. One example of a chosen model is the homogeneous model as done in this thesis.

The rj-MCMC algorithm consists of seven steps, which can be observed in the workflow diagram Figure 2.2.

1. An initial model M is either decided, for example the homogeneous model, or a random model drawn from a uniform distributions of the parameters.
2. The objective function value for the initial model M , $\phi(M)$, is evaluated.

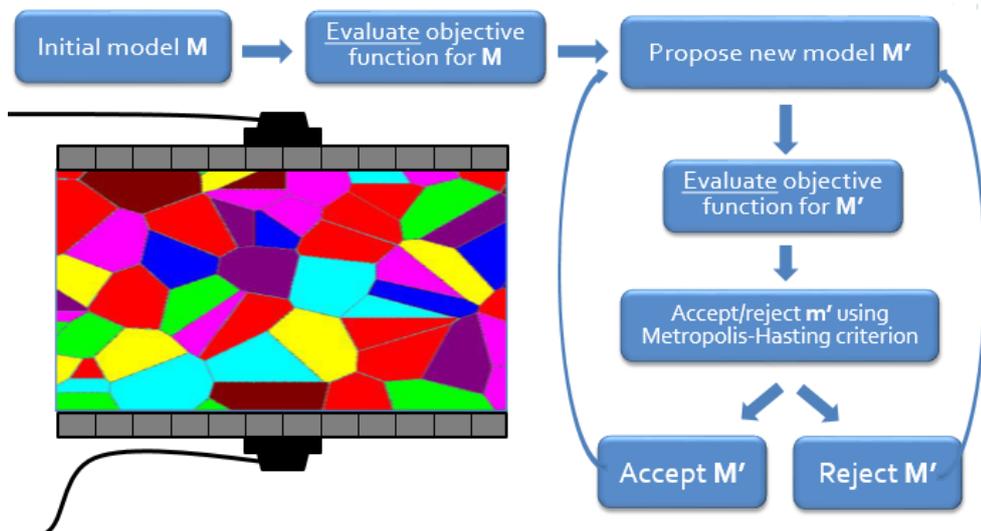


Figure 2.2: Workflow diagram of the steps within the rj-MCMC algorithm

3. A new model M' is proposed by randomly by perturbing the initial model M in some way, this could be a birth of a new cell, death of a cell, moving the position of a seed, change of wave speed within a cell, or a change of the noise parameter, as described below.
4. The objective function value of the proposed model, $\phi(M')$, is calculated.
5. The Metropolis-Hastings criteria (acceptance ratio) $\alpha(M'|M)$ is calculated according to the following equation;

$$\alpha(M'|M) = \min \left[1, \frac{p(M')}{p(M)} \times \frac{p(\mathcal{A}_{i,j}^E|M')}{p(\mathcal{A}_{i,j}^E|M)} \times \frac{q(M|M')}{q(M'|M)} \times |\mathbf{J}| \right] \quad (2.11)$$

where the second term in the bracket is the product of the prior, likelihood, and proposal ratios for M and M' , and the Jacobian of the transformation from M to M' (details below).

6. The proposed model M' is then either accepted or rejected. First generate r , a random deviate from a uniform distribution between 0 and 1. If $\alpha \geq r$ then the change is accepted and M' replaces M as the new current model. If $\alpha < r$ then the change is rejected and the model M is retained as the current model.
7. Iterate from step 3 until sufficient samples have been generated.

As stated, in step 3, a new model M' is proposed by perturbing the current model M in one of 5 ways:

- (i) The wave speed v_i within the cell i is changed subject to

$$v'_i = v_i + X\sigma_v$$

where v'_i is the proposed wave speed in cell i , X is a random variable from the standard normal distribution centred around 0 and variance of 1, and σ_v is the standard deviation of the wave speed.

- (ii) The system noise σ_n is changed subject to

$$\sigma'_n = \sigma_n + X\sigma_n^{prop}$$

where σ_n is the proposed standard deviation on the noise parameter, and σ_n^{prop} is the standard deviation of the proposal distribution for a noise perturbation.

- (iii) The coordinates of a seed $s_i = (x_i, y_i)$ of cell, are changed subject to

$$x'_i = x_i + X\sigma_c$$

and

$$y'_i = y_i + X\sigma_c$$

where x'_i, y'_i are the proposed seed location $s'_i = (x'_i, y'_i)$, where σ_c is the standard deviation of the proposed distribution for a cell move.

- (iv) The birth of a cell, where an additional seed s_{n+1} is randomly added to the set S with velocities assigned by

$$v'_{L+1} = v_{L+1}^* + X\sigma_b$$

where v_{n+1} is the proposed wave speed of the new seed s_{n+1} , v_{n+1}^* is the wave speed of the old model at s_{n+1} and σ_b is the standard deviation on the wave speed proposal for a birth.

- (v) The death of a cell, where a random seed s_i is deleted and the Voronoi is recalculated.

As stated, in step 5 and 6, of the algorithm: The perturbed model M'_1 is accepted or rejected using the Metropolis-Hastings criterion Equation (2.11). This allows for perturbed models M' that are worse than the current model M (have a higher objective function value) to be accepted. That prevents the rj-MCMC getting stuck in a local minimum. For perturbations of *fixed* dimensionality (change of wave speed, noise, or change in seed position), the acceptance parameter $\alpha(M'|M)$ only depends on the likelihoods of the current and proposed models, with the proposed model M' always accepted if its likelihood is greater than the current model M . For the case when the likelihood is less than the current model, it is either accepted or rejected, as described in step 6. In the case of fixed dimensionality, $p(M) = p(M')$ from Equation (2.10). Therefore the prior ratio is equal to 1 as

long as the proposed values are within the bounds of their respective prior:

$$\left[\frac{p(M')}{p(M)} \right]_{\text{fixed}} = \begin{cases} 1 & \text{if } M' \text{ is an acceptable model} \\ 0 & \text{otherwise} \end{cases} \quad (2.12)$$

The proposal probability $q(M'|M)$ expresses the probability to move M to M' , and $q(M|M')$ expresses the probability of the reverse, M' to M . When the dimensionality does not change, $q(M'|M)$ and $q(M|M')$ are symmetrical and their ratio is equal to 1:

$$\left[\frac{q(M|M')}{q(M'|M)} \right]_{\text{fixed}} = \begin{cases} 1 & \text{if } M' \text{ is an acceptable model} \\ 0 & \text{otherwise} \end{cases} \quad (2.13)$$

Lastly the Jacobian is always equal to 1 as shown in [100].

For the cases of perturbations with a change in dimensionality (births and deaths) the acceptance of proposed models is balanced between likelihood, prior, and proposal ratios, such that simpler models (i.e. those with fewer Voronoi cells) have a higher probability. The prior ratio for a birth is

$$\left[\frac{p(M')}{p(M)} \right]_{\text{birth}} = \begin{cases} \frac{L+1}{(I-L)\Delta v} & \text{if } M' \text{ is an acceptable model} \\ 0 & \text{otherwise} \end{cases} \quad (2.14)$$

and for a death is

$$\left[\frac{p(M')}{p(M)} \right]_{\text{death}} = \begin{cases} \frac{(I-L+1)\Delta v}{(L)} & \text{if } M' \text{ is an acceptable model} \\ 0 & \text{otherwise.} \end{cases} \quad (2.15)$$

The likelihood function is evaluated using the following

$$\frac{p(\mathcal{A}_{i,j}^E|M')}{p(\mathcal{A}_{i,j}^E|M)} = \exp - \frac{\phi(M') - \phi(M)}{2}. \quad (2.16)$$

The proposal distributions $q(M'|M)$ and $q(M|M')$ are not symmetrical and are different depending on if the perturbation is a birth or a death. For a birth of the seed s_{L+1} the ratio is

$$\left[\frac{q(M|M')}{q(M'|M)} \right]_{\text{birth}} = \frac{\sqrt{2\pi}(I-L)}{L+1} \sigma_b \exp \frac{(v'_{L+1} - v_i)^2}{2\sigma_b^2} \quad (2.17)$$

where v_{L+1} is the wave speed of the proposed cell and v_i is the present wave speed for the seed s_i . For a death

$$\left[\frac{q(M|M')}{q(M'|M)} \right]_{\text{death}} = \frac{L}{\sigma_b \sqrt{2\pi}(I-L+1)} \exp - \frac{(v'_j - v_i)^2}{2\sigma_b^2} \quad (2.18)$$

where v'_j is the wave speed of the s_j in the new tessellation (after the death of s_i). This gives the acceptance, for the birth of a cell,

$$\alpha(M'|M)_{\text{birth}} = \begin{cases} \min \left[1, \frac{\sigma_b \sqrt{2\pi}}{\Delta v} \exp \frac{(v'_{n+1} - v_i)^2}{2\sigma_b^2} - \frac{\phi(M') - \phi(M)}{2} \right] & \text{if } M' \text{ is an acceptable model} \\ 0 & \text{otherwise} \end{cases} \quad (2.19)$$

and for the death a cell,

$$\alpha(M'|M)_{\text{death}} = \begin{cases} \min \left[1, \frac{\Delta v}{\sigma_b \sqrt{2\pi}} \exp -\frac{(v'_j - v_i)^2}{2\sigma_b^2} - \frac{\phi(M') - \phi(M)}{2} \right] & \text{if } M' \text{ is an acceptable model} \\ 0 & \text{otherwise.} \end{cases} \quad (2.20)$$

For an efficient exploration of the model space, the standard deviations $\sigma_v, \sigma_n, \sigma_c, \sigma_b$ need to be tuned such that they have an acceptance rate in the range of 23% - 44% [143]. With an acceptance rate less than 23% the algorithm is not accepting enough modal perturbations and greater than 44% means the algorithm is accepting too many modal perturbations for efficient convergence of the Markov chain. Furthermore, a delayed rejection scheme is implemented, as used in [41], for perturbations of wave-speed and seed location. If rejection occur for these perturbations, a secondary perturbations with a smaller standard deviation ($\sigma_v^{dr} < \sigma_v$, $\sigma_c^{dr} < \sigma_c$) is implemented. This improves the performance of the algorithm by decreasing the probability of remaining in the current state and avoid persistent rejection [144, 145]. It should be noted that the perturbations in the algorithm is essential in ensuring convergence of the Markov chain.

2.1.7 Sampling the Posterior Distribution

To create reliable posterior probability distributions enough samples need to be generated. This is determined when the Markov chain has converged. However, there are no dependable methods to ensure that convergence has been achieved [41]. In this work, once the acceptance ratios have been tuned to the 23% - 44% range, convergence has been achieved when the objective function value, noise, and number of cells, exhibit stationarity.

Once it is believed that convergence has been achieved, an initial number of samples are discarded. This is known as the burn-in period and allows for the removal of any influence from the initial mode, however, a burn-in is not always necessary [97]. Since each model is a modification of the model before it is clear that the MCMC with i iterations does not have i independent samples. To try and make each sample independent of the others a process of “thinning” the chain is implemented. This is when only every κ^{th} iteration is collected. From these remaining samples posterior probability distributions can be generated from statistical moments at each discrete point. For this thesis, the statistical moments used are the following measures: the *arithmetic mean*, which is ideal for Gaussian posterior distributions but can be affected by outliers or by tails of the distribution; the *median*, a robust statistic that is less affected by outliers than the arithmetic mean; the *mode* or *maximum a posteriori* (MAP), which compared to the other two preserves discrete character of the ensembled Voronoi cells. Therefore

the maps are normally characterised but sharp discontinuities between different wave speeds; and the standard deviation which allows for a direct measure of the uncertainty in the maps.

2.1.8 Imaging a Flaw

Using an FMC dataset obtained from a finite element simulation in conjunction with a flaw imaging algorithm to detect and characterise a flaw within a material. The Total Focusing Method (TFM) is a benchmark flaw imaging algorithm within ultrasonic NDT [9] and has been shown to outperform other techniques such as SAFT and time reversal MUSIC [38, 39]. It assumes that the host material is homogeneous. The basis of the algorithm calculation of the distance from each transmitting element to each pixel in the image domain and then to each receiving element. Since it assumes a homogeneous material, i.e. constant wave speed, the journey time can be estimated and related to a time point on every A-scan. Each pixel can be associated with a given intensity I by summing these amplitudes over the set of transmit/receive pairs. This can be formulated as

$$I_{i,j} = \left| \sum_{tx=1}^p \sum_{rx=1}^p \mathcal{A}_{tx,rx} \left(\frac{\sqrt{(x_{tx} - x_i)^2 + y_j^2} + \sqrt{(x_{rx} - x_i)^2 + y_j^2}}{\bar{c}} \right) \right|, \quad (2.21)$$

where \bar{c} is the homogeneous speed, x_{tx} and x_{rx} are the x -coordinates of the transmitter and receiver elements, and $I_{i,j}$ is the intensity of the pixel at location (x_i, y_j) . This is demonstrated in the Figure 2.3, which shows the ray going from

the transmitter t_x to the pixel located at (x, y) and then to the receiver r_x . The travel time is calculated and the associated time point in the A-scan is shown by the orange arrow.

The host material is very often heterogeneous and if a map of the material properties can be obtained then a modified TFM approach (referred to as TFM+ [41]) can be used. This uses a Multistencil fast marching method (MSFM) [41, 146] to obtain the travel times from each transmitting element to each pixel in the image domain and from each of these pixels to each receiving element. This is formulated as

$$I_{i,j} = \left| \sum_{t_x=1}^N \sum_{r_x=1}^N \mathcal{A}_{t_x, r_x} (\tau^{t_x}(x_i, y_j) + \tau^{r_x}(x_i, y_j)) \right|, \quad (2.22)$$

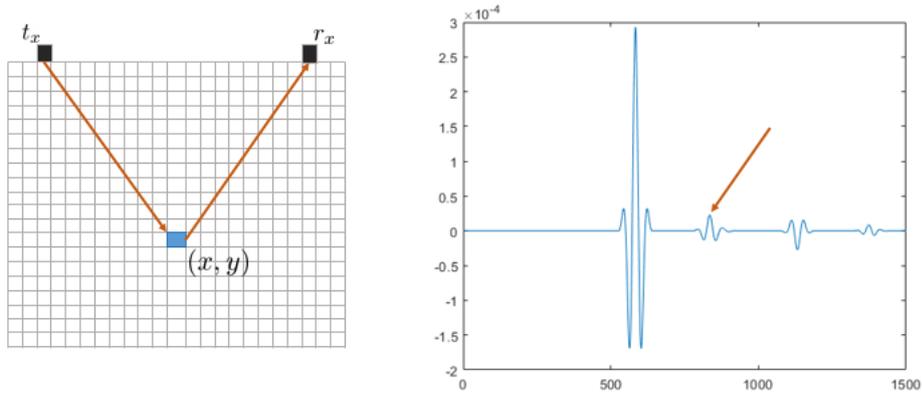


Figure 2.3: Example of a how the TFM image algorithm is calculated.

where $\tau^{t_x}(x_i, y_j)$ is the travel-time from transmitter t_x to pixel (x_i, y_j) and $\tau^{r_x}(x_i, y_j)$ is the travel-time from the pixel to the receiver r_x . This approach not only incorporates the spatially varying wave speeds but also refraction via Fermat's principle. It will be observed later that approximate material maps, obtained from the moments of the posterior distributions, arising from a Bayesian tomography method, are used in conjunction with the TFM+ algorithm.

2.1.9 Quality Metrics

It is useful to quantify the quality of the image reconstructions from the TFM and TFM+ algorithms so that objective comparisons can be made. There are four metrics used within this thesis, signal-to-noise ratio (SNR), flaw location error, ROC curves and the area under the curve (AUC) of these ROC curves.

The SNR is calculated by finding the maximum flaw intensity value (A_{signal}) within a rectangular region Ω in the image domain I around the inclusion, the mean value (A_{noise}) in a rectangle Ω_0 in a region far away from the flaw location, and the formula

$$\text{SNR} = 20 \log_{10} \left(\frac{A_{signal}}{A_{noise}} \right). \quad (2.23)$$

The flaw position is estimated by finding the centre of mass x_c , or max intensity, within a square centred around the known location of a flaw within the image domain I . The location error e_{loc} is then given by the distance between the known inclusion location x^* and this calculated centre of mass (or max intensity)

via

$$e_{loc} = ||x_c - x^*||. \quad (2.24)$$

Receiver operating characteristic (ROC) curves have been a tool to visualise and evaluate classifiers within the medical field [147, 148] and in NDT [41, 149]. One approach within the NDT community is to use a series of experiments and/or simulations to determine the probability of detection [149]. However, the method used in this thesis (and the following definition is from [41]). This uses a quantitative technique to compare different imaging algorithms (TFM and TFM+), or different wave speed maps used in conjunction with the TFM+ algorithm, on the same FMC dataset. The first step is to take the ground truth (a binary true map or TFM+ image using the known wave speed map) and divide this image into a grid (see Figure 5.5 as an example). Each grid cell that contains an intensity above -4dB threshold is used to define the flaw domain Ω_f (there are no other artefacts “false flaws” at -4dB). The number of grid cells assigned to the flaw domain is recorded as n^p and the remaining number of grid cells is recorded as $n^f = n^t - n^p$, where n^t is the total number of grid cells. These quantities are used as the denominators in the calculations of probability of detection and false positive rates. The ROC curves for the images being tested, standard TFM, TFM+ from reconstructed material maps, and TFM+ with the known map (if using a binary map as ground truth), are produced by implementing the following steps:

1. The images are partitioned using the same grid used to define Ω_f

2. The number of grid cells which lie above a threshold D and lie within Ω_f is denoted n^+ . The number of grid cells which meet this threshold but do not lie within Ω_f is denoted by n^- .
3. The probability of detection (PoD) value is calculated as n^+/n^p and the false positive rate (FPR) value is given by n^-/n^f .
4. These calculations are repeated at decreasing thresholds to produce the ROC curve which plots the probability of detection against the false positive rate over the range of selected thresholds.

In this thesis, an initial threshold of $D = -1\text{dB}$ is chosen, the probability of detection (PoD) and false positive rate (FPR) is calculated. This is repeated down to -40dB to obtain the ROC curves. For a perfect classification performance the ROC curve passes through the point (0,1) (the area under the curve (AUC) is equal to 1) indicating 100% PoD with a zero FPR. Normally this does not happen and the AUC lies in the set $[0,1]$.

2.2 Waves in Layered Medium

This section provides a summary of chapter 3 of [150]. The one-dimensional acoustic wave equations are defined and then describe the physics when a wave interacts with a single interface. The equations for a wave travelling through a layered medium are then defined. The equations for the one dimensional velocity

u and pressure p are

$$\rho(z) \frac{\partial u(t, z)}{\partial t} + \frac{\partial p(t, z)}{\partial z} = 0, \quad (2.25)$$

$$\frac{1}{K(z)} \frac{\partial p(t, z)}{\partial t} + \frac{\partial u(t, z)}{\partial z} = 0, \quad (2.26)$$

where ρ is the density of the medium, K is the bulk modulus of the medium, and z is the spatial coordinate along the one-dimension medium. The quantities $c(z) = \sqrt{K(z)/\rho(z)}$ and $\zeta(z) = \sqrt{K(z)\rho(z)}$ represent the local *wave speed* and *impedance* respectively. From these wave equations the left- and right-going waves can be defined as,

$$A(t = 0, z) = \zeta^{-1/2} p_0(z) + \zeta^{1/2} u_0(z), \quad (2.27)$$

$$B(t = 0, z) = -\zeta^{-1/2} p_0(z) + \zeta^{1/2} u_0(z). \quad (2.28)$$

For the transmitted wave function, $f(t)$, Equation (2.27) and Equation (2.28) can be re-written as the right-going wave

$$A = \begin{cases} f(t - z/c), & \text{if } z > 0 \\ 0, & \text{if } z < 0, \end{cases} \quad (2.29)$$

and the left-going wave

$$B = \begin{cases} 0, & \text{if } z > 0 \\ f(t + z/c), & \text{if } z < 0. \end{cases} \quad (2.30)$$

Now consider two homogeneous half-spaces separated by an interface at $z = 0$. Prescribe each half-space with the material properties $\rho_j, K_j, c_j,$ and ζ_j , where $j = 0$ for the left half-space and $j = 1$ for the right half-space. If a wave is generated in the left half-space there will be a right-going incident wave given by $A_0(t, -z) = f(t)$. Once this wave interacts with the interface two waves will be generated. A left-going reflected wave $B_0(t, 0) = Rf(t)$ in the left half-space and a right-going transmitted wave $A_1(t, 0) = Tf(t)$ in the right half-space ($B_1(t, 0) = 0$ since you cannot get a left-going wave in the right half-space), where R and T are the *reflection* and *transmission coefficients* of the interface given by

$$R = \frac{\zeta_0 - \zeta_1}{\zeta_0 + \zeta_1}, \quad T = \frac{2\sqrt{\zeta_0\zeta_1}}{\zeta_0 + \zeta_1}. \quad (2.31)$$

These coefficients satisfy the energy conservation

$$R^2 + T^2 = 1. \quad (2.32)$$

From Equation (2.32), it can deduced that $T = \sqrt{1 - R^2}$.

Now consider a heterogeneous slab, of length L , consisting of N layers with each layer being homogeneous between two homogeneous half-spaces. The j th layer corresponds to the interval $[L_{j-1}, L_j]$ with $L_0 = 0$ and

$$\rho(z) = \begin{cases} \rho_0 & z < 0 \\ \rho_j & z \in [L_{j-1}, L_j) \\ \rho_{N+1} & z > L_N, \end{cases} \quad K(z) = \begin{cases} K_0 & z < 0 \\ K_j & z \in [L_{j-1}, L_j) \\ K_{N+1} & z > L_N, \end{cases} \quad (2.33)$$

with local velocities c_j and impedance ζ_j .

Similar to before, in the left half-space, a right-going source wave $A_0(t, 0) = f(t)$ will generate a left-going wave, via the convolution $B_0(t, 0) = \mathcal{R} * f(t)$ and a right-going wave $A_{N+1}(t, L) = \mathcal{T} * f(t)$, with $B_{N+1}(t, L) = 0$. The coefficients \mathcal{R} and \mathcal{T} denote the *global* reflection and *transmission* coefficients. These can be calculated, in the frequency domain, from the reverse recursive relationships given by

$$\hat{\mathcal{R}}_j(\omega) = \frac{R_j + \hat{\mathcal{R}}_{j+1}(\omega)}{1 + R_j \hat{\mathcal{R}}_{j+1}(\omega)} e^{2i \frac{\omega(L_j - L_{j-1})}{c_j}}, \quad j = 0, \dots, N, \quad (2.34)$$

where ω is the frequency and $R_j = \frac{\zeta_j - \zeta_{j+1}}{\zeta_j + \zeta_{j+1}}$ is the reflection coefficient at the j th interface. The “final” condition at $j = N + 1$ is $\hat{\mathcal{R}}_{N+1} = 0$. This can then be used to determine the global transmission coefficient via

$$\hat{\mathcal{T}}_j(\omega) = \frac{T_j \hat{\mathcal{T}}_{j+1}(\omega)}{1 + R_j \hat{\mathcal{R}}_{j+1}(\omega)} e^{i \frac{\omega(L_j - L_{j-1})}{c_j}}, \quad j = 0, \dots, N, \quad (2.35)$$

where $T_j = \frac{2\sqrt{\zeta_0\zeta_1}}{\zeta_0+\zeta_1}$ is the transmission coefficient at the j th interface. Using Equation (2.34), $\hat{\mathcal{R}}$ can be expanded as a series to form

$$\hat{\mathcal{R}}(\omega) = \sum_{j=1}^N \sum_{k_j=0}^{\infty} \alpha_{k_1, \dots, k_n} \exp\left(2i\omega \sum_{j=1}^N k_j \frac{L_j - L_{j-1}}{c_j}\right), \quad (2.36)$$

where α_{k_1, \dots, k_n} is a coefficient that only depends on R_0, \dots, R_N . Taking the inverse Fourier transform yields

$$\mathcal{R}(t) = \sum_{j=1}^N \sum_{k_j=0}^{\infty} \alpha_{k_1, \dots, k_n} \delta\left(t - 2 \sum_{j=1}^N k_j \frac{L_j - L_{j-1}}{c_j}\right) \quad (2.37)$$

where δ is the Dirac delta function. Each term in the series is associated with a scattering sequence involving reflections and transmissions at different interfaces that determine the value of α_{k_1, \dots, k_n} . Similarly, expanding Equation (2.35) and applying an inverse Fourier transform yields

$$\mathcal{T}(t) = \sum_{j=1}^N \sum_{k_j=0}^{\infty} \beta_{k_1, \dots, k_n} \delta\left(t - 2 \sum_{j=1}^N k_j \frac{L_j - L_{j-1}}{c_j} - \sum_{j=1}^N k_j \frac{L_j - L_{j-1}}{c_j}\right), \quad (2.38)$$

where β_{k_1, \dots, k_n} is a coefficient that only depends on R_0, \dots, R_N since $R_j = \sqrt{1 - T_j^2}$. At this point a forward model of how waves travel through layered media is ready to be constructed.

2.3 Discussion

In this chapter, the components of the inversion method are defined such as: Voronoi tessellation; Pearson Correlation Coefficient; reversible-jump Mark chain Monte Carlo. The metrics to quantify the inversion are given: TFM and TFM+ imaging algorithm; flaw location errors; Signal-to-noise ratio; and receiver operating characteristic (ROC) curves. This is then followed by the mathematics of wave propagation in layered media. The mathematics is then taken into the next chapter were a novel semi-analytical forward model of wave propagation through layered media is developed.

Chapter 3

Development of the Forward Model

This chapter is the start of original content for this thesis. This section details the development of a forward model to be used with the rj-MCMC to complete the inversion scheme used for this thesis.

3.1 Forward Model

A mathematical model of an ultrasonic signal travelling through a 1D layered medium is employed. The advantage of using this ray based model, is that refraction and mode decomposition do not occur. One can view a straight ray through a Voronoi tessellation as a ray going through a layered medium. This forward model is based on the global reflection and transmission Equation (2.37) and Equation (2.38).

3.1.1 Deriving the amplitudes reflected wave and transmitted waves

Deriving a closed form expression for the coefficients α_{k_1, \dots, k_n} in Equation (2.37) is not trivial and so a method for systematically generating these coefficients was devised to gain insight on the behaviour of acoustic waves in multi-layered media.

Denote the amplitude of any partial wave in layer j by I_j . If this wave is travelling to the right then, when it interacts with the j^{th} interface it creates a left going wave in layer j of amplitude $R_j I_j$ and a right going wave in layer $j + 1$ of amplitude $T_j I_j$. If this wave of amplitude I_j is travelling to the left then, when it interacts with the $(j - 1)^{\text{th}}$ interface it creates a left going wave in layer $j - 1$ of amplitude $T_{j-1} I_j$ and a right going wave in layer j of amplitude $-R_{j-1} I_j$.

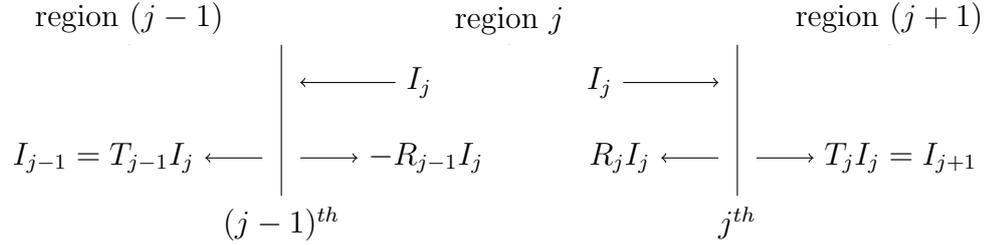


Figure 3.1: Schematic of a left- and right-going wave interaction at an interface.

For example, the incident wave is generated by a point source of amplitude I_0 in the left half space, region $j = 0$, that travels through the half-space until it interacts with the front-wall interface (interface $j = 0$) to form two waves, a left going reflected wave, with amplitude $R_0 I_0$, which travels into the left half space and a right going transmitted wave, with amplitude $T_0 I_0$, which travels into layer

3.

$j = 1$ of the medium. The transmitted wave now becomes the incident wave, with amplitude $I_1 = T_0 I_0$, and creates a left going wave reflected wave with amplitude $R_1 I_1$, and a right going transmitted wave with amplitude $T_1 I_1$ in region $j = 2$; these then both become new incident waves.

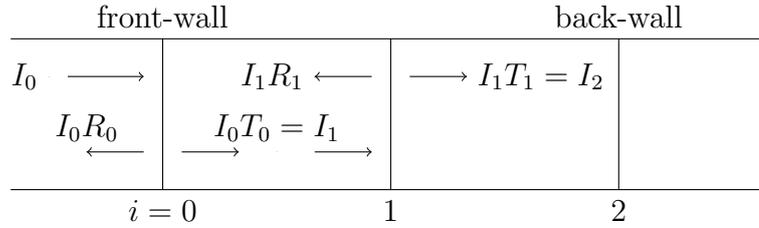


Figure 3.2: Schematic of wave interaction at interfaces between two layers.

This series of waves continue to reverberate throughout the layered medium generating an infinite series of incident waves. Some of the wave energy leaves the system when it is reflected into the left half-space and some when it is transmitted into the right half-space. This process can be represented by a tree graph (see Figure 3.3).

Each vertex of the tree represents a reflection/transmission at a particular interface. The series of reflection and transmission events which a wave encounters is then represented by a walk on this graph which starts at a vertex T_0 (transmission into the layered medium) and terminates at one of the vertices denoted by T_3 (when the wave enters the right hand half-space) or T_0 (when the wave enters the left hand half-space). For the walks which terminate at vertices denoted T_0 , within some finite given time frame, explicit expressions for the α_k in Equation (2.37) can be derived from the product of the set of vertices in each walk. For

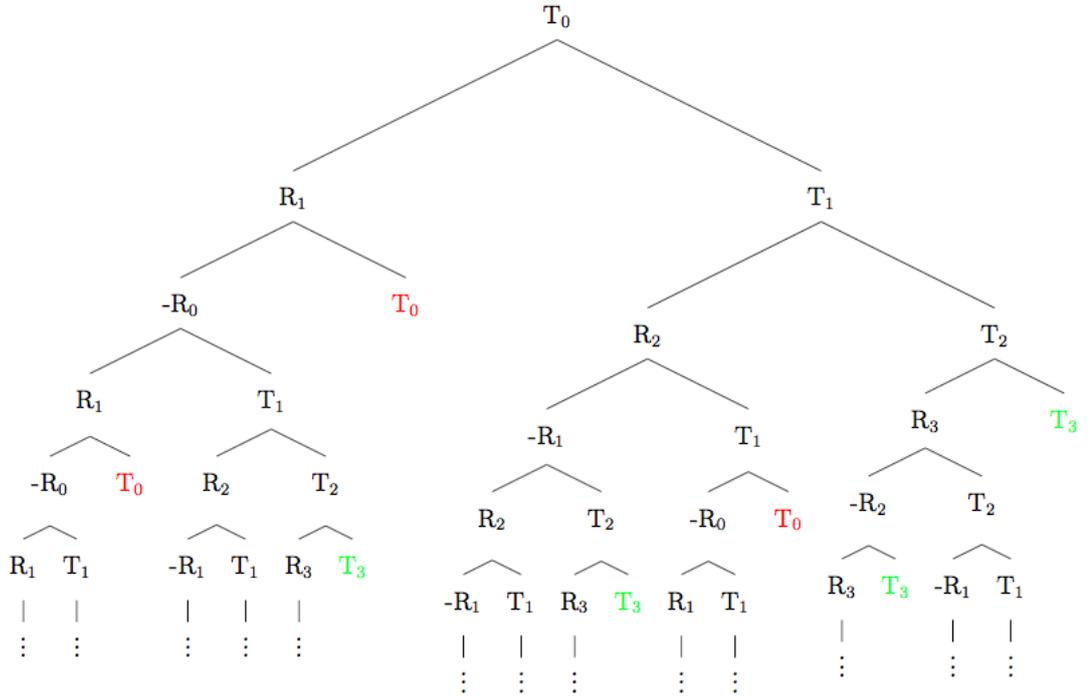


Figure 3.3: Tree graph for a 3 layered medium. When a wave interacts with an interface, the energy is both reflected and transmitted. This is represented by the branching of this tree graph, starting with T_0 transmission into the medium. The vertices of degree 1 (end vertices) are denoted by T_0 (when the wave is transmitted into the left half space) and T_3 (when the wave is transmitted into the right half space). The amplitude of the wave is then the product of the vertex labels along a walk on this graph.

walks that terminate at vertices denoted by T_3 , in some finite given time frame, explicit expressions for the β_k in Equation (2.38) can be derived from the product of the set of vertices in each walk. To calculate the time of arrival the number of times each layer is traversed is tallied up and the travel times are summed. This data is represented by a matrix with one row of amplitudes and another row with the corresponding time of arrivals. This would give an A-scan with a se-

ries of point amplitudes (Dirac delta functions from Equation (2.38)). Therefore, each amplitude is convolved with the Ricker Wavelet as shown in Figure 3.4 and brought together to give an A-scan.

This model was numerically evaluated and the resulting wave signature was then compared to a wave signature arising from a finite element simulation of a wave in the same layered medium. The finite element package OnScale [141] was used to simulate a plane wave traversing a 2D layered model. Absorbing boundary implemented on the domain's vertical edges and free boundaries were used at the front-wall and back-wall interfaces.

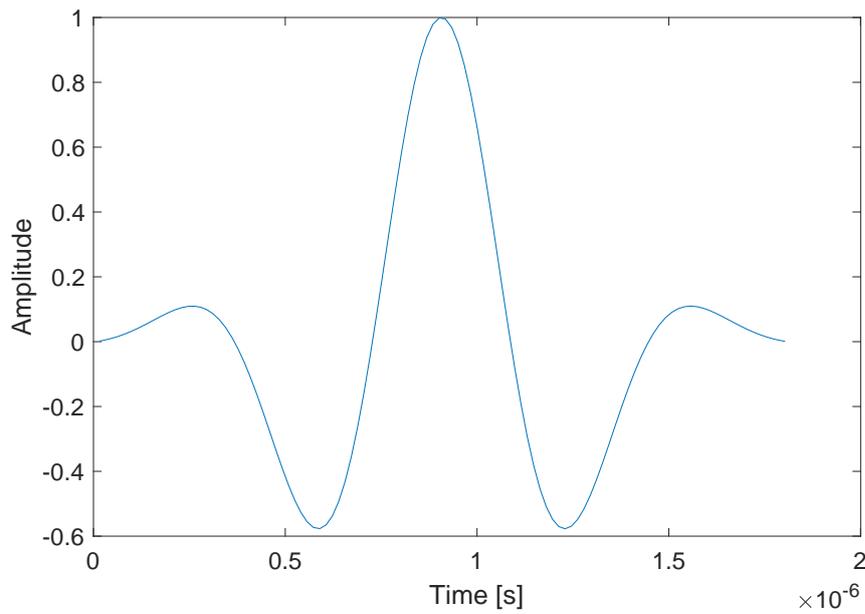


Figure 3.4: Figure of a single wave packet with amplitude 1, that is convolved with the wave amplitudes to create an A-scan.

3.

Layer Number l	Material Velocity (m/s)	Material Density (kg/m ³)
1	6420	2690
2	5700	8280
3	5900	7703
4	6100	4480
5	5010	8930

Table 3.1: Material property list. Layer length is 10mm for each layer.

The domain geometry consisted of five layers, see Table 3.1. The left and right half spaces were assigned the properties of water and air respectively. The resulting reflected wave as given by Equation (2.37) is shown in Figure 3.5.

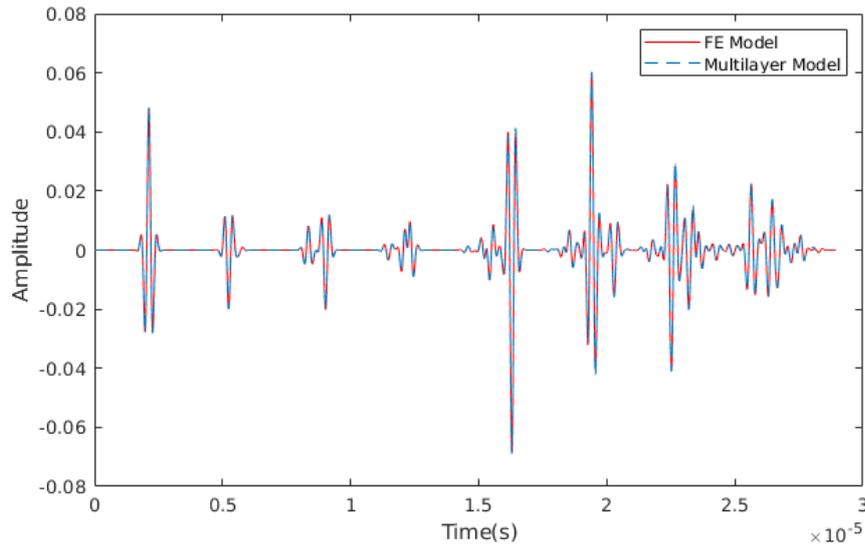


Figure 3.5: Time domain signal comparison between a finite element simulation (On-Scale) and the numerical evaluation of Equation (2.37) of an ultrasonic wave propagating through a five layer medium, see Table 3.1. The time range is determined by the time taken for the ultrasonic wave to reflect off the back-wall (interface $j=5$) and return to the receiver at the front-wall (interface $j=0$).

It can be observed that the two modelling approaches demonstrate excellent agreement with a Pearson Correlation Coefficient of 0.998. The front wall reflection has been omitted as they are not comparable between the models; the finite element model was excited at the $j = 0$ interface and recorded the wave as it was transmitted from the transducer element, whereas our model only simulated the initial reflected wave off the interface. Unfortunately this model, was not computationally efficient as the volume of data required grew exponentially with the number of layer interactions (maximum path length) within the medium and grew beyond the data limits of MATLAB at 24 interactions. As this model will form the forward solver for an inverse problem formulation, computational efficiency is essential. This led to the development of an approximated model that is more computationally efficient but still retains enough information to produce a comparable wave signature.

3.1.2 Approximation of the Reflected & Transmitted Wave

To approximate the analytic model, a restriction to a limited number of significant peaks within a certain time range is implemented. This approach dramatically decreases the number of calculations required whilst still producing a wave signature which captures the most important wave-layer interactions.

Recall Equation (2.34)

$$\hat{\mathcal{R}}_j(\omega) = \frac{R_j + \hat{\mathcal{R}}_{j+1}(\omega)}{1 + R_j \hat{\mathcal{R}}_{j+1}(\omega)} e^{2i \frac{\omega(L_j - L_{j-1})}{c_j}}, \quad j = 0, \dots, N,$$

an expression for the reflection coefficient $\hat{\mathcal{R}}_j(\omega)$, at layer j , of a wave travelling through a randomly layered medium in the frequency domain (ω). A backward recursive relation, starting with $\hat{\mathcal{R}}_{N+1}(\omega) = 0$. By taking the first two terms of the Taylor series expansion of $R_j \hat{\mathcal{R}}_{j+1}$ (it can safely be assumed this is between $[-1,1]$), this can be approximated by

$$\begin{aligned}\hat{\mathcal{R}}_j(\omega) &= (R_j + \hat{\mathcal{R}}_{j+1}(\omega))(1 - R_j \hat{\mathcal{R}}_{j+1}(\omega))e^{2i\omega t_j} + \mathcal{O}[(R_j \hat{\mathcal{R}}_{j+1}^2(\omega))] \\ &= ((1 - R_j^2) \hat{\mathcal{R}}_{j+1}(\omega) - R_j \hat{\mathcal{R}}_{j+1}^2(\omega) + R_j)e^{2i\omega t_j} + \mathcal{O}[(R_j \hat{\mathcal{R}}_{j+1}^2(\omega))],\end{aligned}\quad (3.1)$$

where $t_j = (L_j - L_{j-1})/c_j$ is the time it takes for the wave to travel through layer j and $t_0 = 0$. Then

$$\hat{\mathcal{R}}_N(\omega) = R_N e^{2i\omega t_N},$$

and

$$\begin{aligned}\hat{\mathcal{R}}_{N-1}(\omega) &\approx (1 - R_{N-1}^2) \hat{\mathcal{R}}_N(\omega) e^{2i\omega t_{N-1}} - R_{N-1} \hat{\mathcal{R}}_N^2(\omega) e^{2i\omega t_{N-1}} + R_{N-1} e^{2i\omega t_{N-1}} \\ &= (1 - R_{N-1}^2) R_N e^{2i\omega(t_N + t_{N-1})} - R_{N-1} R_N^2 e^{2i\omega(2t_N + t_{N-1})} + R_{N-1} e^{2i\omega t_{N-1}}.\end{aligned}$$

Continuing in this manner it can be shown that

$$\begin{aligned}\hat{\mathcal{R}}_j(\omega) &= \alpha_1^j e^{2i\omega(\mathbf{K}_1^j \cdot \mathbf{t})} + \dots + \alpha_{m_j}^j e^{2i\omega(\mathbf{K}_{m_j}^j \cdot \mathbf{t})} \\ &= \sum_{k=1}^{m_j} \alpha_k^j e^{2i\omega(\mathbf{K}_k^j \cdot \mathbf{t})},\end{aligned}\quad (3.2)$$

3.

where α_k^j is a coefficient that depends only on R_j, \dots, R_N . The travel times for each of the layers is stored in the vector \mathbf{t} , the number of times the wave travels through each layer is stored in the vector \mathbf{K}_k^j , and m_j is the number of distinct partial wave arrival times for $N-j$ layers (under this Taylor series approximation). Both α_k^j and \mathbf{K}_k^j are complicated, nested functions and are difficult to write them as a closed form expression. The vector can be discretised by a fixed time step Δt and therefore can be rewritten as

$$\begin{aligned}
 \mathbf{K}_k^j \cdot \mathbf{t} &= K_{k,0}^j t_0 + \dots + K_{k,N}^j t_N \\
 &= K_{k,0}^j \Delta t q_0 + \dots + K_{k,N}^j \Delta t q_N \\
 &= (K_{k,0}^j q_0 + \dots + K_{k,N}^j q_N) \Delta t \\
 &= Q_k^j \Delta t
 \end{aligned}$$

where $q_j \in \mathbb{N}$ is the number of time steps in layer j . Then taking the inverse Fourier transform of Equation (3.2)

$$\mathcal{R}_j(t) = \sum_{k=1}^{m_j} \alpha_k^j \delta_{Q_k^j \Delta t} \tag{3.3}$$

where the travel time in layer j has been discretised by $t_j = Q_k^j \Delta t$.

Using Equation (3.1) and Equation (3.2) gives us the following recursive relation

$$\hat{\mathcal{R}}_j(\omega) = (1-R_j^2) \sum_{k=1}^{m_{j+1}} \alpha_k^{j+1} e^{2i\omega(\mathbf{K}_k^{j+1} \cdot \mathbf{t} + t_j)} - R_j \sum_{k=1}^{m_{j+1}} \sum_{p=1}^{m_{j+1}} \alpha_k^{j+1} \alpha_p^{j+1} e^{2i\omega((\mathbf{K}_k^{j+1} + \mathbf{K}_p^{j+1}) \cdot \mathbf{t} + t_j)} + R_j e^{2i\omega t_j}. \quad (3.4)$$

Note that the the double summation comes from

$$\hat{\mathcal{R}}_{j+1}^2(\omega) = \left(\sum_{k=1}^{m_{j+1}} \alpha_k^{j+1} e^{2i\omega(\mathbf{K}_k^j \cdot \mathbf{t})} \right)^2 = \sum_{k=1}^{m_{j+1}} \sum_{p=1}^{m_{j+1}} \alpha_k^{j+1} \alpha_p^{j+1} e^{2i\omega((\mathbf{K}_k^{j+1} + \mathbf{K}_p^{j+1}) \cdot \mathbf{t} + t_j)}.$$

Taking the inverse Fourier transform gives

$$\mathcal{R}_j(t) = (1-R_j^2) \sum_{k=1}^{m_{j+1}} \alpha_k^{j+1} \delta_{(Q_k^{j+1} + q_j)\Delta t} - R_j \sum_{k=1}^{m_{j+1}} \sum_{p=1}^{m_{j+1}} \alpha_k^{j+1} \alpha_p^{j+1} \delta_{(Q_k^{j+1} + Q_p^{j+1} + q_j)\Delta t} + R_j \delta_{q_j \Delta t}. \quad (3.5)$$

Thus, working backwards iteratively from layer N produces an expression for $\mathcal{R}_0(t)$. Then an approximation of the A-scan of a wave travelling through the multi-layered medium is given by the convolution

$$\mathcal{R}_0(t) * f(t),$$

where $f(t)$ is the function describing the source wave.

The above analysis only considers the reflection coefficient, coefficient waves that are transmitted and received at the $j = 0$ element. This model can therefore be used for modelling waves being transmitted/received by ultrasonic array transducers where the same element is performing both the transmissions of the wave and its reception. This is referred to as the Pulse-echo (PE) set-up.

Similar to the reflection case, Equation (2.35),

$$\hat{\mathcal{T}}_j(\omega) = \frac{T_j \hat{\mathcal{T}}_{j+1}(\omega)}{1 + R_j \hat{\mathcal{R}}_{j+1}(\omega)} e^{i \frac{\omega(L_j - L_{j-1})}{c_j}}, \quad j = 0, \dots, N,$$

a backward recursive relation which describes the transmitted wave through a randomly layered medium in the frequency domain using the following. Starting with $\hat{\mathcal{T}}_{N+1}(\omega) = 1$ and where $T_j = \sqrt{\zeta_j \zeta_{j+1}} / (\zeta_j + \zeta_{j+1})$ is the transmission coefficient based on the mechanical impedances of the two adjacent layers [150]. Taking the first two terms of the Taylor expansion of Equation (2.35) leads to

$$\hat{\mathcal{T}}_j(\omega) = \sum_{k=1}^{m_j} \beta_k^j e^{i\omega(\mathbf{K}_k^j \cdot \mathbf{t})}, \quad (3.6)$$

where β_k^j is a coefficient that depends on T_j, \dots, T_N and R_j, \dots, R_N , and \mathbf{K}^j is the vector storing the number of times the wave travels through each layer. Now substituting Equation (3.2) and Equation (3.6) into the first two terms of the Taylor expansion of Equation (2.35) gives the following recursive relation

$$\hat{\mathcal{T}}_j(\omega) = T_j \sum_{k=1}^{m_{j+1}} \beta_k^{j+1} e^{i\omega(\mathbf{K}_k^{j+1} \cdot \mathbf{t} + t_j)} - T_j R_j \sum_{k=1}^{m_{j+1}} \sum_{p=1}^{n_{j+1}} \beta_k^{j+1} \alpha_p^{j+1} e^{i\omega((\mathbf{K}_k^{j+1} + 2\mathbf{L}_p^{j+1}) \cdot \mathbf{t} + t_j)}.$$

Taking the inverse Fourier transform then

$$\mathcal{T}_j(t) = T_j \sum_{k=1}^{m_{j+1}} \beta_k^{j+1} \delta_{(Q_k^{j+1}+q_j)\Delta t} - T_j R_j \sum_{k=1}^{m_{j+1}} \sum_{p=1}^{n_{j+1}} \beta_k^{j+1} \alpha_p^{j+1} \delta_{(Q_k^{j+1}+Q_p^{j+1}+q_j)\Delta t}. \quad (3.7)$$

As before the method works by iterating backwards from layer N to obtain $\mathcal{T}_0(t)$. An approximation of the A-scan of a wave travelling through the multi-layered medium is then given by the convolution

$$\mathcal{T}_0(t) * f(t),$$

where $f(t)$ is the function describing the transmitted wave pulse. This model can be used for modelling waves being transmitted and received by different ultrasonic array transducers where one element is performing the transmissions of the wave and is being received by all elements on an array on the other side of the material. This is referred to as the Pitch-catch or Through-transmission (TT) set-up.

3.1.3 Pseudo-code for Approximating the Reflection & Transmitted Wave

Before the waves can be generated a few inputs are needed: the number the layers; the layer wave speeds; the layer lengths; a wavelet; and a timestep value from the wavelet. With these the wave can be generated.

3.

From a computational perspective, Equation (3.3) for $\mathcal{R}_{j+1}(t)$ takes the form of a $2 \times m_{j+1}$ matrix $[\mathcal{R}_{j+1}]$ where the first row stores the time steps Q_k^{j+1} and the second stores the corresponding amplitudes α_k^{j+1} for $k = 1, \dots, m_{j+1}$; working iteratively backwards from $[\mathcal{R}_N]$ to form the matrix $[\mathcal{R}_0]$. This is achieved by substituting the matrix $[\mathcal{R}_{j+1}]$ into equation (3.5). The first term

$$(1 - R_j^2) \sum_{k=1}^{m_{j+1}} \alpha_k^{j+1} \delta_{(Q_k^{j+1} + q_j)\Delta t}, \quad (3.8)$$

is obtained by adding the value q_j to all terms in the first row (the discretised arrival times) and multiplying the second row (the amplitudes of the reflected waves) by $(1 - R_j)$ giving a $2 \times m_{j+1}$ matrix. For the second term

$$-R_j \sum_{k=1}^{m_{j+1}} \sum_{p=1}^{m_{j+1}} \alpha_k^{j+1} \alpha_p^{j+1} \delta_{(Q_k^{j+1} + Q_p^{j+1} + q_j)\Delta t}, \quad (3.9)$$

add each $Q_p^{j+1} + q_j$ to each Q_k^{j+1} term in the first row. That is, starting with $k = 1$ (Q_1^{j+1}) term in the first row. Create m_j terms by add this in turn to each of the $p = 1, \dots, m_{j+1}$ terms, that is $\{Q_1^{j+1} + Q_1^{j+1}, Q_1^{j+1} + Q_2^{j+1}, \dots, Q_1^{j+1} + Q_{m_{j+1}}^{j+1}\}$. Repeat this for $k = 2$ (Q_2^{j+1}), and so on until $k = m_{j+1}$. Giving rise to $m_{j+1}(m_{j+1} + 1)/2$ terms and add q_j to each of these terms. As there will be terms with the same arrival times these sum together (similar to a binomial expansion), hence why it is not m_{j+1}^2 terms. For the second row, start with $k = 1$ (α_1^{j+1}) and multiply this by each of the $p = 1, \dots, m_{j+1}$ terms, that is $\{\alpha_1^{j+1} \alpha_1^{j+1}, \alpha_1^{j+1} \alpha_2^{j+1}, \dots, \alpha_1^{j+1} \alpha_{m_{j+1}}^{j+1}\}$. Repeat this for $k = 2, \dots, m_{j+1}$, giving rise to each of the $m_{j+1}(m_{j+1} + 1)/2$ terms. Multiply each of these terms by $-R_j$.

Finally, the last term

$$R_j \delta_{q_j \Delta t} \quad (3.10)$$

forms the 2×1 matrix $[q_j, R_j]^T$. These three matrices are then concatenated to form the $2 \times m_j$ matrix $[\mathcal{R}_j]$, where $m_j = m_{j+1} + m_{j+1}(m_{j+1} + 1)/2 + 1 = (m_{j+1} + 2)(m_{j+1} + 1)/2$. This process is repeated until $[\mathcal{R}_0]$ is formed. The last step is to multiply each amplitude by a wave function $f(t)$ to create the A-scan \mathcal{A} .

The above process can be summarised by the following steps;

1. Input the number of layers, layer wave speeds, layer lengths, wavelet, and timestep value.
2. Start with the 2×1 matrix $[\mathcal{R}_j] = [q_j, R_j]^T, j = N$
3. Repeat steps (a) and (d) below $N + 1$ times where $j = N, \dots, 0$ the layer number.
 - (a) Calculate Equation (3.8): multiply each term in the second row (wave amplitudes) of $[\mathcal{R}_{j+1}]$ by $(1 - R_j^2)$ and add q_j to each term in the first row (wave arrival times) giving a $2 \times m_{j+1}$ matrix $[\mathcal{R}_j]_1$.
 - (b) Calculate Equation (3.9): multiply each k^{th} term by each p^{th} term in the second row of $[\mathcal{R}_{j+1}]$ and multiply each new term by $-R_j$. Add each k^{th} term to each p^{th} term in the first row of $[\mathcal{R}_{j+1}]$ and add q_j to each new term, giving a $2 \times m_{j+1}(m_{j+1} + 1)/2$ matrix $[\mathcal{R}_j]_2$.
 - (c) Calculate Equation (3.10): create a 2×1 matrix $[\mathcal{R}_j]_3 = [q_j, R_j]^T$.

(d) Concatenate these three matrices $[\mathcal{R}_j]_1, [\mathcal{R}_j]_2, [\mathcal{R}_j]_3$ to form a $2 \times m_j$ matrix, $[\mathcal{R}_j]$, where $m_j = (m_{j+1} + 2)(m_{j+1} + 1)/2$.

4. Create a vector which contains a single wave for each amplitude at the calculated time of arrival.
5. Sum each of these vectors to give the A-scan for the given model.

This model was implemented in MATLAB and the resulting A-scan was compared with finite element simulation, OnScale [141]. For example, 5 layers composed of the materials listed in Table 3.1, were traversed by a plane wave. On a visual comparison these compare (see Figure 3.6). To quantify their similarity, the Pearson Correlation Coefficient (PCC) was used and this had a value of 0.945.

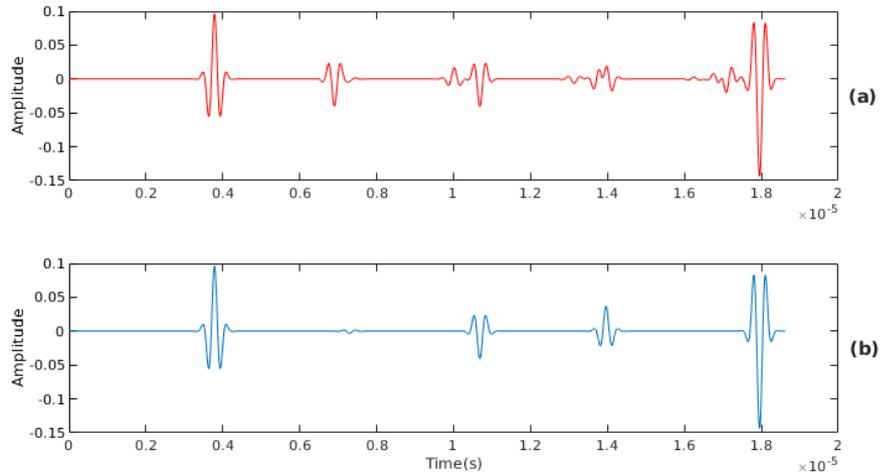


Figure 3.6: Time domain signal comparison of a finite element simulation (plot (a)) and the analytical model given by equation Equation (3.5) (plot(b)). The medium consists of 5 layers, see Table 3.1. The transmitted pulse is given by a Ricker Wavelet

3.

As the recursion relationship iterates from $j = N + 1$ to $j = 0$ it can be seen that m_j grows exponentially. Therefore to keep the computation time low and minimise memory requirements, only the most significant (highest value) amplitudes within a given time frame were kept. The time frame was determined by the length of time it takes for the wave to travel the depth of the medium twice and the number of peaks was limited to twice the number of layers as any more did not provide an increase in correlation, as seen in Figure 3.7 and the computational time grew linearly with the number of peaks, as seen in Figure 3.10.

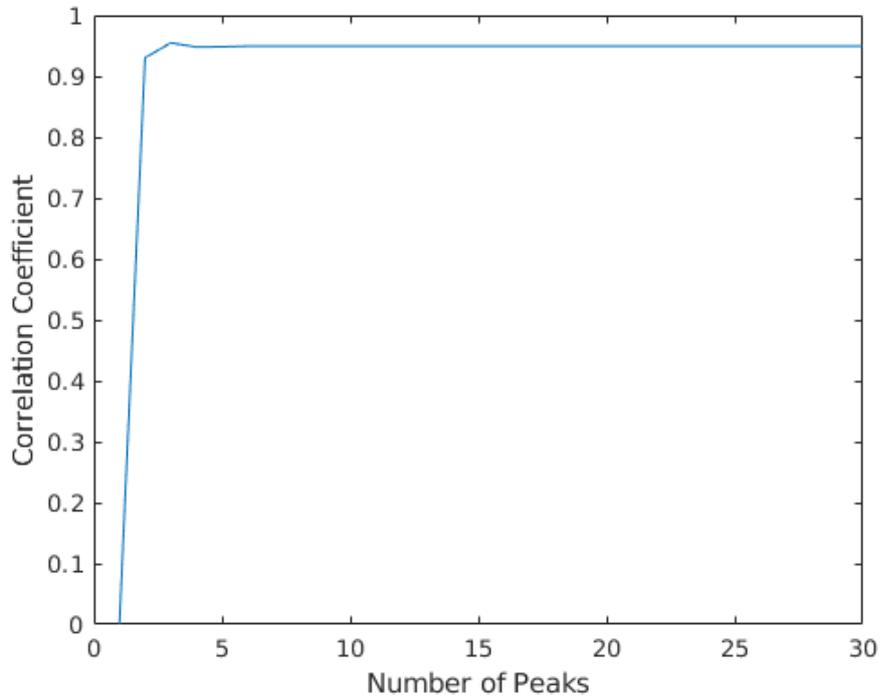


Figure 3.7: Plot of the Pearson correlation coefficient between the A-scan arising from the finite element simulation and that arising from the approximate model given by equation Equation (3.5) varying the number of peaks within a given time frame. Other parameters identical to those in Figure 3.6

3.

Similar to the reflection case, Equation (3.7) is used to form a matrix $[\mathcal{T}_j]$; working iteratively backwards from $[\mathcal{T}_N]$ to form the matrix $[\mathcal{T}_0]$. However, for the transmission case, the matrix $[\mathcal{R}_{j+1}]$ is required. The first term

$$T_j \sum_{k=1}^{m_{j+1}} \beta_k^{j+1} \delta_{(Q_k^{j+1}+q_j)\Delta t}$$

is obtained by adding the value q_j to all terms in the first row and multiply the second row by T_j to form a $2 \times m_{j+1}$ matrix. For the second term

$$-T_j R_j \sum_{k=1}^{m_{j+1}} \sum_{p=1}^{n_{j+1}} \beta_k^{j+1} \alpha_p^{j+1} \delta_{(Q_k^{j+1}+Q_p^{j+1}+q_j)\Delta t}$$

add each $Q_p^{j+1} + q_j$ to each Q_k^{j+1} term in the first row for $p = 1, \dots, m_{j+1}$, giving rise to $m_{j+1} + n_{j+1}$ terms and make the corresponding amplitudes for each term equal to $-T_j R_j \beta_k^{j+1} \alpha_p^{j+1}$ in the second row giving rise to a $2 \times (m_{j+1} + n_{j+1})$ matrix. These two matrices are then concatenated to form the $2 \times m_j$ matrix $[\mathcal{T}_j]$.

The process is the same as the reflection case with the additional following steps in step 2;

- (e) Multiply each term in the second row (wave amplitudes) by T_j and add q_j to each term in the first row (wave arrival times) giving a $2 \times m_{j+1}$ matrix.

3.

- (f) Multiply each k^{th} term by each p^{th} term in the second row of $[\mathcal{T}_{j+1}]$ and multiply each new term by $-T_j R_j$. Add each k^{th} term to each p^{th} term in the first row of $[\mathcal{T}_{j+1}]$ and add q_j to each new term, giving a $2 \times (m_{j+1} + n_{j+1})$ matrix.

This model was implemented in MATLAB and compared with a corresponding finite element simulation, OnScale [141]. For example, 5 layers composed of materials listed in Table 3.1. On a visual comparison these compare well (see Figure 3.8) and have a Pearson correlation coefficient of 0.9095.

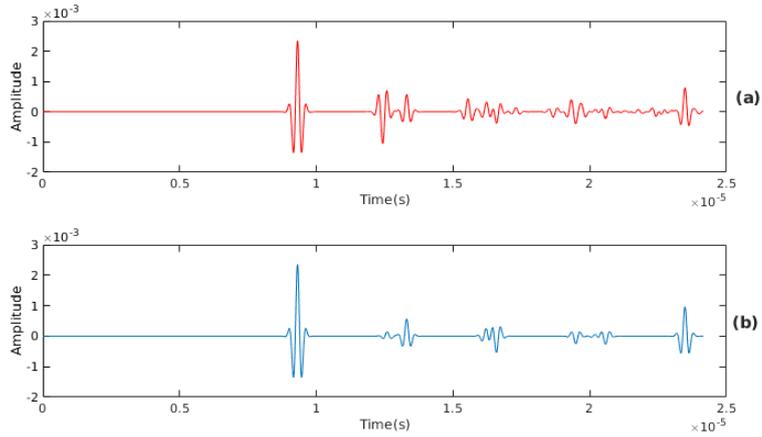


Figure 3.8: Time domain signal comparison of a finite element simulation (plot (a)) and the analytical model given by equation Equation (3.7) (plot(b)). The medium consists of 5 layers, see Table 3.1. The transmitted pulse is given by a Ricker Wavelet

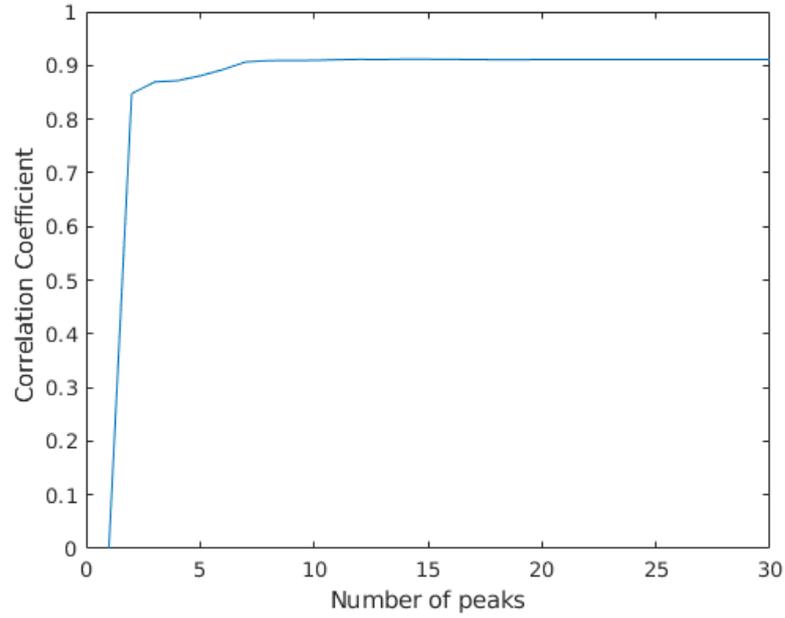


Figure 3.9: Plot of the Pearson correlation coefficient between the finite element simulation and the approximate model given by Equation (3.7) varying the number of peaks within a given time frame. Other parameters identical to those in Figure 3.6

3.2 Discussion

In this chapter, the semi-analytical model that will be used within the inversion scheme has been developed. First a detailed model was developed that proved to be too computationally expensive. A new model was developed that was more computationally efficient. The pseudo code for this semi-analytical model is then given. In the next chapter, this semi-analytical model is used to reconstruct a locally isotropic randomly heterogeneous material.

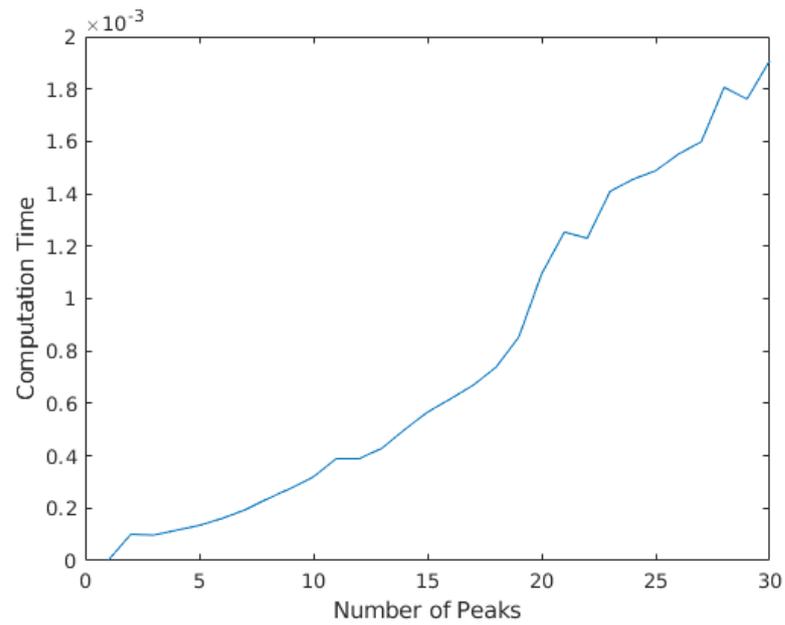


Figure 3.10: Plot of the computation time of the approximate model given by Equation (3.7) varying the number of peaks retained within a given time frame. Other parameters identical to those in Table 3.1

Chapter 4

Tomographic Reconstruction of a Randomly Heterogeneous Material

In this chapter, the methods and tools described in Chapter 2 is brought together to attempt a Full waveform inversion (FWI) of a heterogeneous material with a flaw present. The first case examines is a synthetically generated random material consisting of locally isotropic materials represented by a Voronoi diagram generated using 100 seeds. This case was studied in [41] which used time of flight (ToF) data within an rj-MCMC algorithm to obtain an approximation to the material wave speed map. That work proved that the rj-MCMC framework can successfully obtain an approximation of the material map given sufficient ToF data. When used in conjunction with an imaging algorithm, known as TFM+, the authors showed that it can improve the signal-to-noise ratio (SNR), the flaw location error, and the probability of flaw detection (measured by the area under

the curve generated by ROC curves). Time of flight uses a single data point and discards a rich data set of potential information contained in the full recorded wave. The aim of this thesis is to use the full wave, and by using more information show that an improved approximation of the material map can be obtained from the rj-MCMC algorithm, which, when used to correct for the delay laws within the TFM+ algorithm, should provide better flaw reconstructions compared to the one generated with the use of ToF.

4.1 Randomly Locally Isotropic Heterogeneous Material

The synthetically generated random material found in [41], with dimensions 64mm x 40mm, material wave speeds varying from 4000 to 7000 m/s, and a constant density of 7890kg/m³ was used to create the observed data (see Figure 4.1). To provide some ground truth, data from a finite element simulation, run in OnScale [141] was considered. A finite element model was used to generate a full matrix capture (FMC) dataset (an array of size $p \times p \times t$, where $p = 32$ is the number of array elements, and t is the number of time steps recorded by the receiver). A through-transmission (two arrays one transmitting and another receiver on the other side of the medium) set up was employed, with arrays with a 2mm pitch, and a centre frequency of 1.5MHz. An embedded 4mm diameter side drill hole (SDH) centred around (31mm, 20mm), was also included in the simulation to facilitate flaw imaging comparisons.

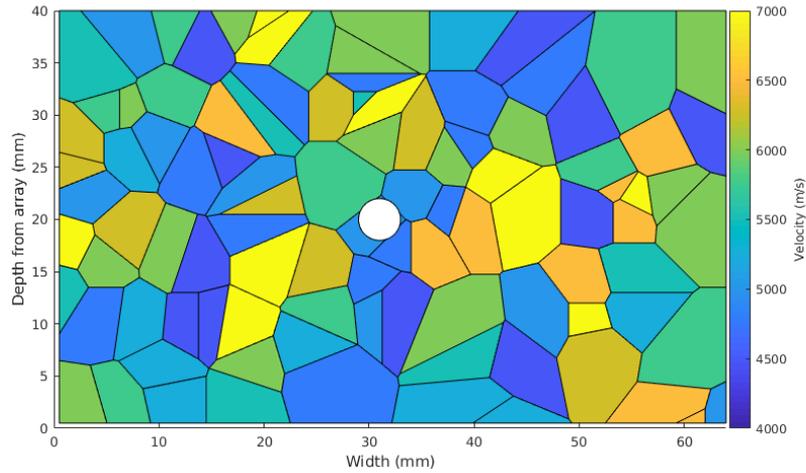


Figure 4.1: Synthetically generated random material diagram with material wave speeds varying from 4000 to 7000m/s. The average wave speed is therefore 5500m/s and this gives rise to an average wavelength of 3.8mm at 1.5MHz. The average Voronoi cell diameter is 5mm and, given the high mechanical impedance mismatches, this material is highly scattering. The white disc denotes a side drilled hole flaw of diameter 4mm simulated as a void. The density is assumed to be constant throughout the host medium.

To enable imaging algorithm comparisons, the standard Total Focusing Method (TFM), Equation (2.21), which assumes that the host material is homogeneous was used to get a reconstruction of the flaw (Figure 4.2). It can be observed that the SDH flaw cannot be detected. This is due to the high degree of scattering within the material as the average wavelength (3.8mm) is commensurate with the average cell diameter (5mm) and there exists substantial inter-grain mechanical impedance mismatches.

When the known wave speed map is used in the TFM+ imaging algorithm (Equation (2.22)) then the flaw reconstruction which results is shown in Figure 4.3. Now the SDH flaw has been detected where it could not be using the

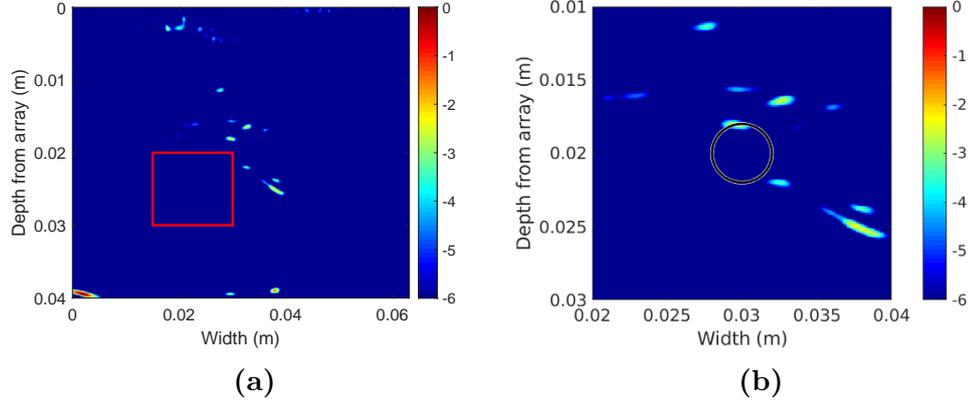


Figure 4.2: Plot (a) shows the TFM results assuming host material is homogeneous. Plot (b) is an enlarged image around the SDH flaw region. Both are plotted with a 6dB dynamic range. No flaw is detected. The black circle indicates where the flaw should be. The red square indicates the noisy region used in the SNR calculation (Equation (2.23)).

standard TFM. There is also another artefact at (0.02, 0.015), this is because all scattering is focused by the known map and this corresponds to the interface with a high variation in wave speed in Figure 4.1.

To quantify the improvement, two factors were looked at: the signal to noise ratio (SNR) and the flaw position error. The SNR was calculated using Equation (2.23) and the maximum flaw intensity value (A_{signal}) within a rectangular region $\Omega=(21-41\text{mm}, 10-30\text{mm})$ centred around the inclusion, and the mean value (A_{noise}) in the rectangle $\Omega_0=(15-30\text{mm}, 20-30\text{mm})$, located some distance from the flaw location (see the red rectangles in Figure 4.2). The error in the flaw position was calculated by finding the centre of mass in the image intensity domain within a 6mm square centred around the known location of the inclusion. The location error is then given by the distance between the known inclusion

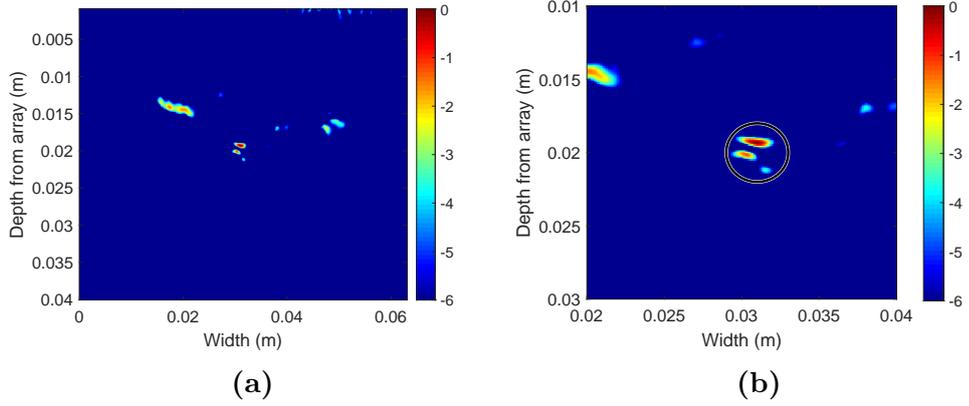


Figure 4.3: Plot (a) shows the TFM+ result accounting for the known material map. Plot (b) is an enlarged image around the flaw region and plotted with a 6dB dynamic range. A flaw is clearly detected. The black circle indicates where the flaw should be.

location and this calculated centre of mass, which gave an error of 0.98mm. The difference in SNR value and flaw location error can be found in Table 4.1 where it can be observed that there is an improvement for both metrics in the TFM+ reconstruction.

The next step is to attempt to obtain an approximation of the material wave speed map using the Full Waveform Inversion (FWI) rj-MCMC algorithm that allows for the detection of the flaw.

Table 4.1: Comparison of the flaw image quality metrics of SNR and Flaw Position Error between Figure 4.2 and Figure 4.3.

	TFM	TFM+
SNR [dB]	13.3	16.3
Flaw Position Error [mm]	1.7	0.98

4.2 Improving Flaw Detection using FWI rj-MCMC

In this section, the heterogeneous material map ?? is reconstructed using the set for the inversion tools described in Chapter 2: a *Voronoi tessellation* to parametrise the spatial domain; the *observed data* generated from OnScale [141]; the bespoke semi-analytical model from Chapter 3 as the *forward model*; an adjusted Pearson Correlation Coefficient (PCC) for the *objective function*; and the rj-MCMC the selected as the *optimiser*.

The rj-MCMC was run for a total of 100,000 iterations, with the first 60,000 being discarded (referred to as the burn-in period). The remaining models were sampled at an interval $\kappa = 100$. A homogeneous model was used to initialise the algorithm, with constant wave speed of 5500 m/s. The algorithm had the following limits on the prior (uniform) distributions: the number of cells must lie in the range of 10 to 300; the wave speed for a Voronoi cell must lie in the range of 5500 ± 1500 m/s; and the noise parameter must lie in the range 0.01 and 1. The proposal distributions σ for wave speed, cell position, birth, and noise perturbations were Gaussian and centred on the values from the previous iteration. The standard deviation for the proposal distributions were: on *wave speed* 100 m/s, and 50 m/s for the case of a delayed rejection (σ^{dr}); on *cell position* 10%, and 2% for the delayed rejection; on the wave speed for a new cell generated by a *birth* perturbation 500 m/s centred at the wave speed value at that spatial position in the previous iteration; on the *noise* parameter 0.05 from the existing noise value. These are summarised in Table 4.2

Table 4.2: Prior and proposal distributions set for the simulation.

Perturbation	Range	σ	σ^{dr}
Wave speed	5500 ± 1500 m/s	100 m/s	50m/s
Seed position	-	10%	2%
Birth	-	500 m/s	-
Noise	0.01 - 1	0.05	-

The average objective function value, over the full 1024 rays, as the algorithm iterates through the samples, is shown in Figure 4.4. This shows a steady decrease until the 60,000th iteration where it levels off at a value of 0.4. The number of cells as the algorithm iterates through the samples is shown in Figure 4.5 (a). This increased steadily until it reached about 250 cells just after the 60,000 iterations. These factors determined the burn-in period used to generate the posterior distributions. The distributions of the number of cells for the last 40,000 iterations is shown in Figure 4.5 (b) which shows a normal distribution with a left tail. This left tail corresponds to the iterations around 60,000 in Figure 4.5 (a).

From the moments and characterisations of the posterior distribution of the wave speed at each point in the domain, a material map was constructed in Figure 4.6, which shows the MAP (maximum a posteriori). The mean and median of the posterior distributions have a similar appearance but with a smoother gradient between wave speeds regions. Using this approximation of the material wave speed map with the TFM+ algorithm, the flaw reconstructions obtained are shown in Figure 4.7. Using the same method as described in Section 4.1, an SNR value was calculated as 15.0dB, and a flaw location error of 0.95mm was obtained.

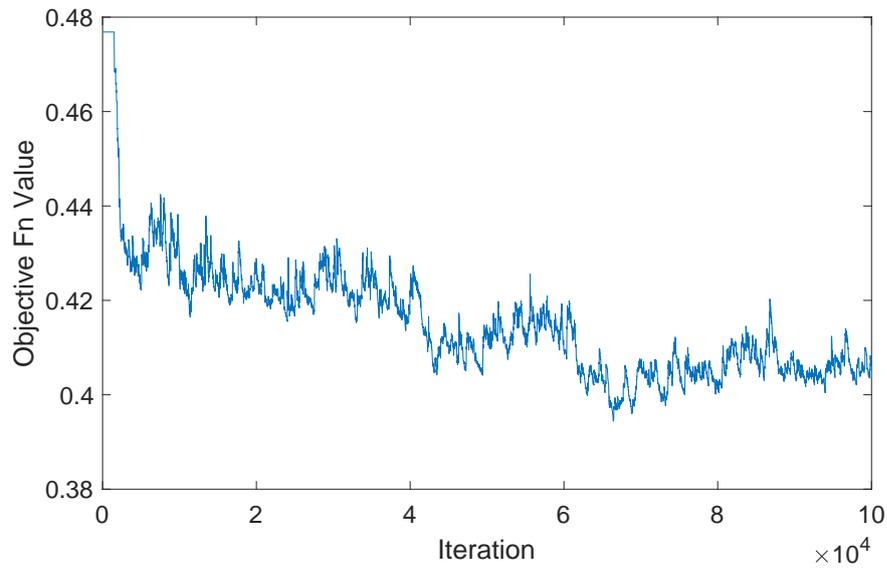


Figure 4.4: The average objective function value as the algorithm iterates through each sample.

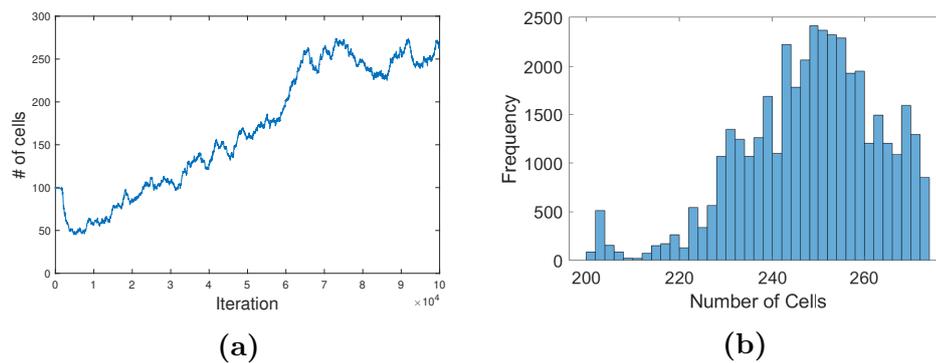


Figure 4.5: Plot (a) is the number of cells at each iteration of the rj-MCMC algorithm. Plot (b) is the distribution of the number of cells for the last 40,000 cells (after the chosen burn-in period of 60,000).

This has a comparable SNR value to that obtained using the known map and a comparable flaw location error. Importantly, it outperforms the standard TFM imaging shown in Figure 4.2 as detailed in Table 4.3.

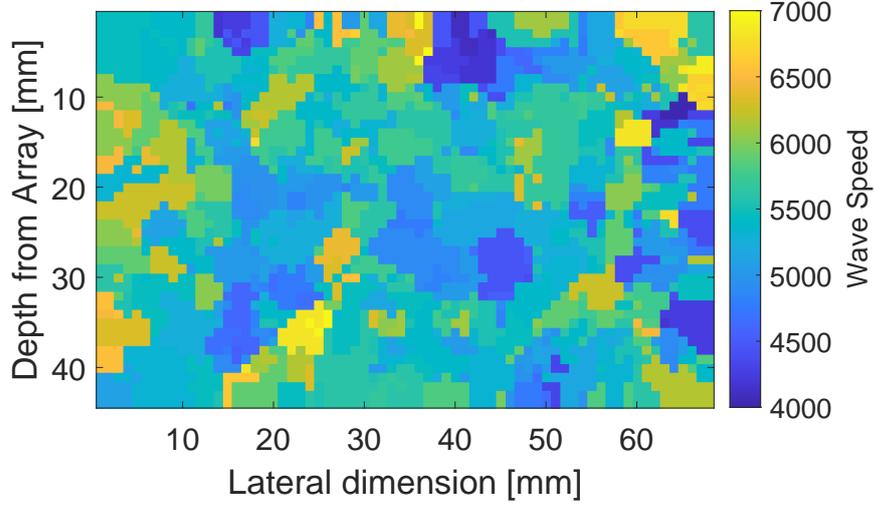


Figure 4.6: The material map obtained from the MAP of the posterior distribution on wave speeds generated by the rj-MCMC (material wave speeds vary from 4000 to 7000m/s.)

Table 4.3: Comparison of the flaw image quality metrics of SNR and Flaw Position Error between Figure 4.2, Figure 4.7, and Figure 4.3.

	TFM	TFM+ FWI	TFM+
SNR [dB]	13.3	15.0	16.3
Flaw Position Error [mm]	1.7	0.95	0.98

The TFM+ image from the known map (Figure 4.3 (a)) was used as the ground truth in the ROC curve generation [41]. This was done by dividing the TFM+ image into a grid. Each grid cell with an image intensity over -4dB was defined as being part of the flaw domain Ω_f and the number of cells in the flaw domain is recorded as n^p . The total number of cells not in the flaw domain is denoted by n^f . The standard TFM (Figure 4.2) and the TFM+FWI image (Figure 4.7) are then partitioned using the same grid. The probability of detection (PoD)

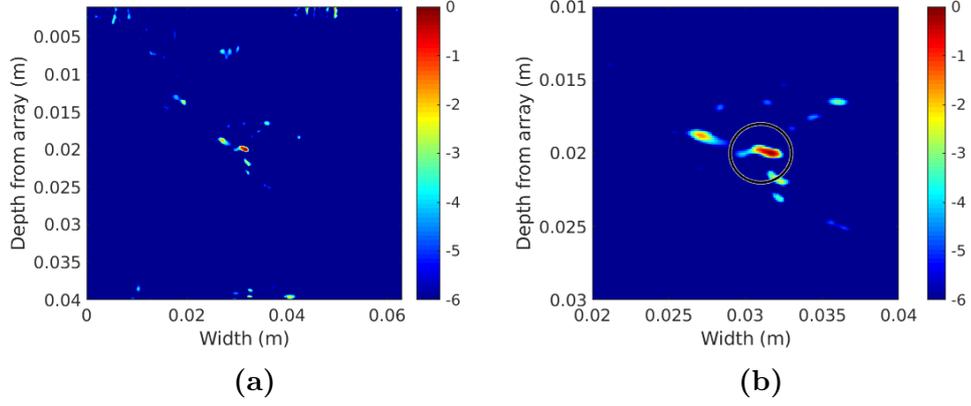


Figure 4.7: Plot (a) is the TFM+ image arising from the MAP of the posterior distribution Figure 4.6. Plot (b) has been cropped around the flaw region. The black circle indicates the known location of the flaw.

value is calculated by n^+/n^p and the false positive rate (FPR) value is given by n^-/n^f , where n^+ is the number of grid cells which have a value above some threshold D and lie within the flaw domain Ω_f and n^- is the number of grid cells with an image intensity value above this threshold that are not in Ω_f . This is repeated from the initial threshold of $D = -1\text{dB}$ down to -40dB and Figure 4.8 was obtained. It clearly shows that the flaw reconstruction TFM+ FWI greatly outperforms the standard TFM, with an area under the curve (AUC) value of 0.872 (Table 4.4).

Table 4.4: Comparison of the AUC values from Figure 4.8.

Method	AUC
TFM	0.65808
TFM+ FWI	0.87185

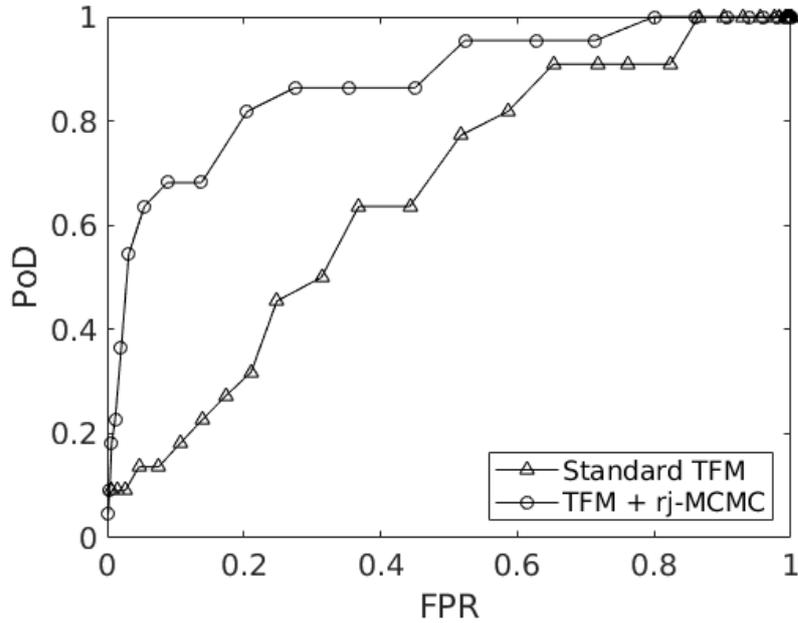


Figure 4.8: ROC curve generated from the standard TFM image (Figure 4.2) and the TFM+FWI image (Figure 4.7) using the TFM+ image from the known map (Figure 4.3) as the ground truth

From these results, it can be deduced that the objective of improving flaw detection has been met. A flaw is now clearly detectable. The next step is to compare these results with the results from [41] to see if this method performs better than ToF.

4.3 Comparing Full Waveform Inversion to Time of Flight

In [41], there were two methods of obtaining a reconstruction of the material wave speed map. The first used a through transmission (TT) set-up (pitch-catch), and the second used a pulse-echo (PE) set up. The first comparable metric given is

the error in the flaw location. As in Section 4.1 the distance between the location of maximum flaw image intensity and the known centre of the flaw are used to quantify the error. Using this method, the FWI TFM+ image has a flaw location error of 1.2 mm which is comparable to the ToF (TT) TFM+ image and better than the ToF (PE)+ TFM+ image, as shown in Table 4.5.

The next metric that is used to compare the abilities of each of the material reconstruction methods is the area under the curve (AUC) of the associated receiver operating characteristics (ROC) curves. The generation of the ROC curves is described in Subsection 2.1.9 and Section 4.2. The perfect imaging algorithm would reconstruct the flaw in precisely the correct position, with the correct size and with no other “false” flaws appearing in the image. The AUC can take values between zero and one and in such instances it would be one. Therefore, the higher the values of the AUC are the better the flaw imaging algorithm. The reconstructed material maps from the Full Waveform Inversion (FWI) method as the basis for the image processing algorithm (TFM+) to calculate the AUC(FWI). This value is then compared with that obtained from the standard TFM AUC(TFM), and values from the literature where two different Time of Flight (ToF) methods to reconstruct the material maps pulse-echo AUC(ToF -

Table 4.5: Comparison of the flaw location errors between the reconstruction in Figure 4.7 and the results found in [41].

Basis of TFM+	Flaw location error [mm]
FWI	1.2
ToF (TT)	1.2
ToF (PE)	1.8

PE) and through-transmission AUC(ToF - TT) [41]. As before the TFM+ arising from the known map was used as the ground truth to generate the ROC curves. In [41], the AUC value was 0.539 for the standard TFM, 0.651 for ToF (PE), and 0.703 for ToF (TT), as summarised in Table 4.6. FWI, with an AUC of 0.872, outperforms both forms of ToF. However, when comparing the standard TFM, the results differ. Therefore, an adjustment was implemented to compare like for like, by multiplying the FWI AUC by $0.538/0.658$ to give a value of 0.71343, which still outperforms both the ToF methods.

The next metric is the area under the curve (AUC) of the associated ROC curves. The AUC value for (a) the standard TFM, the TFM+ using the reconstructed maps from, (b) through transmission set-up for time of flight ToF (TT), (c) pulse-echo set up for time of flight ToF (PE), and (d) the full waveform inversion, are all compared. As before the TFM+ arising from the known map was used as the ground truth to generate the ROC curves. In [41], the AUC value was 0.539 for the standard TFM, 0.651 for ToF (PE), and 0.703 for ToF (TT), as summarised in Table 4.6. FWI, with an AUC of 0.872, outperforms both forms of ToF. However, when comparing the standard TFM, the results differ. Therefore, an adjustment was implemented, to compare the AUC from TFM+ FWI and TFM+ (TT) and (PE), by multiplying the FWI AUC by $0.538/0.658$ to give a value of 0.71343, which outperforms both the ToF methods.

Table 4.6: Area under the curve for the ROC curves generated from ToF [41].

Method	AUC
TFM	0.53850
TFM+ (TT)	0.70342
TFM+ (PE)	0.65071

From the AUC and flaw location error, it can be seen that the FWI method for material reconstruction performed just as well as the ToF TT method. Both these methods had the same array set up. The PE method, which is more industrially relevant, uses a single phased array which a FWI approach has not been developed yet.

4.4 Posterior Distribution vs a Randomly Generated Heterogeneous Map

To confirm that the MAP of the posterior distribution on wave speeds as shown in Figure 4.6 is a good representation of the synthetically generated random material in Figure 4.1, and does not improve the flaw reconstruction by chance, a new random Voronoi diagram containing 100 seeds (such that the Voronoi cell sizes are approximately the same size), and wave speeds varying from 4000 - 7000 m/s uniformly, was generated. Using this randomly generated map within the TFM+ algorithm, the flaw reconstructions can be found in Figure 4.9. From this figure it can be seen that a flaw is not detectable.

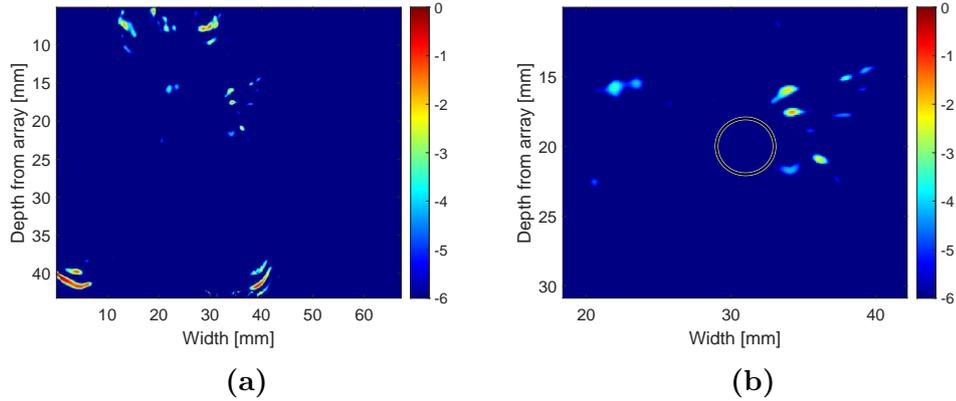


Figure 4.9: TFM+ image arising from the randomly generated heterogeneous material, plot (a). Plot (b) has been cropped around the flaw region. The black circle indicates where the flaw should be.

When looking at the quantitative measures of how this flaw reconstruction performs, it has an SNR value of 15.7dB, a flaw location error of 2.2 mm (when measuring from the centre of mass (CoM)), a flaw location error of 8.1 mm when measuring from the maximum intensity (MI) value, and a ROC AUC value of 0.622 (see Figure 4.10). These values are summarised and compared to the performance of the standard TFM and TFM+ FWI in Table 4.7, where it can be evaluated that the flaw reconstruction using the random Voronoi in the TFM+ algorithm does not perform any better than the standard TFM and in some measures does worse.

From this, it can clearly be seen that the flaw reconstruction using the characteristics of posterior distribution from the FWI has achieved its objective of improving flaw detection and characterisation. However, it did not perform any better than ToF TT approach.

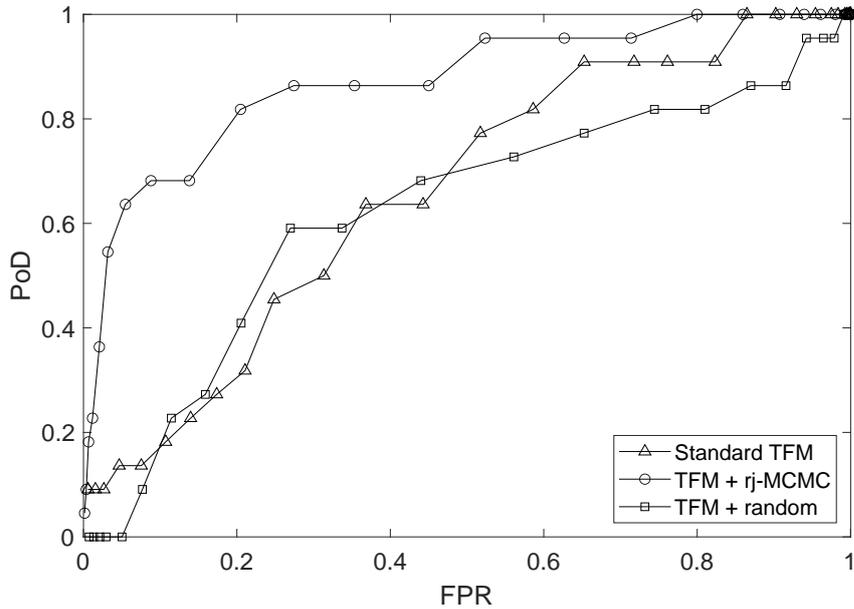


Figure 4.10: ROC curves using the TFM+ from the known map as the ground truth. Comparing the standard TFM, TFM+ FWI, and TFM+ from the random Voronoi.

Table 4.7: Comparing the qualitative metric (SNR, flaw error, and AUC) between the standard TFM, TFM+ FWI, and TFM+ random material map.

Method	SNR	Flaw error (CoM) [mm]	Flaw error (MI) [mm]	ROC AUC
TFM	13.3	1.7	2.1	0.65808
FWI TFM+	15.0	0.95	1.2	0.87185
Random TFM+	15.7	2.2	8.1	0.62212

4.5 Discussion

The work done in [41] gave the foundation for the work done in this chapter. It provided a scenario where a flaw cannot be detected using the standard TFM imaging algorithm but can be detected with TFM+ imaging algorithm with the known wave speed map, and with an approximation of the wave speed map from

the FWI-rj-MCMC inversion algorithm. This led to two clear objectives for this chapter. The first objective was to improve flaw detection and characterisation through full waveform inversion (FWI). This objective was met successfully. The second objective was to outperform the ToF based material map reconstruction combine with the TFM+ imaging algorithm from [41]. This objective was not met, however, the FWI approach did not perform any worse than ToF.

To gain insight into the FWI algorithm it seemed sensible to start by using the semi-analytical model (which is the forward model in our inverse problem setup) as the observed data (also known as the inverse crime). This meant that what was being compared in the objective function had identical physics. This allowed for additional scrutiny of the algorithm at every step, the code used to produce the semi-analytical model, and the objective function.

On doing so it was discovered that the initial peak of the A-scan dominates the objective function value and the tail of the A-scan was obscured so this essentially equivalent to the time of flight method. This explains why the work presented in this chapter successfully reconstructs the flaw but does not improve upon ToF.

Chapter 5

Tomographic Reconstruction of a Layered Medium

A forward model for the time domain transmitted ultrasound wave travelling through a piecewise constant material has been developed for one spatial dimension in Chapter 2. This is analogous to the case of an incident plane wave whose direction of propagation is orthogonal to the interfaces of a two dimensional layered medium. For this reason a layered medium consisting of four parallel, isotropic layers, of infinite extent in the lateral direction, was studied (see Figure 5.1). The material properties can be found in Table 5.1. The shear wave and density were chosen to be constant throughout the medium to simplify the model. The reconstruction of this medium could be the initial steps to solving the inverse imaging problem of carbon fibre composite materials that are increasingly used within industry [24].

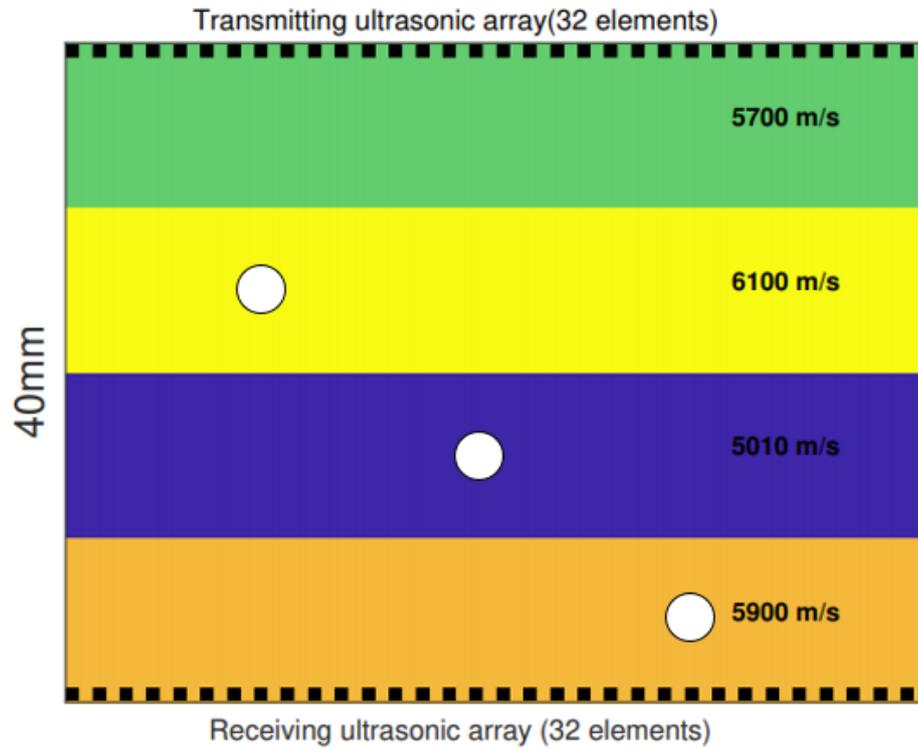


Figure 5.1: Schematic of a layered medium with 4 layers each 10mm in thickness, with wave speeds 5700m/s, 6100m/s, 5010m/s, and 5900m/s respectively. Two 32 element ultrasonic phased arrays with element pitch 2mm are shown at the upper and lower surfaces as shown. Three flaws (inclusions), with 4mm diameter, positioned at (16mm, 15mm), (32mm, 25mm), and (48mm, 35mm), each with longitudinal wave speed of 5300 m/s and 2260kg/m³ density.

Table 5.1: Material properties of the layered medium shown in Figure 5.1

Layer	Longitudinal Wave Speed [m/s]	Shear wave speed [m/s]	Density [kg/m ³]
1	5700	3200	7890
2	6100	3200	7890
3	5010	3200	7890
4	5900	3200	7890

5.1 TFM of Layered Medium

To illustrate the benefits of recovering the medium properties, a situation was chosen where the standard imaging algorithm (TFM) (where the material is assumed to be homogeneous) fails to satisfactorily image a flaw but the TFM+ (where the heterogeneous material map is accounted for) improves the image to a more satisfactory level. Three 4mm diameter inclusions were embedded at locations (16mm, 15mm), (32mm, 25mm), and (48mm, 35mm), each with longitudinal wave speed 5300m/s and density 2260kg/m³, in the sample depicted in Figure 5.1. The FMC data was simulated with the finite element package On-Scale [141] using a phased array in the pulse-echo format. The array had a pitch (element spacing) of 2mm, and emitted a Ricker wavelet function (see Figure 3.4) with a centre frequency of 1.5Mhz. The simulation was run on a domain with element size $\lambda/15$, where the wavelength λ is calculated using the minimum wave speed present in the model (here the shear wave speed). The time duration for which the model was run was 50 μ s chosen to obtain A-scans that would have sufficient duration for the tomographic inversion; the wave takes roughly 8 μ s to traverse the layered medium and this ensured that multiple reflections were captured. The other parameters within the finite element package are functions of these parameter values. Free boundary conditions were implemented at the backwall and absorbing boundary conditions were used in the lateral direction.

The standard TFM Equation (2.21) (with a constant velocity of 5677m/s, given by the mean wave speed for the four layers) was applied to the FMC data and the resulting image is shown in Figure 5.2. Delay law corrections were then applied using the known map to generate the TFM+ image (Equation (2.22)) shown in Figure 5.3.

Comparing the images of the inclusions in Figure 5.2 and Figure 5.3, it can be observed that the flaw closest to the back wall is better resolved when knowledge of the layered material is incorporated. In the standard TFM image this inclusion appears to have the same amplitude as the surrounding noise, whereas in the TFM+ image the flaw is clearly visible. By taking into account the varying wave speeds, the TFM+ algorithm more accurately locates the correct time stamp within the A-scans. This means the peaks within the A-scans associated with the flaw are calculated and the flaw is more accurately imaged. To quantify the improvement two factors were looked at: the signal to noise ratio (SNR) and the flaw position. The SNR was calculated using Equation (2.23) where the maximum flaw intensity value (A_f) was selected from the rectangular region Ω drawn around the inclusion, (44-54mm, 32-38mm), and the mean noise value (A_n) in the rectangle Ω_0 of the same dimensions in the mirrored position, (12-20mm, 32mm-38mm) - see the red rectangles in Figure 5.2. The flaw position is calculated by finding the maximum intensity within a 6mm square centred around the known location of the inclusion within the image. The location error is then given by the distance between the known inclusion location and this calculated maximum intensity. These values are summarised in Table 5.2. It can be

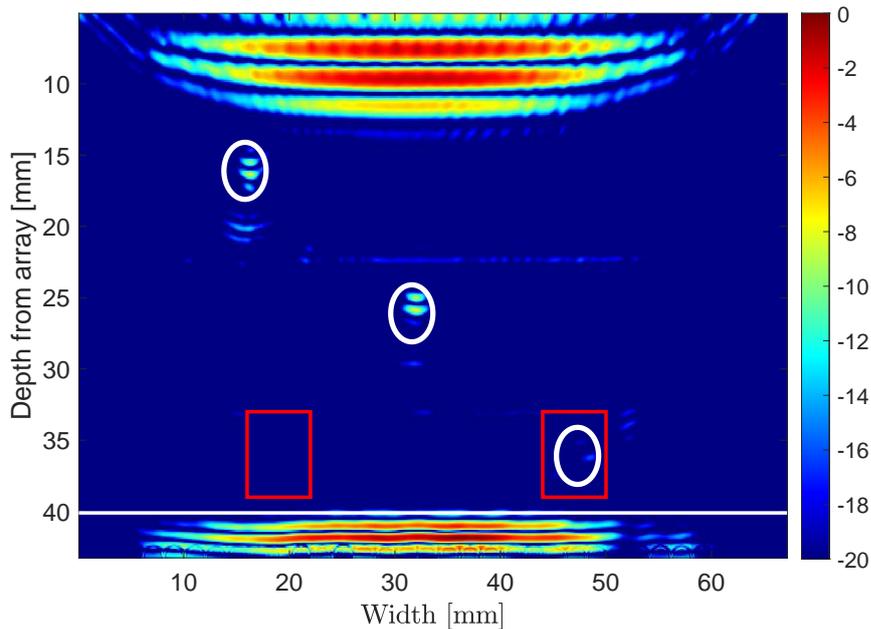


Figure 5.2: The TFM image created using a homogeneous host material wave speed of 5677m/s, arising from a finite element FMC dataset for the heterogeneous material shown in Figure 5.1. Three inclusions (white circles), with 4mm diameter were embedded at locations (16mm, 15mm), (32mm, 25mm), and (48mm, 35mm). The white line indicates the back wall and the red rectangles indicate the regions used to calculate the SNR values.

Table 5.2: Comparison of the flaw image quality metrics of SNR and Flaw Position Error.

	TFM	TFM+
SNR (dB)	17.5	23.6
Flaw 1 Position Error [mm]	0.283	0.985
Flaw 2 Position Error [mm]	1.523	0.633
Flaw 3 Position Error [mm]	1.315	1.300

observed that there is an improvement of 6.1dB in the SNR with the TFM+ and marginal improvement of 0.89mm, and 0.015mm in the flaw positions of the middle and bottom flaw, with the standard TFM having a better flaw location

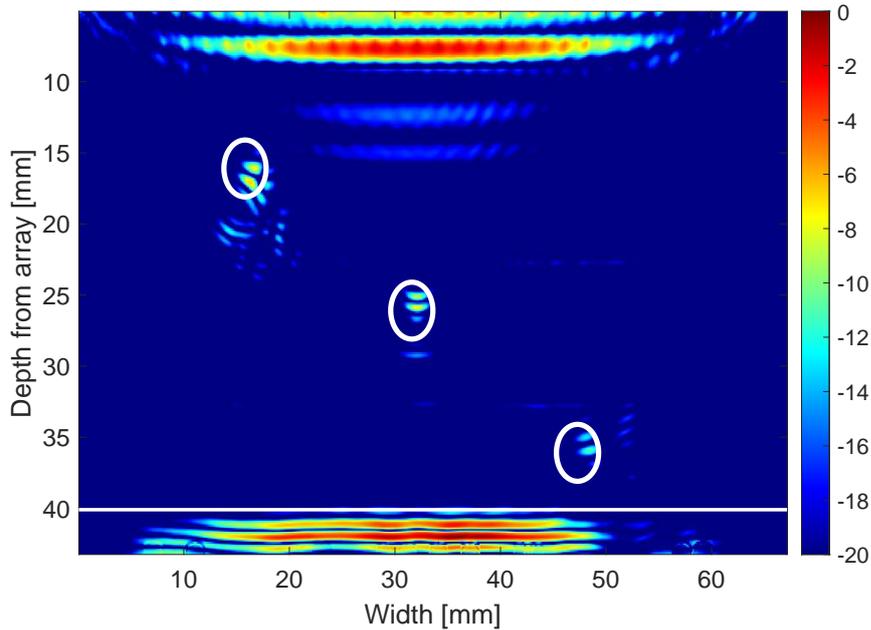


Figure 5.3: The TFM+ image arising from the finite element FMC data for the scenario shown in Figure 5.1, with three inclusions with 4mm diameter embedded at locations (16mm, 15mm), (32mm, 25mm), and (48mm, 35mm), indicated by the white circles, with material constants, longitudinal wave speed 5300m/s and density 2260kg/m³.

error for first flaw. These flaw location errors, for both the standard TFM and TFM+ with the known map, are within the flaw domain. However, it is possible for the flaw location to outperform the TFM+, as with the top flaw here and in Subsection 6.2.1 and Subsection 7.3.2. This issue arises from the simplicity of the medium. It could, therefore, be inferred that this is not a good metric of quality, however, it is part of the framework for improvement in flaw detection, position, and characterisation that would hopefully be implemented for more complex heterogeneous material

As the material is piecewise homogeneous, the wave speed changes between each layer are relatively small, and the ultrasonic array has a large aperture and many elements, then the TFM imaging does a reasonably good job in this case. As such, there is not a great deal of scope for the TFM+ imaging approach to improve upon. Nevertheless, this first scenario will be used as a test case to assess the level of imaging improvement gained when the TFM+ algorithm uses the reconstructed material maps from the rj-MCMC simulations.

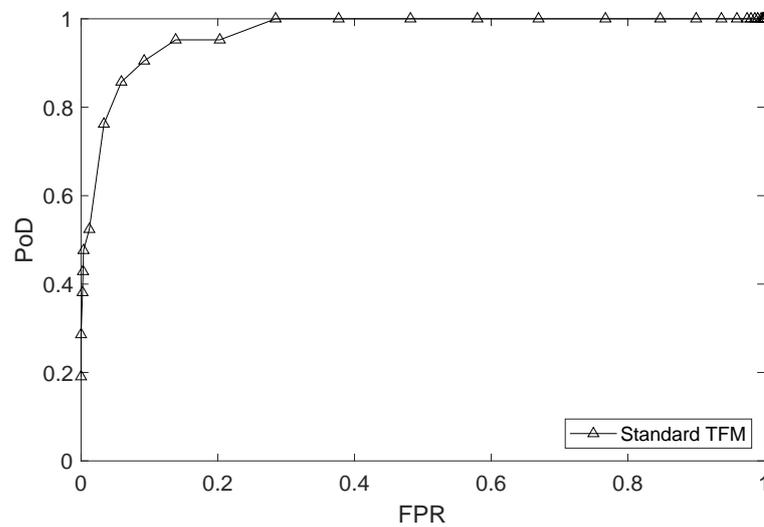


Figure 5.4: ROC curve for the images arising from the standard TFM (Figure 5.2) using the image from the TFM+ (Figure 5.3) as the ground truth.

Table 5.3: Comparison of Area under the curve between the basic TFM and TFM+ with known map using the binary map as the ground truth.

Method	AUC
TFM	0.8304
TFM+	0.9012

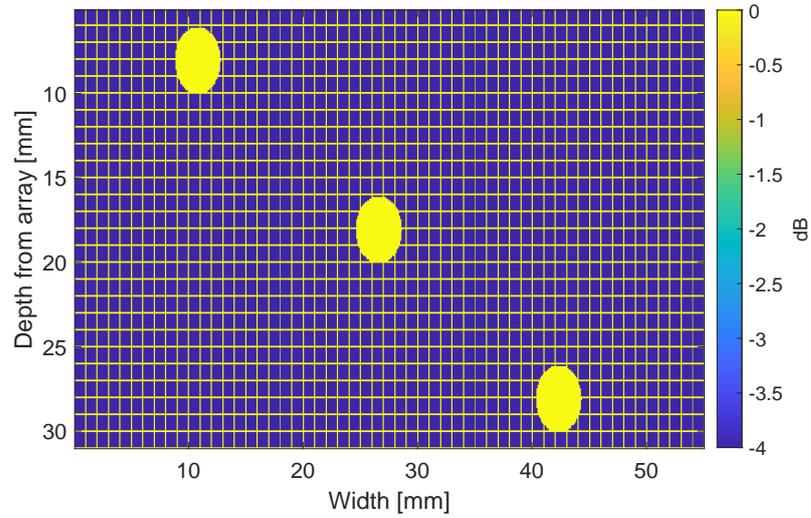


Figure 5.5: A binary map, divided into a grid, of the three flaws present within the medium used as the ground truth for the ROC curves.

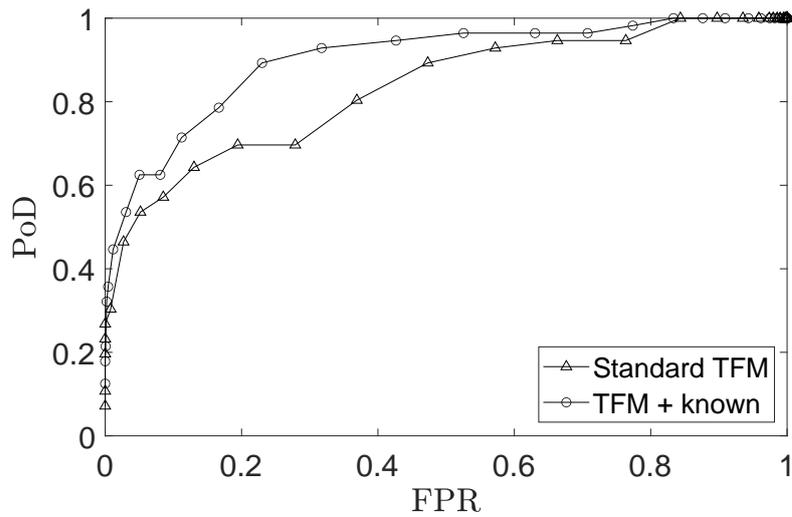


Figure 5.6: ROC curves for the images arising from comparison of the standard TFM (Figure 5.2) and TFM+ (Figure 5.3) with the known map to a binary map representing the exact locations and dimensions of the flaws.

The ROC curve (see Figure 5.4) provides the probability for detection of the inclusions present in the images arising from the standard TFM Figure 5.2 using TFM+ Figure 5.3 with the known map as the ground truth. It can be observed that there is very little room for improvement with an AUC of 0.969. The reason for this is that the main difference in the inclusions between standard TFM and TFM+ is the intensity amplitude and not the shape of the inclusion. The TFM imaging algorithms do not image the inclusions but the reflections caused by the inclusions. Each inclusion, within a TFM image, is represented by three components. These three components are the reflection from the upper surface of the inclusion, the reflection from the lower surface of the inclusion, and an internal reflection within the inclusion. An alternative approach is to use a binary map (Figure 5.5) as the ground truth for the ROC curves. In these binary maps, the whole inclusion is depicted by a dB of 0. This means that the greater intensity present in the TFM+ algorithm has an improved PoD, as one would expect. The respective area under the curves (AUC) are recorded in Table 5.3. The question then arises if a reconstructed map is able to correctly identify the inclusions more successfully than the standard TFM.

5.2 Setting up and Initiating the rj-MCMC Algorithm

The first step in reconstructing a medium is to first generate the experimental or observed data. This could be data collected from an ultrasonic array in a lab, data from a finite element simulation, or data generated from the forward model

that will be used as part of the inversion algorithm (colloquially known as the inverse crime). For this chapter, the data is generated within MATLAB using the forward model in the transmission (pitch-catch) set up (where two arrays are placed on either side of the sample, directly opposite each other).

To generate the data there are three necessary inputs. Firstly, a waveform file which determines the transmitted wave. In this work the Ricker wavelet function, (a predefined waveform) with centre frequency 1.5MHz, generated in OnScale was used (see Figure 3.4). Secondly, the timestep used by OnScale when generating the waveform; a timestep of 1.688120e-08 seconds was used here. Thirdly, a file containing cell seed positions and cell wave speeds to enable the generation of a Voronoi tessellation that will subsequently attempt to reconstruct. To create a Voronoi tessellation comprising of four horizontal layers, eight Voronoi cells were used (since the algorithm used to generate the Voronoi tessellation required a minimum of 5 cells). There are two in each layer and the cell seeds both have the same depth position (y-coordinate) and cell wave speed are set as in Table 5.1. Four of the cell seeds have the same lateral position (x-coordinate) and the other four again have the same lateral position but mirrored from the first four to create a Voronoi tessellation in the form of a rectangular grid (see Figure 5.7). From this geometry a MATLAB algorithm determined the raypaths for all combinations of sources and receivers positioned as in Figure 5.1. For each raypath, the number of layers that the wave passed through, the associated layer thicknesses, and wave speeds of each layer were fed into the forward model algorithm, generating an FMC of forward model generated A-scans.

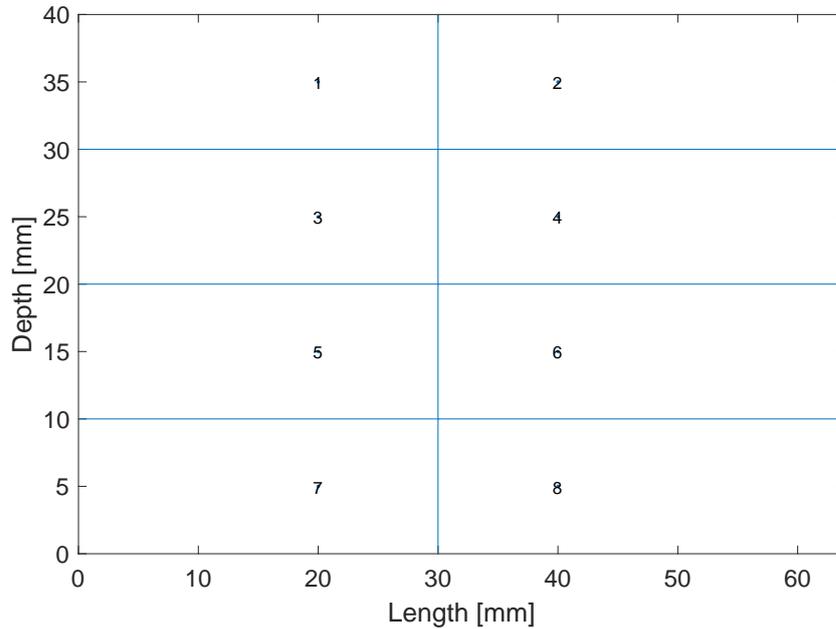


Figure 5.7: The eight cell Voronoi tessellation used to create the homogeneous initial Voronoi and the known 4 layer medium.

The *rj-MCMC* algorithm is initiated using a few set up files. These are: a *sources* file containing the positions of the source array elements (in this example the array was placed at a depth of 0mm and started at 1mm laterally and each element was placed at 2mm intervals (pitch), with element 32 being placed at 63mm); a *receiver* file containing the positions of the receivers (this had matching lateral positions as the source files but at a depth of 40mm), and an *observed ray* file containing which raypaths are to be modelled (this allows for only part of the dataset to be used if required). From these files, raypaths are generated and stored in a *raypath* file which contains the coordinates of each step of all raypaths. A *general input* input file containing all the prior and proposal distributions, and the MCMC run parameters. Lastly, (optionally) an *initial Voronoi* file which is

the initial Voronoi tessellation. A random Voronoi tessellation can also be used for the initial geometry and is generated using the prior distributions on the material parameters. Lastly, there are files that contain the FMC data, timestep value, and the transmitted waveform. Once these files are input all that is required is the number of iterations the rj-MCMC algorithm will use.

5.3 Results using the PCC Objective Function

To begin, an inversion from the full FMC dataset was attempted using the Pearson Correlation Coefficient (PCC) objective function as presented in Subsection 2.1.3. A homogeneous (a Voronoi tessellation composed of eight Voronoi cells each having a wave speed of 5600 m/s) was used to initialise the rj-MCMC algorithm. Unfortunately, the algorithm was unable to reconstruct a layered medium and the resulting material maps obtained by examining moments of the posterior distribution on wave speeds were homogeneous (see Figure 5.8).

On further investigation, this was in part due to the initial received wave packet dominating the correlation coefficient calculation. The objective function was insensitive to the part of the A-scan which contained the reflected waves arising from multiple reflections by the layer interfaces. To increase the objective function sensitivity to these reflections the logarithm of the A-scans was used and the results are presented in the next section.

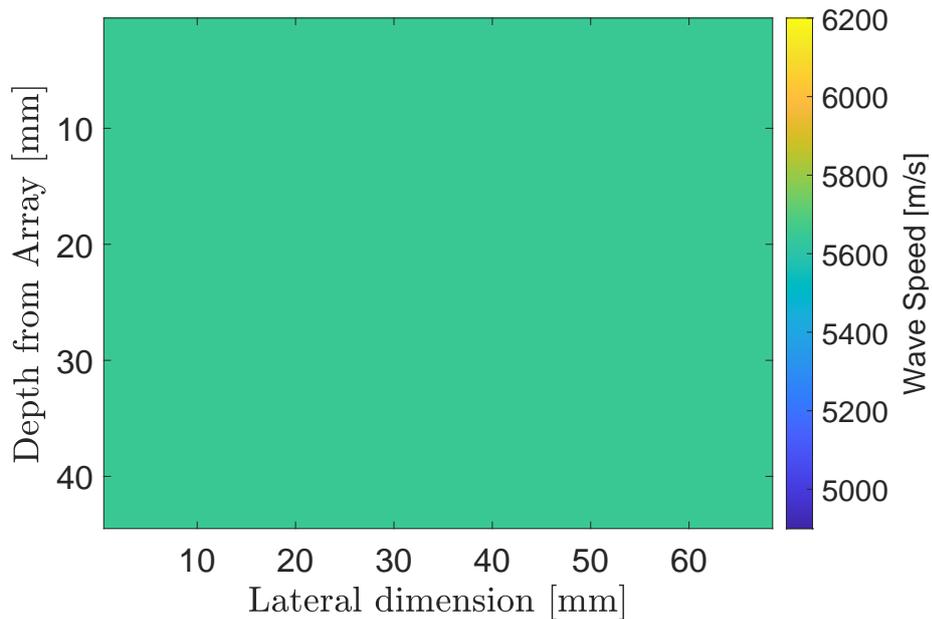


Figure 5.8: Material map approximation constructed from the median of the posterior distribution on wave speed at each point in the domain generated using the PCC based objective function within the rj-MCMC framework.

5.4 Reconstruction of a Layered Medium using 32 rays from the Full Matrix Capture

The medium to be reconstructed is a layered medium and this would suggest that, in theory, a single vertical ray would be sufficient to invert for this 1D structure [117, 118, 119]. However, to keep the algorithm generally applicable to 2D scenarios and for example, evenly layer structures, the case where 32 vertical rays is studied and examined whether these are sufficient to reconstruct the 2D layered structure.

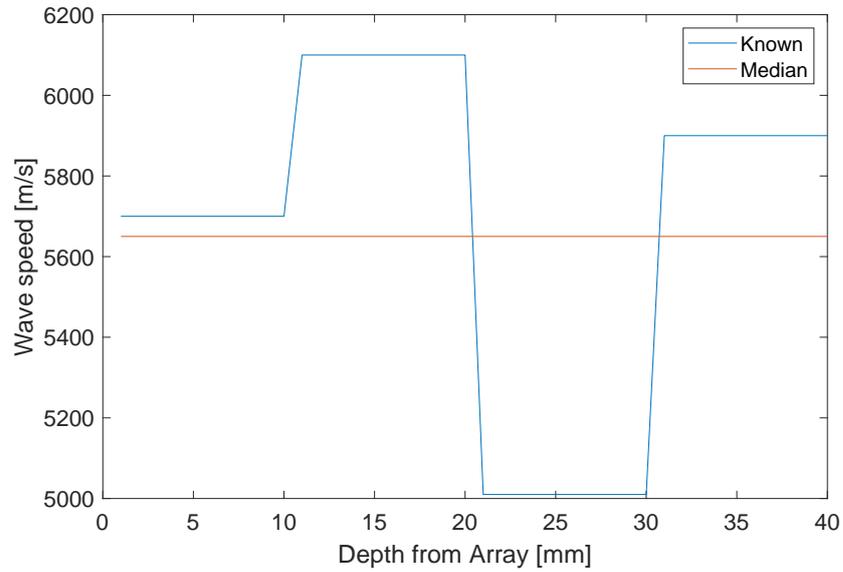


Figure 5.9: Wave speed profiles of the known map and the reconstructed map shown in Figure 5.8 along the vertical cross-section at $x = 32\text{mm}$.

Two scenarios were considered. Firstly, a pulse-echo set-up, where the transmitting array is also the receiving array. Secondly, a pitch-catch set-up, where the top surface array transmits and the bottom surface array receives. This scenario is repeated with the rj-MCMC noise parameter constrained. These two different pitch-catch scenarios are denoted as Case A (with a free noise parameter) and Case B (with constrained noise parameter).

5.4.1 Reconstruction using a Pulse-Echo set up

Using the vertical rays in a Pulse-echo set-up requires the use of the reflected wave as given by Equation (3.5). Once again the rj-MCMC algorithm was unable to obtain a posterior distribution map that contained any features of the layered medium. As before the first reflected wave dominated the objective function.

The reflection model Equation (3.5) is a 1D model and therefore can only be applied to the vertical rays that transmit and receive on the same element. Rays that transmit from one element and received on another would require a 2D model. It is theoretically possible to mathematically define the raypath within a layered medium, however, for a random heterogeneous model, as expected within

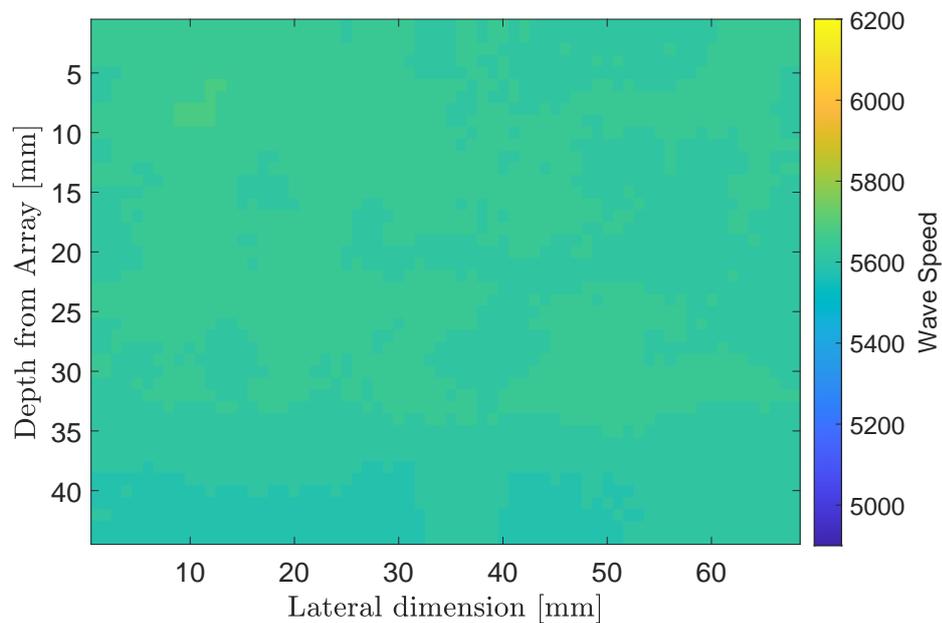


Figure 5.10: Material map generated from the median posterior distribution from the scenario where 32 rays were used with the pulse-echo set-up Equation (3.5).

the rj-MCMC algorithm, it is not possible. This is a limitation in the work that would need to be overcome for the work to be applied in experimental cases. There are situations during ultrasonic NDT that access to only one side of the material is possible e.g. the examination of a pipe. In the through transmission (TT) set-up it is possible to model all A-scans from each transmit/receive pair of phased array elements. This allows for greater coverage of the medium being examined and greater reconstructions. In certain instances (such as those found at the manufacturing stage of components) it is possible to have access to both sides of a component. For this reason this is the set up used for the remainder of the thesis.

5.4.2 Reconstruction using a Pitch-Catch set up - Case A

For the pitch-catch case when only 32 A-scans (vertical rays) were utilised. The rj-MCMC was run for a total of 3,000,000 samples, where the first 1,500,000 were discarded (the burn-in period), this was to ensure the initial homogeneous model did not affect the posterior distribution. The remaining models were sampled at an interval of $\kappa = 100$. A homogeneous model was used to initialise the algorithm with each Voronoi cell having a wave speed of 5600m/s, as this was the median of the prior distribution on wave speeds. The algorithm had the following limits on the prior (uniform) distributions: the number of cells must lie in the range of 5 to 150; the wave speed for a Voronoi cell must lie in the range 5600 ± 700 m/s; and the noise parameter must lie in the range 0.002 and 0.5. The proposal distributions

(σ) for wave speed, cell position, birth, and noise perturbations were Gaussian and centred on the values from the previous iteration. The standard deviation for the proposal distributions were: on *wave speed* 200m/s, and 50m/s for the case of a delayed rejection (σ^{dr}); on *cell position* 8% and 2% for the delayed rejection; on the wave speed for a new cell generated by a *birth* perturbation 150m/s at that point in the previous iteration; on the *noise* parameter 0.008 from the existing noise value. These are summarised in Table 5.4.

With these proposal distributions, the algorithm had an overall acceptance rate of 60.8%. The acceptance rates were, 88.8% for the wave speed, 53.4% for the position, 53.4% for the birth and death of a cell, and 47.6% for the noise. These acceptance rates are high; the literature recommends that for efficient sampling of the model space these should lie within the range of 23% to 40% [41]. This high acceptance rate remained robust to changes in the standard deviations of the proposal distributions.

The average objective function value, Equation (2.1), for each model as the algorithm iterates through the samples, is shown in Figure 5.11, with an initial objective function value of 0.449 and it levels off around 0.27 relatively quickly.

Table 5.4: Prior and proposal distributions set for the simulation for 32 rays Case A.

Perturbation	Range	σ	σ^{dr}
Wave speed	5600 \pm 700 m/s	200 m/s	50m/s
Seed position	–	8%	2%
Birth	–	150 m/s	–
Noise	0.002 - 0.5	0.008	–

The number of Voronoi cells was limited to 150 (see Figure 5.12 (a)). No matter what the upper limit of the number of cells Voronoi cells was set to, the algorithm increased to this upper limit, similar to Figure 5.15. This suggests that the method is not converging, most likely, due to the fact that the problem is non-unique and the 32 vertical rays does not provide enough information to constrain the problem uniquely. The moments of the posterior distributions, Figure 5.13, show that no texture or material characteristic were reconstructed. The mean and median Figure 5.13 (a) and (b) are effectively homogeneous maps around the mean wave speed of 5600m/s. The maximum a posteriori (MAP) distribution (c) is a heterogeneous map with a wave speed of predominately 6200m/s. After observing these material reconstructions it was proposed that the noise parameter should be limited. This is done in the next subsection and referred to as Case B.

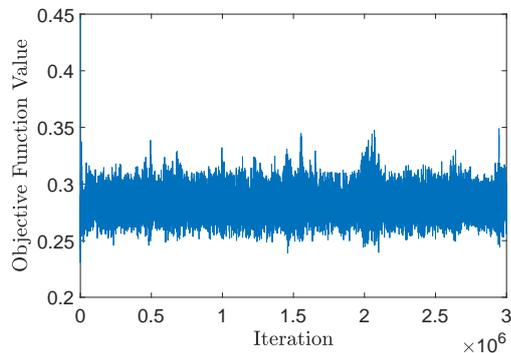


Figure 5.11: The average objective function value as the algorithm iterates through each sample. For the simulation using 32 rays in the pitch-catch setup - Case A

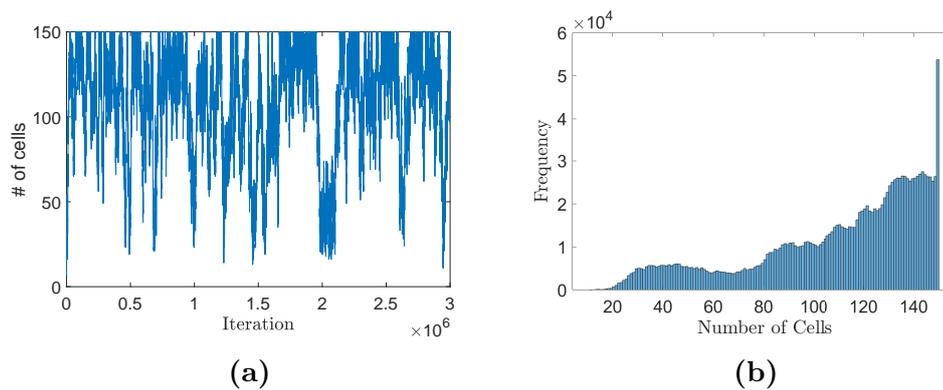


Figure 5.12: Plot (a) is the number of cells at each iteration of the rj-MCMC algorithm and plot (b) is the distribution of the number of cells across all iterations. The maximum number of cells being limited to 150. For the simulation using 32 rays in the pitch-catch setup - Case A

5.4.3 Reconstruction using a Pitch-Catch set up Without the Noise Variable - Case B

In this section, the noise parameter is limited. The main function of the noise parameter is to parametrise uncertainty and prevents data overfitting or underfitting. Since the simulated FMC data $\mathcal{A}_{i,j}^E$ was created from the same forward model used in the inversion algorithm the data is noiseless. Therefore the level of uncertainty is minimum.

The algorithm was run again with the same parameters but this time the noise parameter had a range between 0.0001 and 0.001, with a standard deviation of 0.015. Now that the standard deviation of the noise parameter was far greater than the range it will have 0% acceptance rate and always be near zero. With

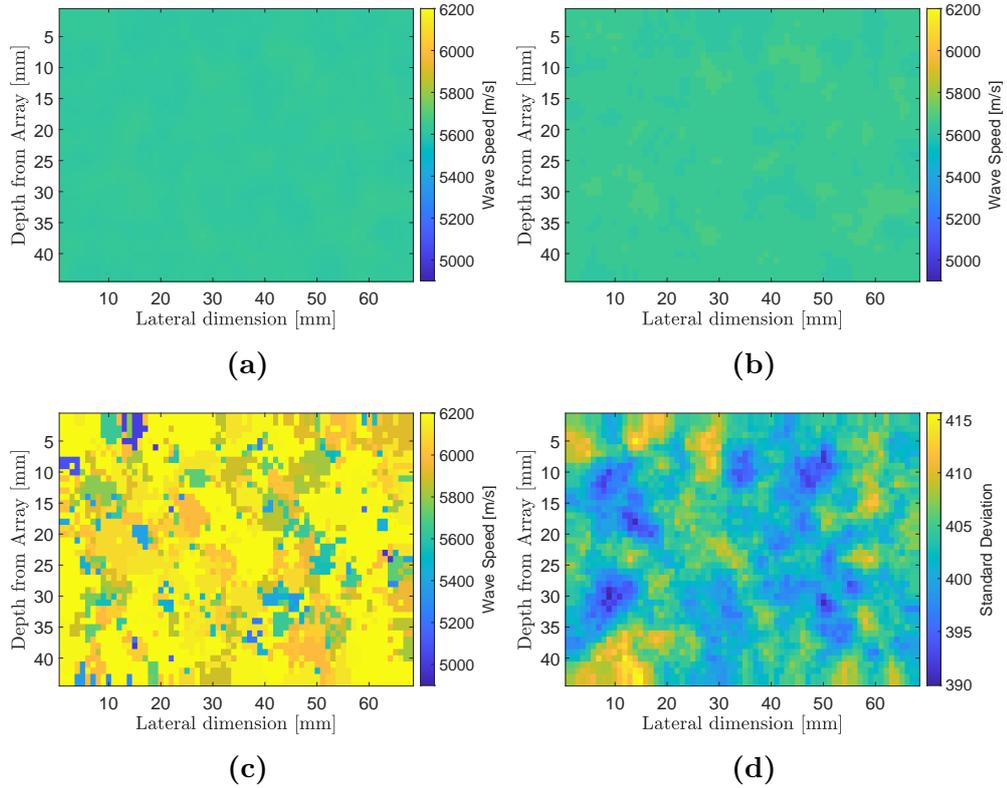


Figure 5.13: Material map reconstructions of the 4 layered medium for the simulation using 32 rays Case A. Plot (a) is the mean of the posterior distribution on wave speed at each point in the domain whilst plot(b)-(d) show the median, MAP, and the standard deviation respectively.

these priors, the algorithm had an overall acceptance rate of 21.2%. The acceptance rates were: 40.9% for the wave speed; 37.7% for the position; 14% for the birth and death of a cell.

The objective function value, Equation (2.1), for each model as the algorithm iterates through the samples, is shown in Figure 5.14; it exhibits an initial objective function value of 0.45 and it levels off around 0.2 relatively quickly. This is approximately 0.07 less than when the noise parameter is not limited in Case A.

The number of Voronoi cells was limited to 150 (see Figure 5.15 (a)). Similar to Figure 5.12 the algorithm always reached the upper limit of the number of cells. This shows that the algorithm is attempting to use all of the degrees of freedom to fit the data and that the method is not converging.

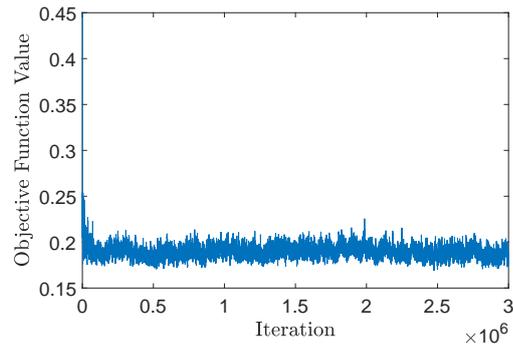


Figure 5.14: The average objective function value as the algorithm iterates through each sample. For the simulation using 32 rays Case B

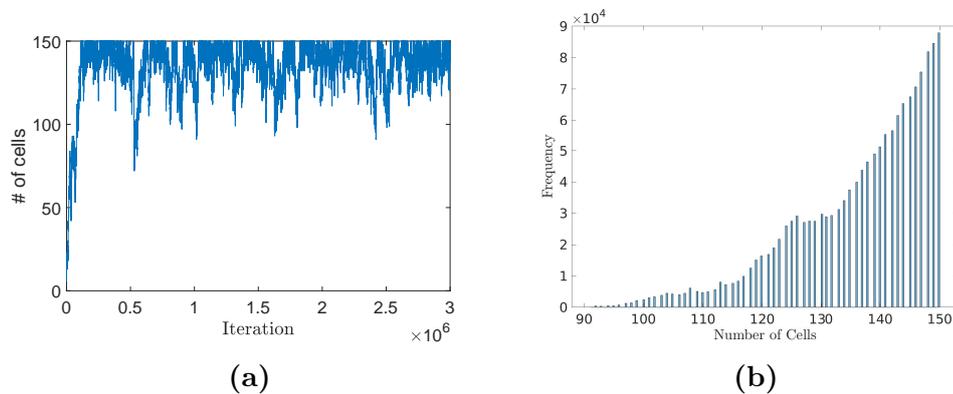


Figure 5.15: Plot (a) is the number of cells at each iteration of the rj-MCMC algorithm and plot(b) is the distribution of the number of cells across all iterations. The maximum number of cells being limited to 150. For the simulation using 32 rays Case B

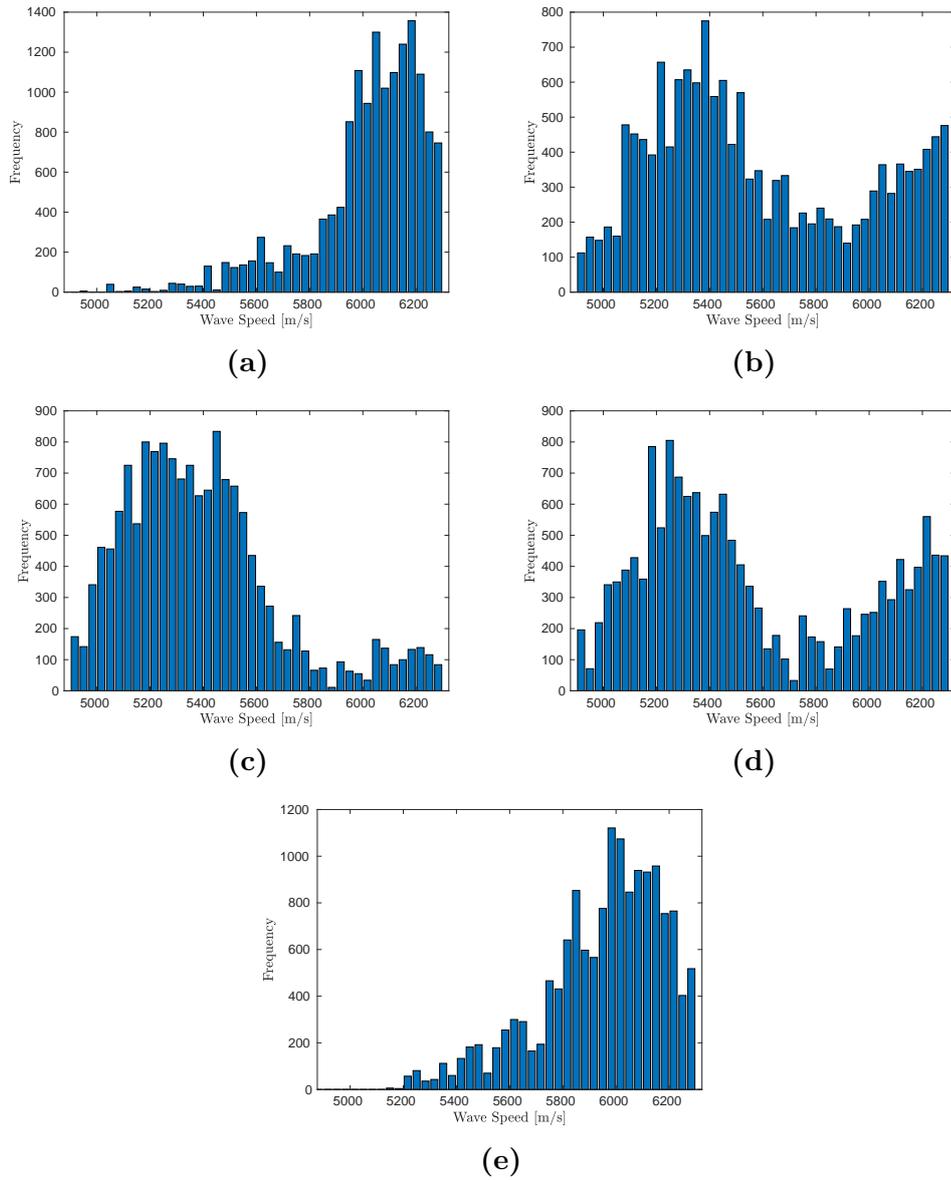


Figure 5.16: Posterior distributions on the wave speeds, recorded at every 100th iteration, at 5 points in the spatial domain: (a) at $x = 16\text{mm}$, $y = 15\text{mm}$, in the first layer, (b) at $x = 32\text{mm}$, $y = 20\text{mm}$, on the interface between first and second layer, (c) at $x = 22\text{mm}$, $y = 23\text{mm}$, in the second layer, (d) at $x = 23\text{mm}$, $y = 27\text{mm}$, on the interface between the second and third layer, (e) at $x = 20\text{mm}$, $y = 35\text{mm}$, in the third layer, as shown in Figure 5.17

However, the histograms shown in Figure 5.16 depict the posterior distributions of wave speeds at 5 different points in the spatial domain (see points (a)-(e) in Figure 5.17(d)). Plot (a) shows the distribution at point $x = 16\text{mm}$, $y = 15\text{mm}$, a point in the middle of the first layer of the posterior maps; plot (b) the distribution at the point $x = 32\text{mm}$, $y = 20\text{mm}$, located on the interface between the first and second layer; plot (c) the distribution at the point $x = 22\text{mm}$, $y = 23\text{mm}$ located in the second layer of the posterior distribution maps; plot (d) the distribution at the point $x = 23\text{mm}$, $y = 27\text{mm}$ located on the interface between the second and third layer; and plot (e) the distribution of wave speeds at the point $x = 20\text{mm}$, $y = 35\text{mm}$ in the third layer. The points (a),(c), and (e), show truncated normal distributions indicating that the algorithm is confident with the wave speeds in these regions. With the distributions being normal at these points the mean and the median will both give robust reconstructions. It should be noted that in (c) there is a heavy tail and this is most likely due to a higher variation in the exploration of the model space at this point and leads to the higher standard deviation within this layer as shown by Figure 5.17(d) and Figure 5.20. Figure 5.16(b) and (d), show bi-modal distributions and this will be due to the algorithm being unable to decide which wave speed to assign to this pixel. This is indicated by the high standard deviation at these points in Figure 5.17 (d).

From the moments of the posterior distributions of the wave speed at each point in the domain, material maps were constructed. Figure 5.17, shows (a) the mean, (b) the median, (c) the MAP (maximum a posteriori), and lastly (d) the

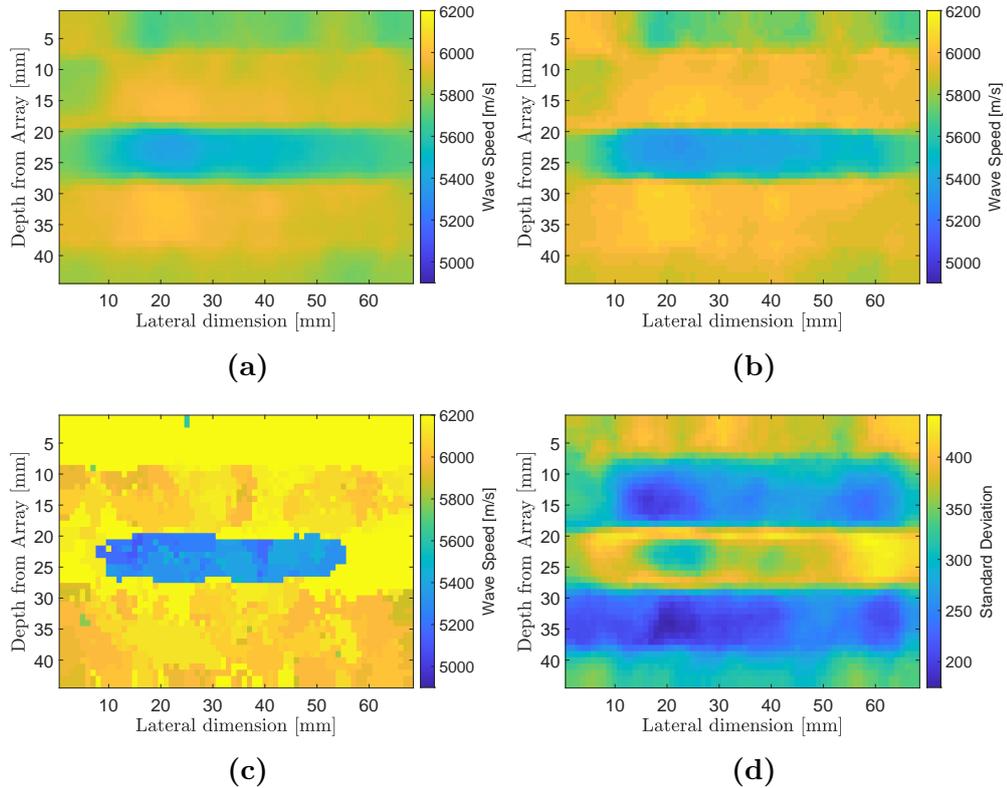


Figure 5.17: Material map reconstructions of the 4 layered medium for the simulation using 32 rays Case B. Plot (a) is the mean of the posterior distribution on wave speed at each point in the domain whilst plot(b)-(d) show the median, MAP, and the standard deviation respectively.

standard deviation map of the posterior probability distribution. Four distinct layers can be observed with a similar fluctuation in wave speeds between the layers and so the methodology has recovered the texture of the material; however, the layer thicknesses and wave speeds have significant errors and this is highlighted in Figure 5.18 which depicts a vertical cross section through each of the maps. The first layer has a thickness of approximately 5mm and has some variations in the wave speed in the horizontal (x) direction within this layer and these

fluctuations differ between the different reconstructed maps. The top left hand corner of the median and the mean reconstructed maps has an area of higher wave speed with the region below having a lower wave speed than the rest of the layer. Features like this are not surprising since only the vertical rays are taken into account. There could be only a single ray going through the Voronoi cell depicting this area, whereas, an equivalent area in the centre could have multiple rays passing through the same Voronoi cell. The MAP has a wave speed of approximately 6200m/s in the top layer, while the mean and median estimate 5750m/s (closer to the known wave speed). There is also a high level of standard deviation in this region. The reconstruction of the second layer has a thickness of approximately 12mm and exhibits wave speeds ranging between 5850m/s and 6150m/s. The third layer, with thickness approximately 8mm, is faster than the known wave speed with a range of 5400m/s to 5600 m/s. The reconstruction of the last layer has a thickness of approximately 14mm, an estimated wave speed range of 5900m/s to 6000m/s for the median and mean, and 6200m/s to 6300m/s for the MAP. Figure 5.17 (d) plots the standard deviation at each pixel in the reconstruction shown in (b). The uncertainty present is strongest at the interfaces between the layers. however, this is expected as the wave speed value of these pixels move between the two wave speed values in the adjacent regions (see the bi-modal distribution in Figure 5.16(b) and (d)) and this can in fact be used to tell us about the achieved resolution of the reconstructions.

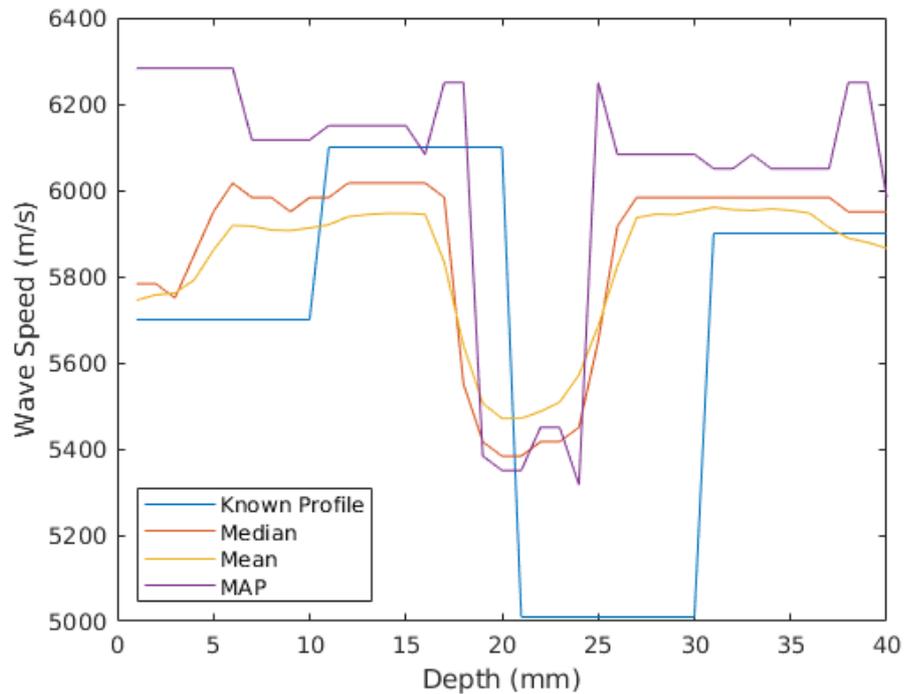


Figure 5.18: Wave speed profile of moments of the posterior distribution of wave speeds (Figure 5.17) along the vertical cross-section at $x = 32\text{mm}$

The difference in wave speed between the reconstructions from the known map is shown in Figure 5.18 which plots a vertical cross-section of the medium at the halfway point (32mm) of the wave speed profile. The variation of the wave speed in the reconstructed material maps in the lateral direction (x -coordinate) along the length of the material at each millimetre depth is shown in Figure 5.19 with the wave speed profile at $x = 32\text{mm}$ where plot (a) shows the mean, plot (b) shows the median, and plot (c) shows the MAP. The blue boxes denote the interquartile range and the red crosses denote the outliers. This shows that the wave speeds are relatively consistent along the lateral direction of the medium, apart from in the MAP profile between 18mm and 25mm where there is a large

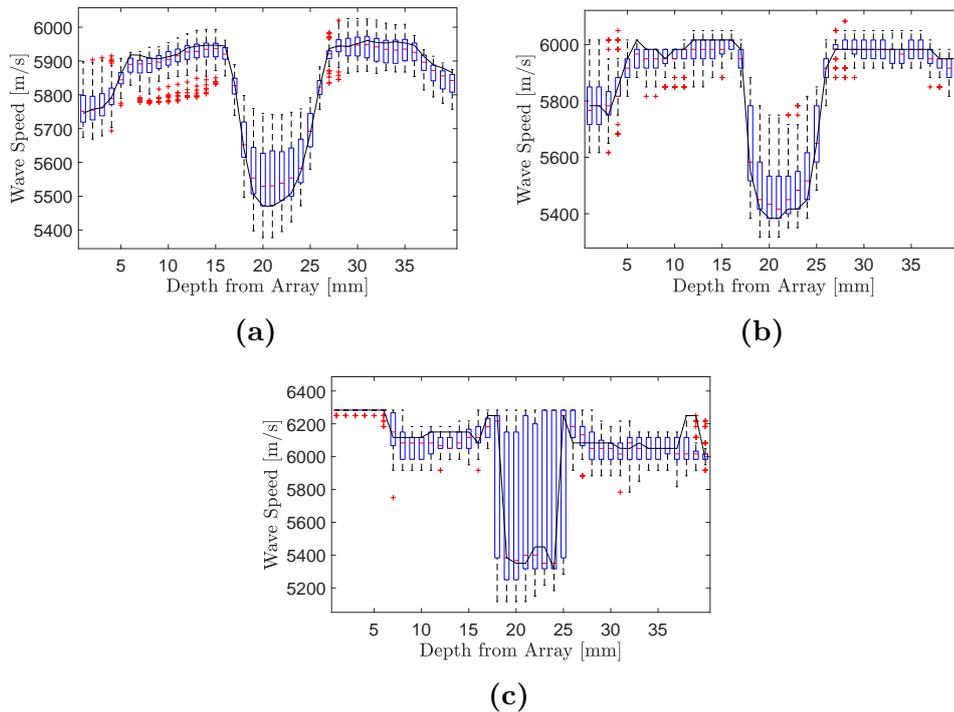


Figure 5.19: Wave speed profiles at $x = 32\text{mm}$ and box plots of the wave speed variation at every mm depth of the medium. Plot (a) is generated using the mean map, plot (b) the median, and plot (c) the MAP from Figure 5.17.

variation in the wave speeds in layer 3 due to the high wave speeds at the sides of the reconstructions. The differences in location of the layer interfaces, between the reconstructions from the known map is shown in Figure 5.20, which plots the standard deviation along the depth of the medium at $x = 32\text{mm}$; a high standard deviation should indicate a layer interface. From this it can be observed that there is an overall high standard deviation, however, there are still peaks at 17mm and 25mm which indicated layer interfaces.

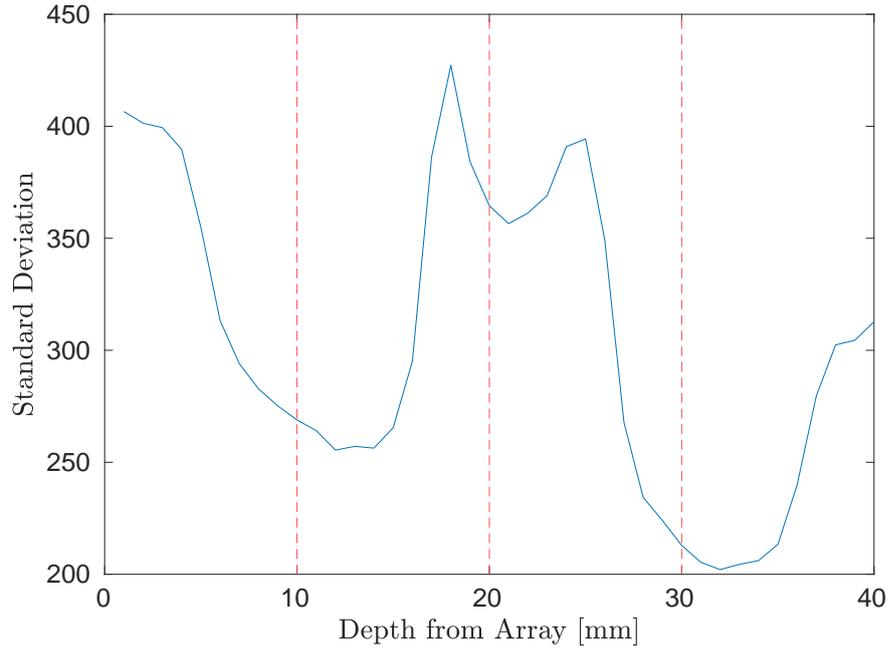


Figure 5.20: A vertical profile of the standard deviation (blue) (Figure 5.17(d)) at $x = 32\text{mm}$ and the known layer interfaces (red dashed).

The results from Figure 5.18, Figure 5.19, and Figure 5.20, shows visually that layer texture can be reconstructed but there is significant room for improvement. Before attempting to improve on the material reconstruction the quality of the reconstruction is quantified and then tested for its improvement for flaw detection as part of the TFM+ algorithm.

To quantify the accuracy of the reconstructed maps in Figure 5.17, the Frobenius norm

$$\varepsilon(M^*, M) = \|M^* - M\|_F \quad (5.1)$$

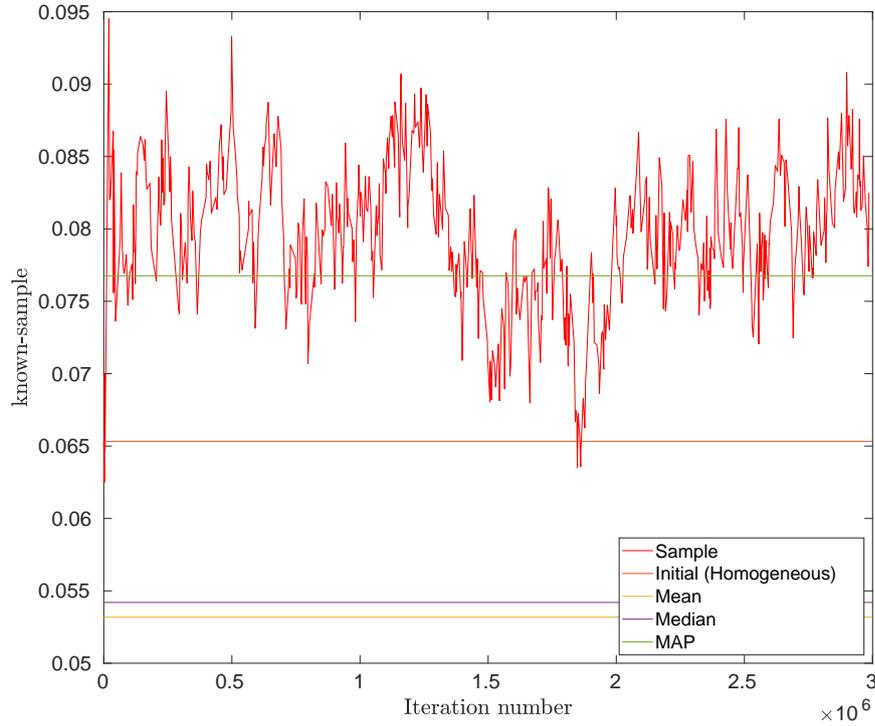


Figure 5.21: Plot of the Frobenius norm (ε) between the known map (M^*) and each 1000th Voronoi diagram (red). Also, the Frobenius norm versus: Initial - homogeneous - case (M^H , $\varepsilon^H = 0.65$, blue), Mean posterior (\overline{M} , $\overline{\varepsilon} = 0.053$, orange), Median posterior (M^m , $\varepsilon^m = 0.054$, yellow), MAP (M^{map} , $\varepsilon^{map} = 0.077$, purple), as shown in Figure 5.17.

of the difference matrix between the known map (M^*) and the initial (homogeneous) (M^H), mean (\overline{M}), median (M^m), and MAP (M^{map}), posterior distributions were calculated and plotted in Figure 5.21. It can be observed that the mean of the posterior distribution is closest to the known map, $\overline{\varepsilon} = 0.053$, with the median error $\varepsilon^m = 0.054$ being just slightly higher. The maximum a posteriori fares the worst, $\varepsilon^{map} = 0.077$, most likely due to the incorrect high wave speed region reconstructed at the top of the domain (see Figure 5.17(c)). The

figure also plots the Frobenius norm (denoted ε^k) of the difference between every 1000th accepted Voronoi tessellation (M^k) and the known map . Interestingly, it can be observed that the Frobenius norm error associated with the homogeneous map $\varepsilon^H = 0.065$ is close to the minimum value; this however highlights the fact that the Frobenius norm is an averaged measure of the similarity between two matrices and does not account for the local textures which are better captured by the reconstructions.

An alternative approach which does capture the degree to which the material texture has been reconstructed is the norm,

$$\hat{\varepsilon}(\hat{M}^*, \hat{M}) = \left(\sum_{i=0}^{N-1} \sum_{j=0}^{(N/2)+1} (\hat{M}^*(i, j) - \hat{M}(i, j))^2 \right) - (\hat{M}^*(0, 0) - \hat{M}(0, 0))^2 \quad (5.2)$$

where N is the dimension of the Fourier transform, \hat{M}^* is the two dimensional spatial Fourier transforms of M^* and \hat{M} the corresponding transform for M^H , \bar{M}, M^m, M^{map} . This is normalised with respect to the value $\hat{\varepsilon}(\hat{M}^*, 0)$. From Figure 5.22, it can be observed that the median of the posterior distribution is closest to the known map $\hat{\varepsilon}^m = 0.0082$, with the mean not as close, $\hat{\varepsilon} = 0.0098$ and the MAP fares the worst $\hat{\varepsilon}^{map} = 0.0122$, most likely due to the incorrect high wave speed region reconstructed at the top of the domain. The homogeneous map with a value of $\hat{\varepsilon}^H = 0.017$ indicates that texture is encapsulated in this metric.

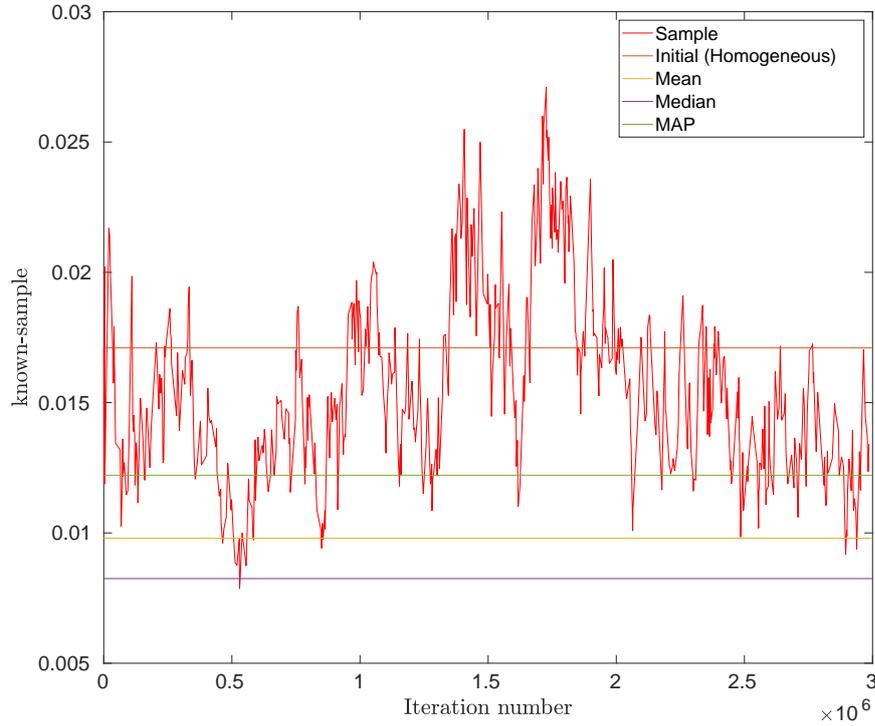


Figure 5.22: Plot of the error ($\hat{\varepsilon}$) of the two dimensional Fourier transform of the known map and each accepted 1000th sample ($\hat{\varepsilon}^k$, red). Similarly for the; Initial - homogeneous - case ($\hat{M}^H, \hat{\varepsilon}^H = 0.017$, orange), Mean posterior ($\hat{M}, \hat{\varepsilon} = 0.0098$, yellow), Median posterior ($\hat{M}^m, \hat{\varepsilon}^m = 0.0082$, purple), and MAP ($\hat{M}^{map}, \hat{\varepsilon}^{map} = 0.0122$, green), as shown in Figure 5.17.

These results agree with what is observed visually. The median of the posterior distribution gives the best reconstruction with the layers being more defined. The MAP being reconstruction is adversely affected the bimodal nature of the posterior distribution and the method can select the “wrong” modal. This could, in some cases, be an aid as it has the potential to give greater definition at layer interfaces. In all cases there is still room for improvement in the reconstructions with regards to wave speed and layer thicknesses.

5.4.4 TFM+ from the Reconstructed Material Map Using 32 Rays

The success of the material reconstructions can also be measured by the improvement in flaw detection and characterisation using the TFM+ algorithm. The flaw reconstruction, using the median posterior distribution map from Figure 5.17 (b), is displayed in Figure 5.23. It can be observed that the TFM+ reconstruction is unsatisfactory and visually appears worse than the standard TFM; there is only one satisfactorily detectable inclusion and there is a shift in the back wall. The shift in position of the backwall is due to the thicker layers within the reconstruction with a high velocity and the layer with a low wave speed being thinner than the known map thus creating a higher mean velocity. The SNR is 16.1, the flaw location error is 1.924mm (for the lowest flaw), and the ROC AUC is 0.839. The SNR and flaw location error both perform worse than the standard TFM. This indicates that a poor reconstruction can actually perform worse than the homogeneous map within the imaging algorithm. There is only a marginal improvement in the ROC AUC measure (see Table 5.5). The ROC curves are shown in Figure 5.24 which shows comparable results between the standard TFM and the TFM+ generated from the median posterior distribution material map. These values are summarised in Table 5.5

There is clearly scope to improve the reconstructed material maps and obtain a successful TFM+ reconstruction. One route to achieving this is to use more information and employ the A-scans from all transmit-receive pairs; not just the

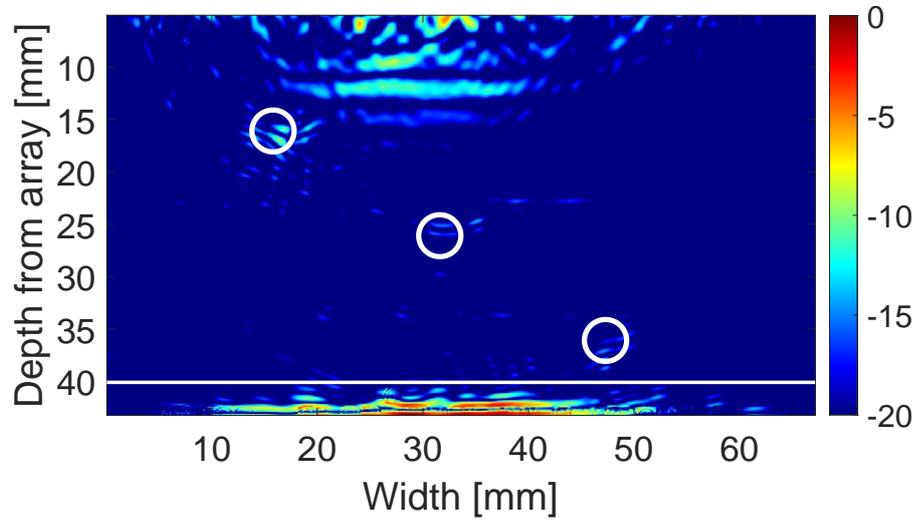


Figure 5.23: A TFM+ image of the three circular inclusions shown in the schematic in Figure 5.1 and indicated by the three white circles. The horizontal is the position of the back wall. The TFM+ algorithm Equation (2.22) with the material map given by the median of the posterior distribution as shown in Figure 5.17.

Table 5.5: Comparison of quality metric between SNR, Flaw Error and ROC AUC for the simulation using 32 rays Case B.

	TFM	TFM+FWI	TFM+
SNR	17.5	16.1	23.6
Flaw Error [mm]	1.315	1.924	1.300
AUC	0.8304	0.8389	0.9012

vertical rays as used above. Using the FMC data means an increase in calculation time but to balance this the inversion algorithm runs for fewer iterations. This is shown in Section 5.5.

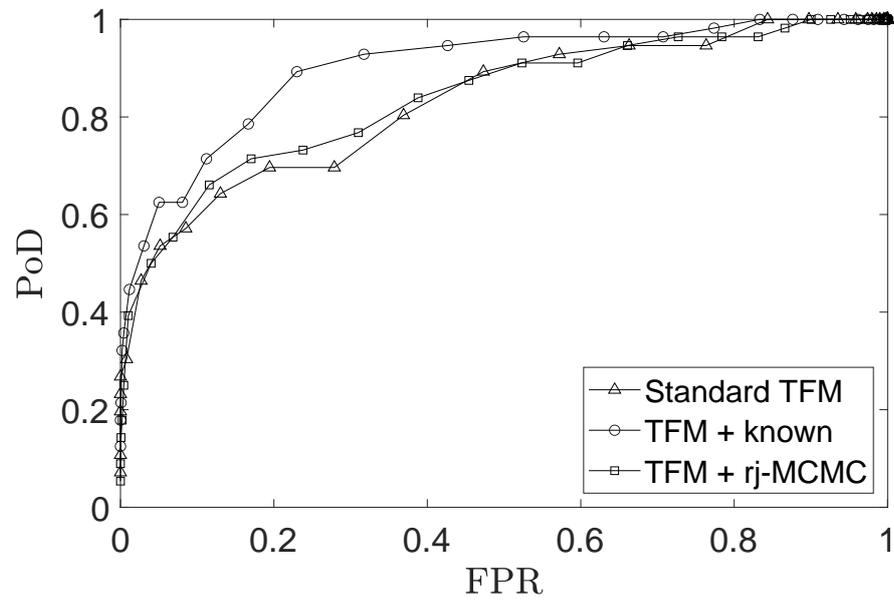


Figure 5.24: ROC curves, from Figure 5.6, including the TFM+ generated from the median posterior distribution from Figure 5.17(b).

5.5 Reconstruction of a Layered Medium using 1024 rays from the Full Matrix Capture

In this section, the case where all 1024 A-scans are utilised for the material map reconstruction. The rj-MCMC was run for a total of 600,000 samples, where the first 400,000 were discarded and the remaining models were sampled at an interval of $\kappa = 100$ (note that the problem is constrained by more data, therefore can reduce the number of iterations the MCMC is run over). The initial model used was composed of eight Voronoi cells (in a rectangular grid pattern as before), each assigned a wave speed of 5600 m/s.

The algorithm had the following limits on the prior uniform distributions: the number of cells must lie in the range 5 to 75 cells; the wave speed for a Voronoi cell must lie in the range 5600 ± 700 m/s; and the noise parameter must lie in the range 0.001 and 0.5. The standard deviation for the proposal distributions (σ) were: on *wave speed* 200m/s, and 50m/s for the delayed rejection (σ^{dr}); on *cell position* 8%, and 2% for the delayed rejection; on the wave speed for a new cell generated by a *birth* 100m/s; on the *noise* parameter 0.03. These are summarised in Table 5.6. The standard deviation for the proposal distributions were chosen such that the overall acceptance rate would lie between 23% and 40%. The limits on the uniform prior on the noise parameter were chosen by setting the lower limit to be lower than the objective function value between the FMC data modelled through the correct map and the observed data of 0.0029 (minor difference due to numerical discrepancies). The upper limit being greater than the objective

function value between the FMC data arising from a homogeneous map and the observed data. With these priors the algorithm had an overall acceptance rate of 29.9%. The acceptance rates were 57.7% for the wave speed, 42.9% for the cell seed position, 9.7% for the birth and death of a cell, and 20.3% for the noise. By using the full 1024 rays possible, the Voronoi cells within the model have more ray paths going through them. This then affects the model parameters even though it is the same underlying model.

Table 5.6: Prior and proposal distributions set for the simulation for 1024 rays.

Perturbation	Range	σ	σ^{dr}
Wave speed	5600 ± 700 m/s	200 m/s	50m/s
Seed position	–	8%	2%
Birth	–	100 m/s	–
Noise	0.001 - 0.5	0.03	–

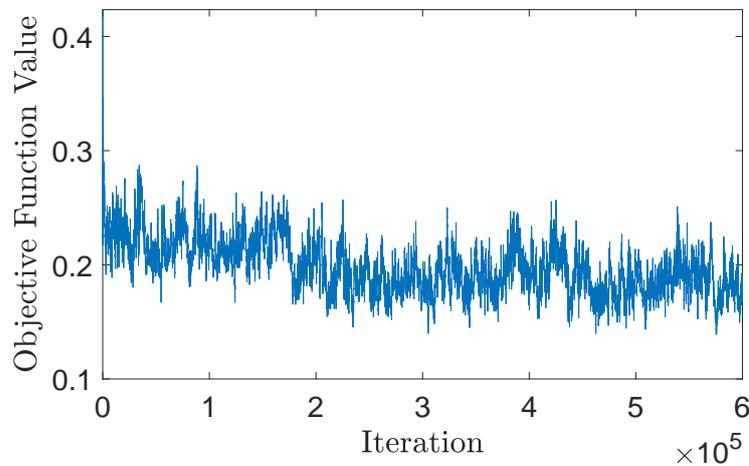


Figure 5.25: The average objective function value as the algorithm iterates through each Voronoi sample for the simulation using the full 1024 rays.

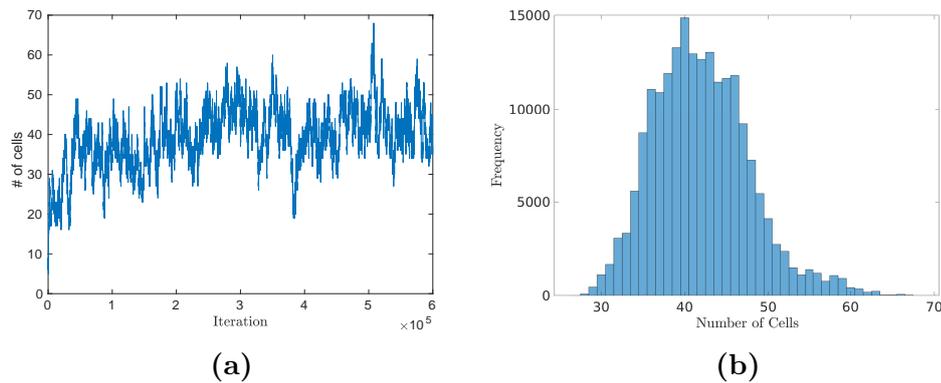


Figure 5.26: Plot (a) is the number of cells at each iteration of the rj-MCMC algorithm and plot(b) is the distribution of the number of cells from the last 200,000 samples. For the simulation using the full 1024 rays.

The average objective function value, Equation (2.1), as the algorithm iterates through the samples, is shown in Figure 5.25. The initial objective function value is 0.424 and it levels off below 0.2 after 200,000 iterations. The number of cells steadily increases for the first 300,000 iterations then stays around 40 cells for the remainder of iterations. This is highlighted in Figure 5.26 which shows the number of Voronoi cells for the last 200,000 iterations, omitting the burn-in, has a skewed normal distribution with a tail to the right and has a peak centred around 40 cells. The number of cells is higher than one may expect for a simple model but it has converged in this case. This can in part be attributed to the fact that, although it is possible to describe this layered material with four Voronoi cells, any deviation in the x co-ordinates will cause the layer interfaces to diverge from the horizontal. However, less precision is required in describing the horizontal layers using multiple cells which balance out these skewing effects.

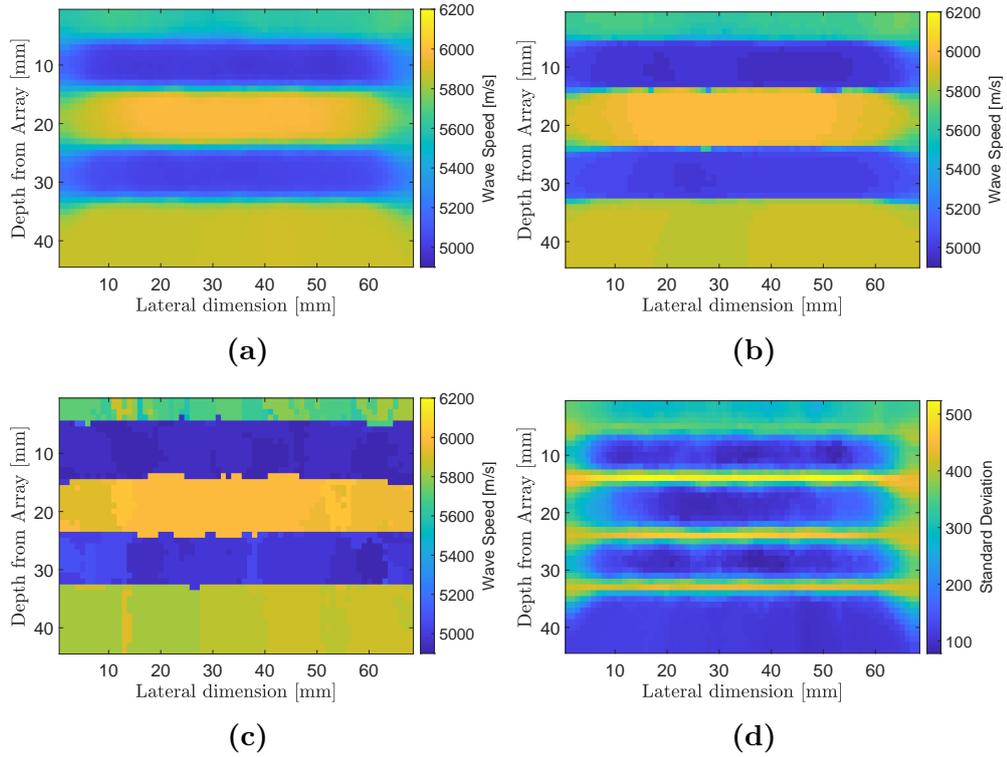


Figure 5.27: Material map reconstructions of the 4 layered medium for the simulation using the full 1024 rays. Plot (a) is the mean of the posterior distribution on wave speed at each point in the domain whilst plots (b)-(d) show the median, MAP, and the standard deviation respectively.

From the moments of the posterior distributions on the wave speed at each point in the domain, material maps were constructed. Figure 5.27, shows (a) the mean, (b) the median, (c) the MAP, and (d) the standard deviation of the posterior probability distributions. Five distinct layers can be observed; three layers well approximate the correct layer thickness and wave speed, however, the first layer is reconstructed as two layers. The reason for this could be in part due to the fact that the semi-analytical data does not contain any internal reflections from the first layer. It is modelled using a backwards recursive relation from

the bottom layer to the top. The wave speeds of the reconstructed layers are approximately 5650m/s, 5000m/s, 6000m/s, 5000m/s, and 5850m/s respectively. Figure 5.27 (d) plots the standard deviation of the posterior distribution on wave speed at each pixel. It shows a high standard deviation at the layer interfaces and lower within the layers. This suggests the problem is poorly constrained at layer interfaces; the pixels here jump between neighbouring layers and typically exhibit bimodal posteriors. Note that there is a high level of standard deviation in the first layer again suggesting this part of the domain is poorly constrained.

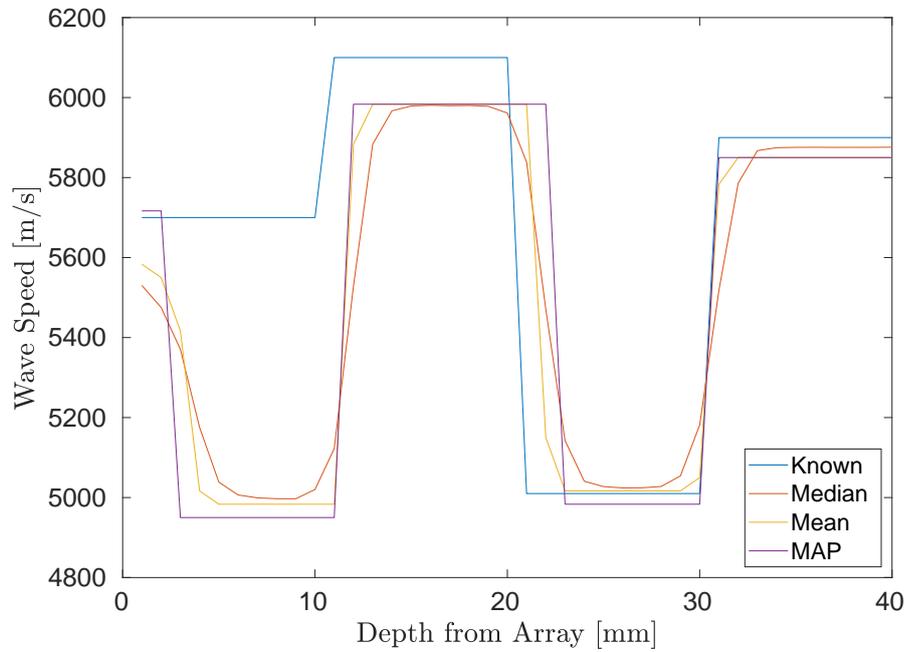


Figure 5.28: Wave speed profile of moments of the posterior distributions (Figure 5.27) along the vertical cross-section at $x = 32\text{mm}$.

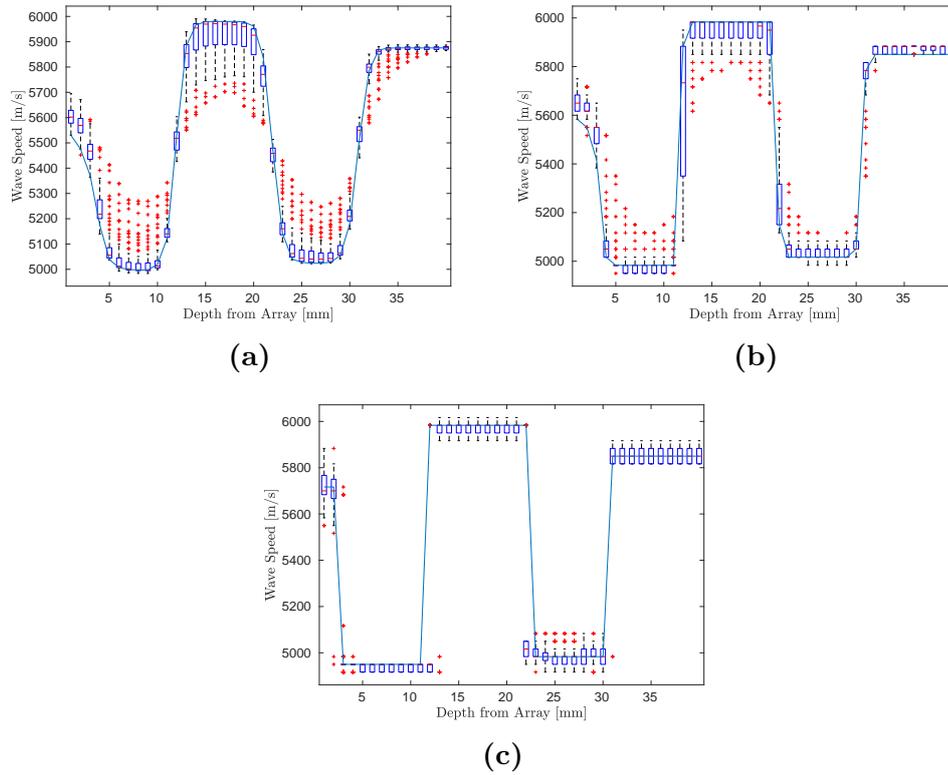


Figure 5.29: Wave speed profiles at $x = 32\text{mm}$ and box plots of the wave speed variation at every mm depth of the medium. Plot (a) is the median, plot(b) is the mean, and plot(c) is the MAP, of the posterior distributions from the rj-MCMC method (see Figure 5.27)

The difference in wave speeds between the reconstructions from the known map is shown in Figure 5.28, which plots a vertical cross-section of the medium at the halfway point (32mm) of the wave speed profile and it shows a consistent wave speed across each layer. The layers are more distinct due to the sharp change in wave speeds. The three profiles (mean, median, and MAP) all depict similar profiles meaning there is a high level of robustness in the posterior distribution.

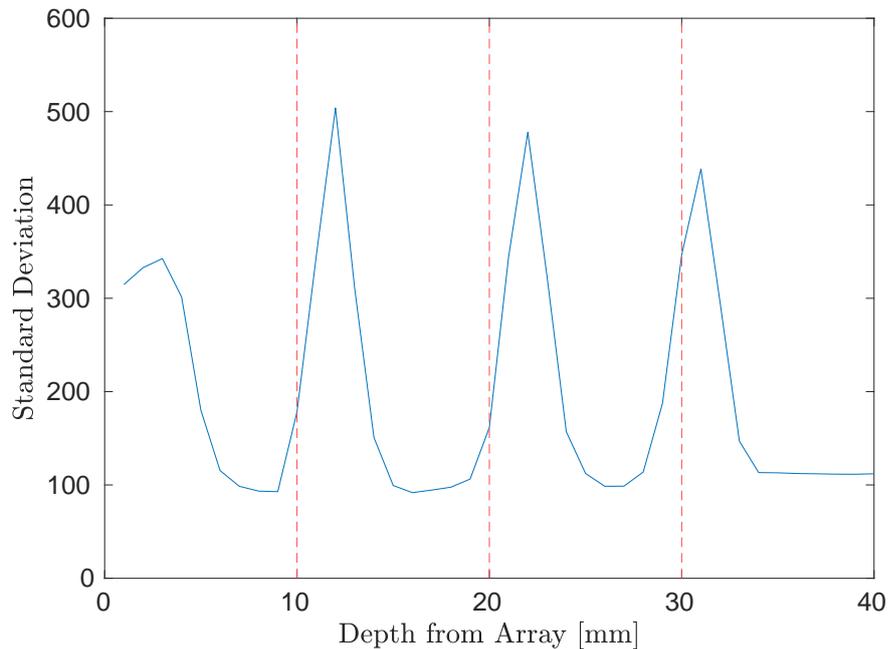


Figure 5.30: A vertical profile of the standard deviation of the posterior distribution (blue) - Figure 5.27(d) - at $x = 32\text{mm}$. The known layer interfaces are shown (red dashed).

This is in contrast to the wave profiles in Figure 5.18 (which were generated using the 32 vertical rays), indicating that with more data the inversion is better constrained and gives better tomographic reconstructions.

The variation of the wave speed along the length of the medium at each millimetre depth, with the wave speed profile at $x = 32\text{mm}$, is shown in Figure 5.29, where plot (a) shows the mean, plot (b) shows the median, and plot (c) shows the MAP, of the posterior distributions. This shows that the wave speeds are consistent along the length of the medium by having a small interquartile range (blue boxes), apart from at the layer interfaces where wave speeds from the ad-

jacent layers may occur. There is a greater amount of outliers (red crosses) in (a) the profile from the mean map, due to the wave speed being closer to the homogeneous value of 5600 m/s at the edges of the reconstruction in Figure 5.27 (a). The reason for this being more present in the mean map rather than the other two indicates that the posterior distributions near the edges are skewed. With a skewed distribution the mean does not accurately represent the centre of a distribution whereas it does in the median and MAP. Note, the edges of the image are poorly constrained within the reconstruction; there is a maximum of 4 rays going through this region, whereas, a region in the centre of the reconstruction could have 16 rays. The interquartile ranges present in these profiles are much smaller and consistent compared to the interquartile ranges present in Figure 5.19. This highlights the need to have multiple rays at different angles to fully constrain the inversion and give a more accurate reconstruction. Figure 5.30 plots the standard deviation along the depth of the medium at $x = 32\text{mm}$. It can be observed that the internal layer interfaces are consistently 10mm apart and are about 1mm off the correct location.

To quantify the accuracy of the reconstructed maps in Figure 5.27, the Frobenius norm error was examined between the two dimensional, spatial Fourier transforms of \hat{M}^* and \hat{M} as given by Equation (5.2). From Figure 5.31, it can be observed that the mean and the median of the posterior distribution map fare the best (the mean has an error of $\hat{\varepsilon} = 0.0049$ and the median has error of $\hat{\varepsilon}^m = 0.0057$) and the MAP fares the worst (with $\hat{\varepsilon}^{map} = 0.0071$).

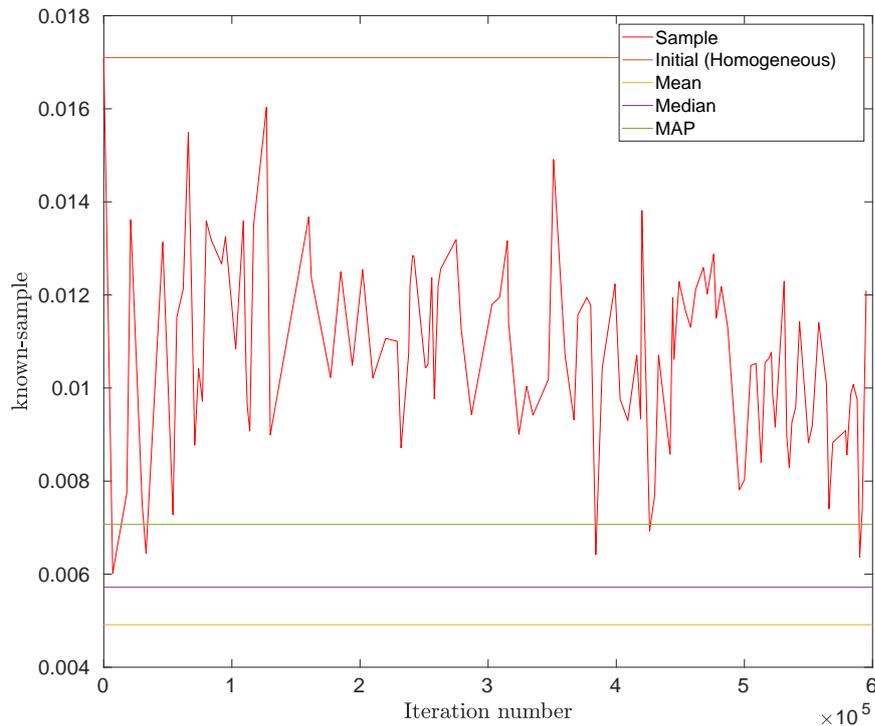


Figure 5.31: Plot of the error ($\hat{\varepsilon}$) between the two dimensional Fourier transforms of the known map and each accepted 1000th sample ($\hat{\varepsilon}^k$, red). Also, the error for the initial - homogeneous - case (\hat{M}^H , $\hat{\varepsilon}^H = 0.017$, orange), Mean posterior (\hat{M} , $\hat{\varepsilon} = 0.0049$, yellow), Median posterior (\hat{M}^m , $\hat{\varepsilon}^m = 0.0057$, median), and MAP (\hat{M}^{map} , $\hat{\varepsilon}^{map} = 0.0071$, green), from Figure 5.27.

By visually inspecting the moments and characteristics of the posterior distributions (see Figure 5.27) and examining the difference in the spatial Fourier transform, it can be concluded that the texture of the material has been successfully characterised; with confidence that the medium is constructed from layers with almost constant thickness. Also the pattern of changing wave speed, from one wave speed to a higher, to a lower, to a higher wave speed, in the longitudinal

direction is also better described. In the lateral direction, the wave speed stays fairly consistent within each layer (see Figure 5.29) where each boxplot depicts the variation of wave speed in the lateral direction.

5.5.1 TFM+ from the Reconstructed Material Map Using 1024 Rays

To examine the success of the reconstructions with respect to enhanced flaw imaging, using the median of the posterior distribution to correct the delay laws in the TFM+ algorithm the resulting flaw reconstructions are displayed in Figure 5.32. It can be observed that there is a shift in the backwall due to the average wave speed of the median posterior distribution (5498m/s) being lower than the known map which equates to a longer travel time through the median map. However, the flaw closest to the backwall is now clearly detected which indicates an improvement from the standard TFM. The inclusion closest to the back wall has a SNR value of 28.4dB and a flaw error of 1.315mm. The ROC curves shown in Figure 5.33 shows that the TFM+ image from the median map has a greater PoD than the standard TFM with an AUC of 0.8713. There is not an improvement in the flaw location error compared to the standard TFM. However, there is improvement in the SNR and a small improvement in the ROC AUC (note, there is not room for significant improvement as the standard TFM performs well in this simple medium). These are summarised in Table 5.7.

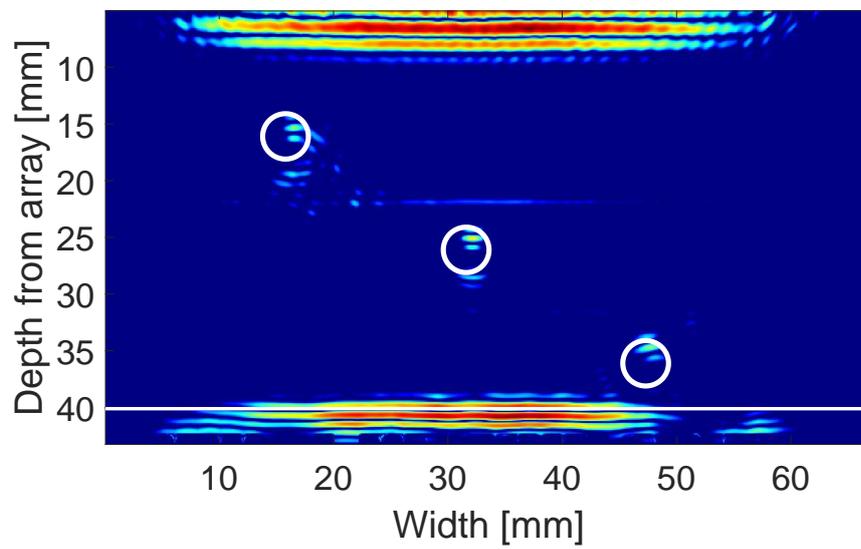


Figure 5.32: TFM+ image (Equation (2.22)) using the Median posterior distribution map, Figure 5.27 (b). The three white circles indicate the inclusions shown in the schematic Figure 5.1 and the white line indicating the position of the backwall.

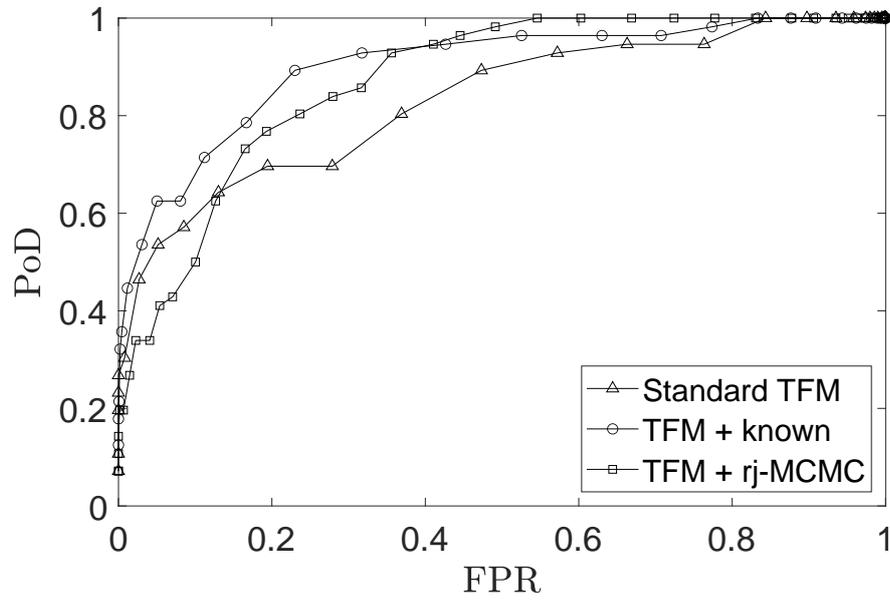


Figure 5.33: ROC curves from Figure 5.6 including the ROC curve from the TFM+ generated from the median posterior distribution from Figure 5.27(b).

Table 5.7: Comparison of quality metric SNR, Flaw Error and AUC between the standard TFM, TFM+ FWI, TFM+ known, for the simulation using the full 1024 rays.

	TFM	TFM+FWI	TFM+
SNR	17.5	28.4	23.6
Flaw Error [mm]	1.315	1.315	1.300
AUC	0.8304	0.8713	0.9012

5.6 Reconstruction of a Layered media with the Addition of Noise

In this chapter so far, the observed data has been created using the same analytical model as the forward model and so this data is effectively noiseless. To check how robust this method is to noise, artificial noise was added to the observed data. This was done in two ways. Proportional noise, where the change in value is dependent on the amplitude at that time point, was added via

$$\mathcal{A}_{i,j}(t) = \mathcal{A}_{i,j}(t) + \mathcal{A}_{i,j}(t)\mathcal{N}(0, e_{N_1}). \quad (5.3)$$

Additive noise, where the change is independent of the data, was added via

$$\mathcal{A}_{i,j}(t) = \mathcal{A}_{i,j}(t) + \mathcal{N}(0, e_{N-2}). \quad (5.4)$$

The FMC dataset from the previous section was modified to have $e_{N_1} = 5\%$, $e_{N_1} = 10\%$, and $e_{N_1} = 20\%$ proportional noise. With the addition of 20% proportional noise the effect on the A-scans can be seen in Figure 5.34. All three simulated data sets were then used as the observed data for the rj-MCMC tomographic inversion algorithm with no other changes to set up or parameters. It can be observed in Figure 5.35 that there are minor variances in the wave speed between the reconstructions, however, 4 layers are still clearly detected with wave speeds still in the same range and the original “noiseless” posterior distribution.

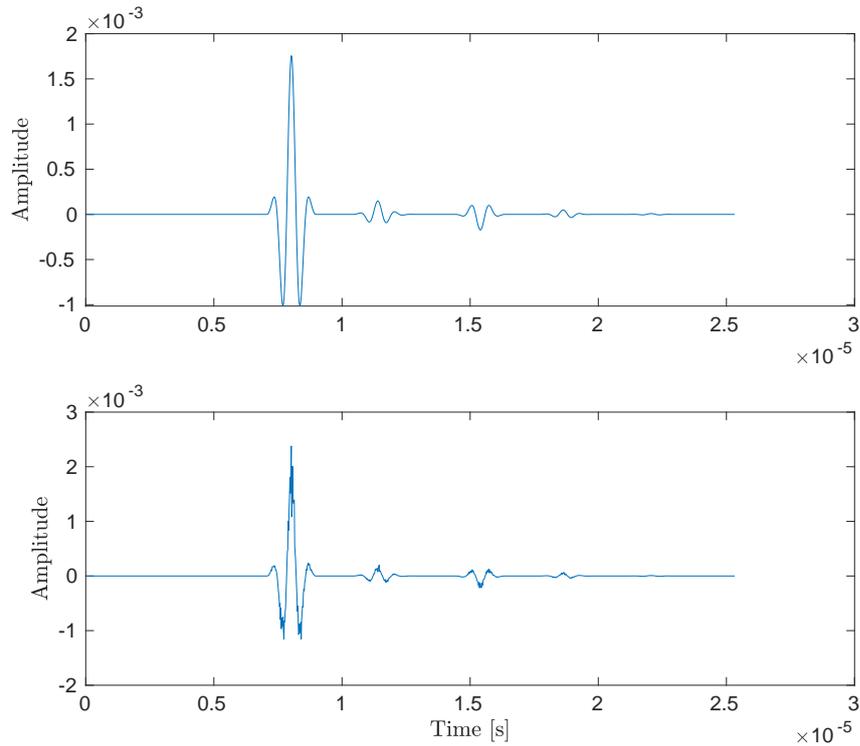


Figure 5.34: A-scan comparison of a wave travelling through 4 layers with no noise (top) and with random proportional noise of 20% variation (bottom).

The changes in these posterior distributions were quantified using the metric in Equation (5.2) to give the $\varepsilon^{\hat{m}} = 0.0057$, $\varepsilon^{m^{\hat{*}0.05}} = 0.0069$, $\varepsilon^{m^{\hat{*}0.1}} = 0.0064$, $\varepsilon^{m^{\hat{*}0.2}} = 0.0078$. This means that proportional noise does effect the reconstruction but not substantially so this method is robust to proportional noise.

Next, examining the case where the data is subjected to additive noise. The amplitude at each point in time was perturbed by a random value from the distribution $\mathcal{N}(0, 0.01e - 3)$. The maximum amplitude of the wave is approximately $1.7e-3$, the equates to an additive noise of around 0.005%. The A-scan with this

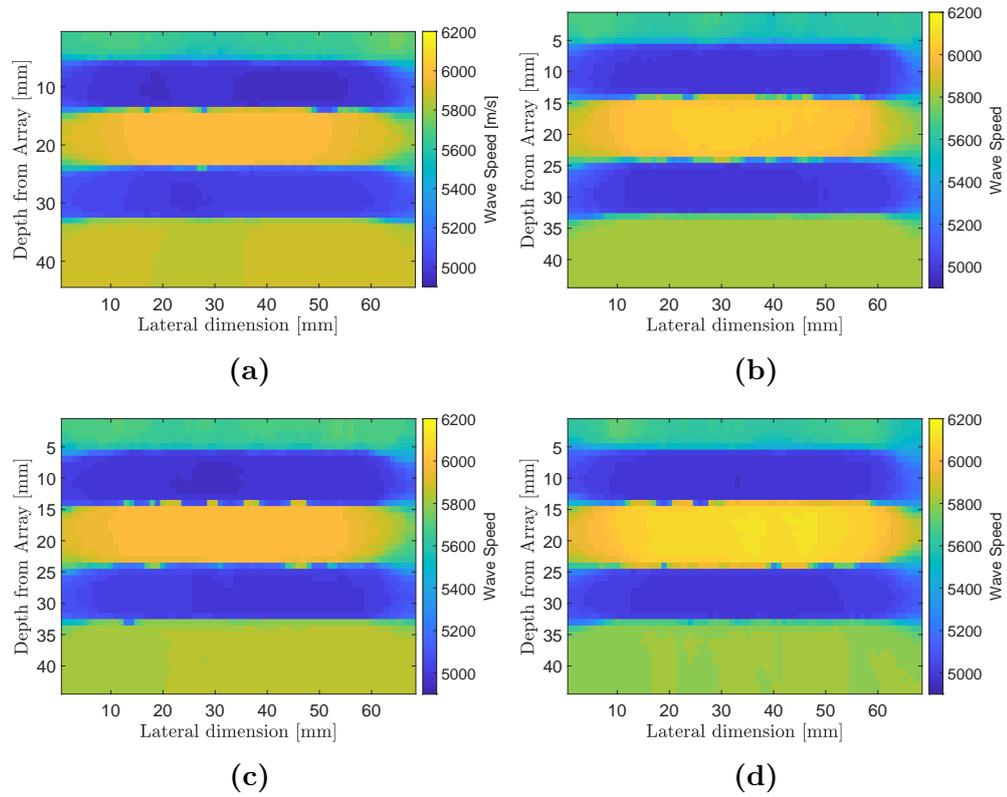


Figure 5.35: Median posterior distributions of the 4 layered medium with varying levels of proportional noise. Plot (a) 0% noise, plot (b) 5% proportional noise, plot (c) 10% proportional noise, and plot(d) 20% proportional noise. The proposal distribution parameters are the same as Table 5.6.

noise can be observe in Figure 5.36. This is similar to the white noise often observed within the experimental system due to electrical interference. With this additive noise the reconstruction was homogeneous with a wave speed of 500 m/s and no features where able to be reconstructed. This is due the objective function taking the log of this data amplifying the effect of the additive noise.

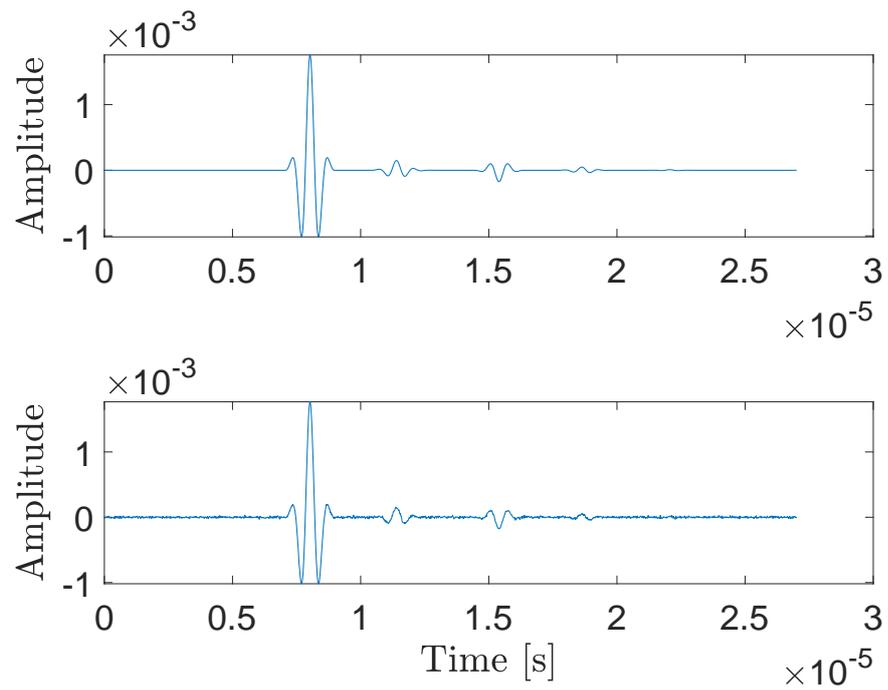


Figure 5.36: A-scan comparison of a wave travelling through 4 layers with no noise and with random additive noise distributed around $0.01e-3$

5.7 Discussion

After being unable to obtain the material map approximation using the Pearson correlation coefficient as the objective function, a new objective function was developed. This new objective function required taking the logarithm of the A-scan. This allowed for the smaller amplitudes to have a greater importance within the objective function. These smaller amplitudes contain the information of the internal reflection within the layered medium.

First, to comply with the one-dimensional assumptions made in the modelling framework, only the vertical rays were considered in the inversion (those travelling perpendicularly to the layer interfaces). In this scenario we were initially unable to reconstruct any material features in either the pitch-catch arrangement (where two arrays were modeled, one on the top boundary to transmit and one on the bottom to record) or the pulse echo arrangement (single array inspection). However, by tightly constraining the noise parameter, we were able to reconstruct some information on the layered texture in the pitch-catch scenario. Moving to consider the full FMC data set (so not imposing restrictions on the incident angles of rays that could be used), the additional data helped to better constrain the problem and reconstructions which better represented both the layered structure and the variation in wave speeds were obtained. These results show that it is possible to do a full waveform inversion as part of an rj-MCMC system. Not only is it possible but the reconstructions from the rj-MCMC algorithm are robust to proportional noise. Using the reconstructed maps as part of the TFM+ imaging

algorithm provided enhanced images of the inclusions. In particular, where there was very little room for improvement, an improvement of 4.1% in probability of detection for the bottom inclusion which was the most difficult to image.

The next step is to invert with simulated data arising from a finite element program which includes more of the physics present in the system (for example mode conversion and refraction) and assess whether our semi-analytic model is enough to reconstruct the material map from this more complex dataset. This is done in the next chapter.

Chapter 6

Tomographic Reconstruction of a Layered Medium from Finite Element Simulated Data

6.1 Adjusting to Simulated Data

Having reconstructed the layered medium (Figure 5.1) where the data has arisen from the forward model used within the inversion framework (inverse crime), it is sensible to challenge this tomographic reconstruction methodology with simulated data that includes more of the physical phenomena observed in practice. For example, the inclusion of different wave mode types (shear/longitudinal),

wave mode conversion at interfaces, and wave refraction, to name but a few. A commercial package (OnScale) was used to run a two dimensional finite element simulation based on numerical integration of the elastodynamic equations [141].

6.1.1 Adjusting to Simulated Data - Planar Wave

The first step is to make sure that the finite element simulation and the A-scans generated from our analytical model have a good agreement. To this end a plane wave travelling through the same two dimensional layered model in Chapter 2 was modelled. A plane wave was emitted from the upper boundary into the layered material whose lateral boundary conditions are periodic. The top half space was assigned the same material properties as the first layer and the bottom half space was modelled as water. A pulse-echo A-scan is shown in Figure 6.1 and five echoes can be observed: the initial wave is the reflection off the front wall; the next three are reflections from the three layer interfaces; the last is the wave reflection off the back wall. Note that there is a small timeshift and a minor difference in amplitude. These are due to the differences in numerical accuracy between the two model implementations and the finite element simulation deforming the wave as it travels through the layers. The first wave pulse represents the recording of the pulse as it is transmitted and is not part of the semi-analytical model.

An A-scan generated in pitch-catch mode is shown in Figure 6.2. It can be observed that there is a significant difference in the amplitudes. On investigation the reason for this is the difference in what quantity each model outputs.

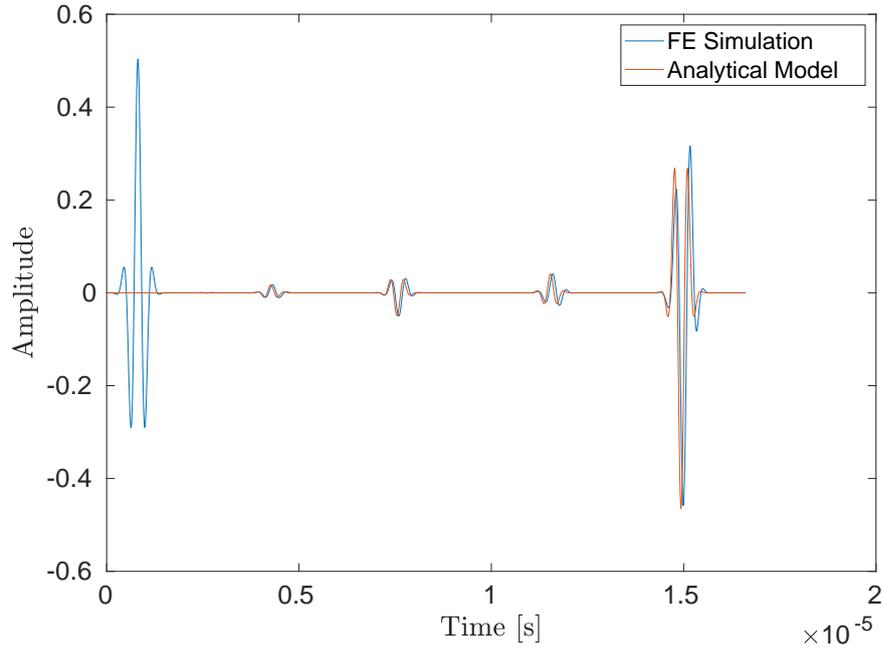


Figure 6.1: An A-scan comparison of a finite element simulated reflected wave (blue line) and the equivalent from the semi-analytic model Equation (3.5) in pulse-echo mode. The wave is travelling through the 4 layered medium shown in Figure 5.1

The model in Chapter 2 has a wave $A(t, z)$ with dimensions $\text{kg}^{\frac{1}{2}}\text{s}^{-\frac{3}{2}}$ and has a transmission coefficient between layer i and layer $i + 1$ given by

$$T_i = \frac{2\sqrt{\zeta_i \zeta_{i+1}}}{\zeta_i + \zeta_{i+1}}.$$

The finite element simulation models a pressure wave $p(t, z)$ (dimensions $\text{kgm}^{-1}\text{s}^{-2}$) which has a transmission coefficient of,

$$\tilde{T}_i = \frac{2\zeta_{i+1}}{\zeta_i + \zeta_{i+1}},$$

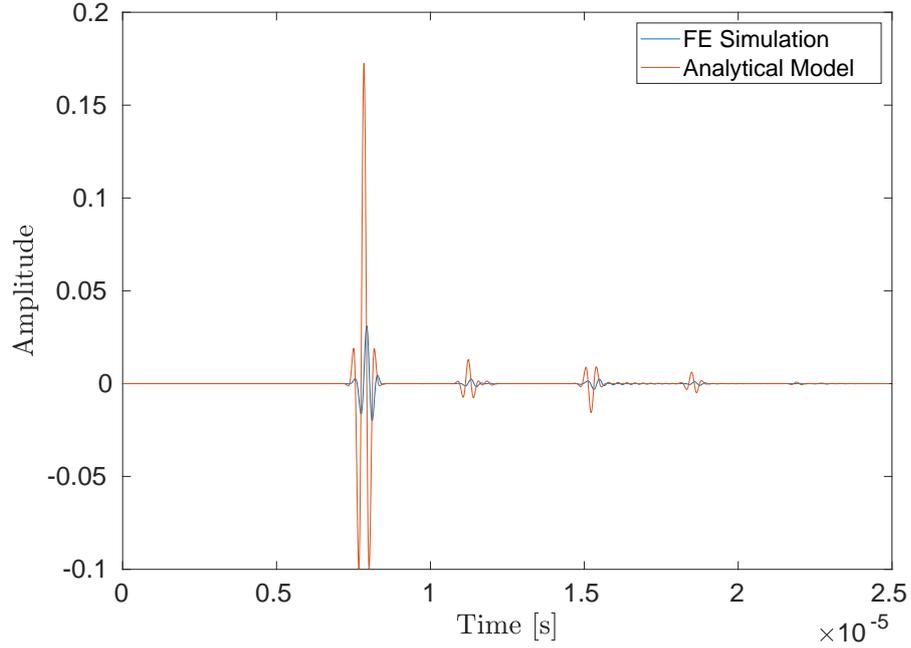


Figure 6.2: An A-scan comparison of a finite element simulated transmitted wave (blue line) and the equivalent from the semi-analytic model (orange line) Equation (3.7) in the Pitch-catch mode. The wave is travelling through the 4 layered medium shown in Figure 5.1

at the same interface. The upper half space, with impedance ζ_0 , contains the downward travelling wave $f(t)$ - the source term. Therefore, in the first layer with impedance ζ_1 , there is a downward travelling wave A_1 which is equal to $T_0 f(t - z/c)$ when using the semi-analytical model. However, when the pressure (and velocity) waves are produced in the finite element simulation, they are $p_0^r =$

$\frac{1}{2}\zeta_0^{\frac{1}{2}}f$ (in the upper half space) and $p_1^r = \frac{1}{2}\zeta_1^{\frac{1}{2}}A_1$ (in the first layer). That is

$$\begin{aligned} p_1^r &= \frac{1}{2}\zeta_1^{\frac{1}{2}}A_1 = \frac{1}{2}\zeta_1^{\frac{1}{2}}T_0f \\ &= \frac{1}{2}\zeta_1^{\frac{1}{2}}\left(\frac{2\sqrt{\zeta_0\zeta_1}}{\zeta_0 + \zeta_1}\right)f \\ &= \frac{1}{2}\zeta_0^{\frac{1}{2}}\left(\frac{2\zeta_1}{\zeta_0 + \zeta_1}\right)f = \left(\frac{2\zeta_1}{\zeta_0 + \zeta_1}\right)p_0^r \\ &= \tilde{T}p_0^r. \end{aligned}$$

In the upper half space there is a downward moving pressure wave p_0^l and an upward moving pressure wave p_0^r . This means there a factor of $\frac{1}{2}\zeta_0^{\frac{1}{2}}$ between the semi-analytical model in Chapter 2 these pressure waves to their finite element counterparts. These are related by the following equations,

$$\begin{aligned} p_0^l &= -\frac{1}{2}\zeta_0^{\frac{1}{2}}B_0, \\ p_0^r &= \frac{1}{2}\zeta_0^{\frac{1}{2}}f. \end{aligned}$$

This can be repeated for the velocity wave $u(t, z)$.

The reflection coefficient, in both formulations are equal in amplitude but differ in sign and $-R = \tilde{R} = \left(\frac{\zeta_1 - \zeta_0}{\zeta_0\zeta_1}\right)$ and $B_0 = Rf$. Hence,

$$\begin{aligned} p_0^l &= -\frac{1}{2}\zeta_0^{\frac{1}{2}}B_0 = -\frac{1}{2}\zeta_0^{\frac{1}{2}}Rf \\ &= \tilde{R}p_0^r, \end{aligned}$$

This explains why there was good amplitude agreement in the reflection (pulse-echo) case in Figure 6.1 but poor agreement in the transmission (pitch-catch) case in Figure 6.2.

The analytical model therefore needs to be adapted such that it uses \tilde{T} and \tilde{R} . For the latter case this is simply done by multiplying by -1 (which has already been applied in Figure 6.1). For the former case, in the pulse-echo scenario every time a wave is transmitted into a layer it then needs to be transmitted back out of that layer for the wave to be received. So the transmission coefficients appear in the amplitude term in product pairs. Note that

$$\begin{aligned} T_i^2 &= 4 \frac{\zeta_i \zeta_{i+1}}{(\zeta_i + \zeta_{i+1})^2} \\ &= \left(\frac{2\zeta_i}{\zeta_i + \zeta_{i+1}} \right) \left(\frac{2\zeta_{i+1}}{\zeta_i + \zeta_{i+1}} \right) \\ &= \tilde{T}_i^r \tilde{T}_i^l \end{aligned}$$

and so the transmission coefficient product pairs are identical in both formulations.

In the pitch-catch case, every amplitude calculation contains the product

$$\begin{aligned}
T_0 T_1 \dots T_N &= \frac{2\sqrt{\zeta_0 \zeta_1}}{\zeta_0 + \zeta_1} \frac{2\sqrt{\zeta_1 \zeta_2}}{\zeta_1 + \zeta_2} \frac{2\sqrt{\zeta_2 \zeta_3}}{\zeta_2 + \zeta_3} \dots \frac{2\sqrt{\zeta_{N-1} \zeta_N}}{\zeta_{N-1} + \zeta_N} \\
&= \frac{2^N \sqrt{\zeta_0 \zeta_1 \zeta_2 \dots \zeta_{N-1} \zeta_N}}{(\zeta_0 + \zeta_1)(\zeta_1 + \zeta_2) \dots (\zeta_{N-1} + \zeta_N)} \\
&= \left(\frac{\zeta_N}{\zeta_0}\right)^{\frac{1}{2}} \frac{2^N \zeta_1 \zeta_2 \dots \zeta_{N-1} \zeta_N}{(\zeta_0 + \zeta_1)(\zeta_1 + \zeta_2) \dots (\zeta_{N-1} + \zeta_N)} \\
&= \left(\frac{\zeta_N}{\zeta_0}\right)^{\frac{1}{2}} \tilde{T}_0 \tilde{T}_1 \dots \tilde{T}_N,
\end{aligned}$$

for a N layer material. Therefore to change the pitch-catch case all that is required is to multiply by the prefactor $\left(\frac{\zeta_N}{\zeta_0}\right)^{\frac{1}{2}}$ so that the analytical model agrees with the finite element model. Note, that when the half-spaces are the same ($\zeta_0 = \zeta_{N+1}$) then this prefactor becomes 1.

When this change has been implemented, the A-scans compare well in the pitch catch case as shown in Figure 6.3. Here, a good match up is observed, with a small timeshift and minor differences in amplitude due to the additional physics in the finite element model.

6.1.2 Adjusting to Simulated Data - Spherical Wave

The previous section considered a plane wave excitation as this essentially transforms the 2D finite element simulation into a 1D problem and a good agreement with the 1D analytical model was found. However, if the Full Matrix Capture (FMC) data is to be used then more typically each array element would fire in

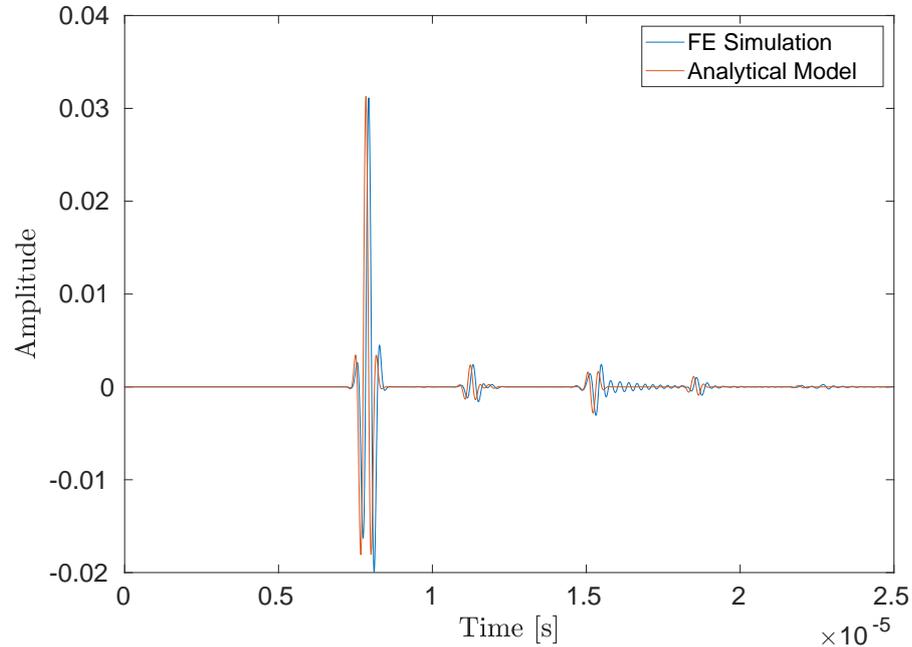


Figure 6.3: An A-scan comparison of a plane wave generated from a finite element simulation (blue line) and the corrected semi-analytical model (orange line) in the Pitch-catch mode. The wave is travelling through the 4 layered medium shown in Figure 5.1

turn and given that each element is smaller than or at least commensurate with the wavelength, an approximately spherical wave emerges; at least in the far field. It is therefore important to compare the semi-analytical model with the A-scan produced from a finite element simulated spherical wave. The same 4 layered heterogeneous spatial domain was examined, with two 32 element arrays (with 2mm pitch), positioned above each other as in Figure 5.1, and an extended lateral domain (now with absorbing boundary conditions) to reduce side wall effects.

Multiplying the semi-analytical model calculations by a factor (max value of the simulated A-scan divided by the max value of the simulated A-scan) that removed the attenuation due to spherical spreading loss, the A-scan in Figure 6.4 for the vertical raypath is produced. It can be observed that there is a relatively good match up (objective function value of). Note that the large wave packet near the end of the A-scan is the wave that has travelled the depth of the medium three times. This will not be included within the inversion algorithm and informs the time range chosen for the subsequent tomography exercise.

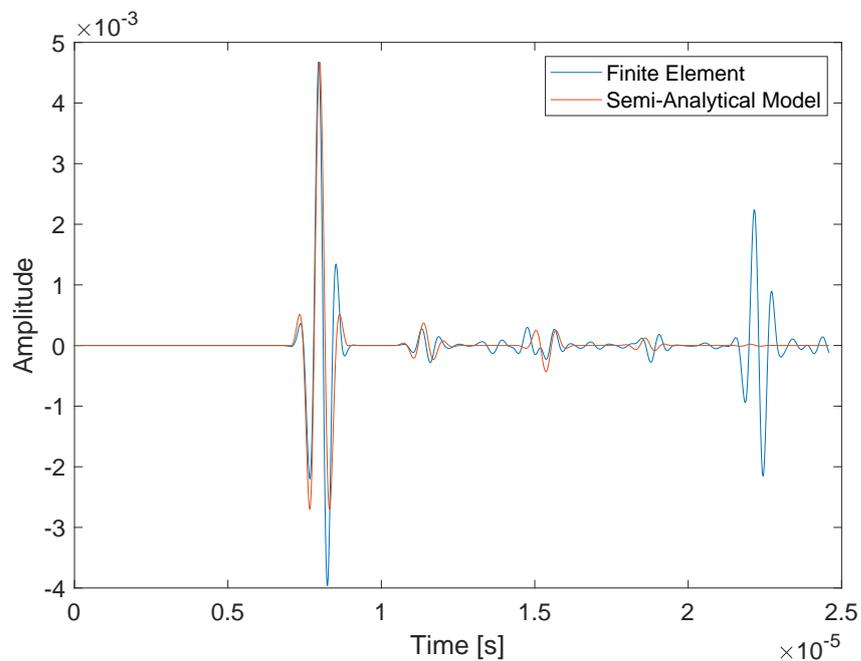


Figure 6.4: An A-scan comparison of a finite element simulation (blue line) and the equivalent semi-analytical model (orange line) for a single element in pitch-catch mode. The transmitter and receiver directly are above each other. The wave is travelling through the 4 layered medium shown in Figure 5.1

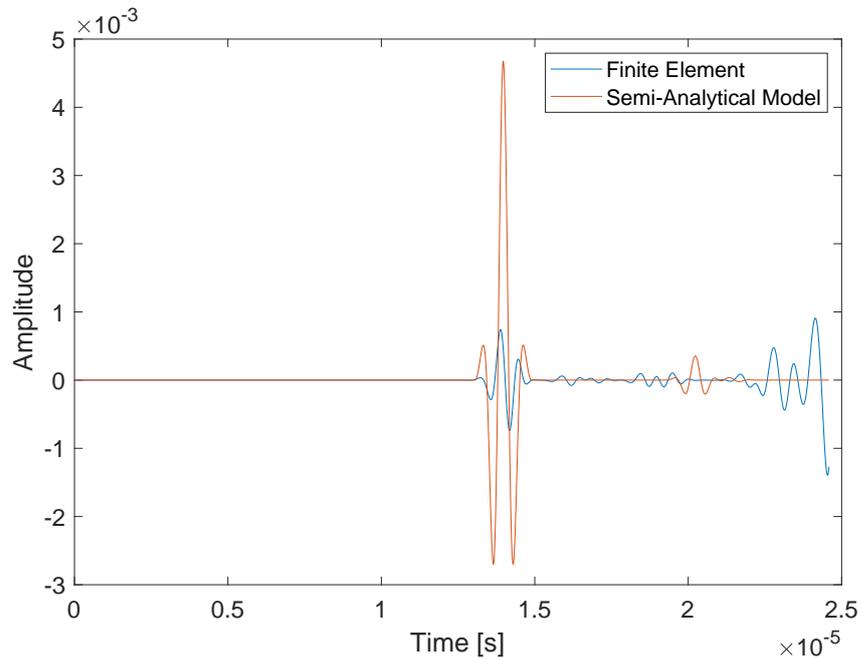


Figure 6.5: An A-scan comparison of a finite element simulation (blue line) and the equivalent semi-analytical model (orange line) from transmitter 1 to receiver 32. The wave is travelling through the 4 layered medium shown in Figure 5.1.

In Figure 6.5 there is a comparison between the finite element simulation and the analytical model for the situation where the transmitting element (transmitter 1) is now diagonally across from the receiving element (receiver 32). The analytical model assumes that the ray travels in a straight line between the transmitter and the receiver but of course this is now not the true as the reflections within each layer means that the ray is composed of a series of segments in a range of directions (see Figure 6.6 - note the figures ignores refraction between layer transmission). The analytical model was given the same scaling as above but there is still an amplitude difference. This difference most likely due to the

wave path from 1 to 32 is the “tail” of the spherical wave and has a much smaller amplitude. There also appears to be a shear wave coming in at the end of the A-scan.

In Figure 6.7 an A-scan comparison for the waves transmitting on element 1 to receiving element 21 is plotted. It can be observed that there is a large wave packet around 18μ secs which is not included in the analytical model. From analysis of plots of the surface displacement, this wave packet appears to be a shear wave that is generated by the transmitting array element. Figure 6.8 plots the displacement, at four points in time, to highlight the generation of a shear wave plot (a) that travels through the medium and then reaches the receiver plot (d). As the analytical model only considers longitudinal waves this accounts in part for the difference between the A-scans in Figure 6.7. To reduce the production of a shear wave by the transmitting element the finite element

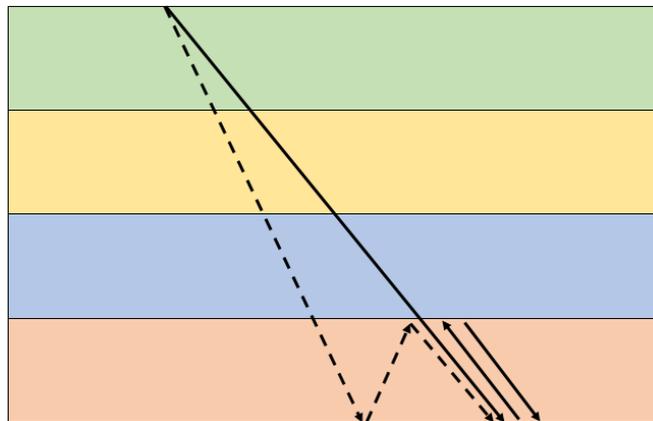


Figure 6.6: Two raypaths travelling through a layered medium with a reflection in the final layer. Solid line indicates the raypath assumed by the one dimensional semi-analytical model. Dashed line indicates the raypath taken by the spherical wave (ignoring refraction) within the finite element simulation.

simulation was changed. Instead of each element in the array having a finite size (one commensurate with the wavelength), a point source, that emits a spherical pressure wave, was used as the transmitter. The displacement of the material caused by a transmission from this point source is shown in Figure 6.9 which plots the acoustic pressure travelling through the medium. After implementing the point source transmitters in the finite element simulation the A-scans produced exhibit smaller amplitude shear waves (see Figure 6.10).

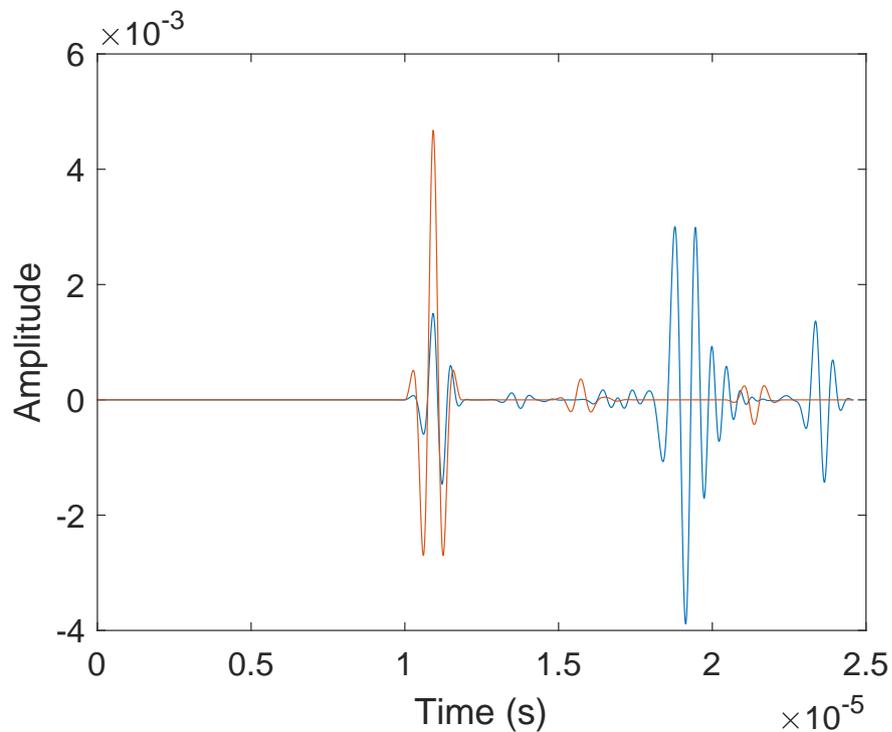


Figure 6.7: An A-scan comparison between a finite element simulation (blue line) and the equivalent A-scan from semi-analytical model (orange line) from transmitter 1 to receiver 21.

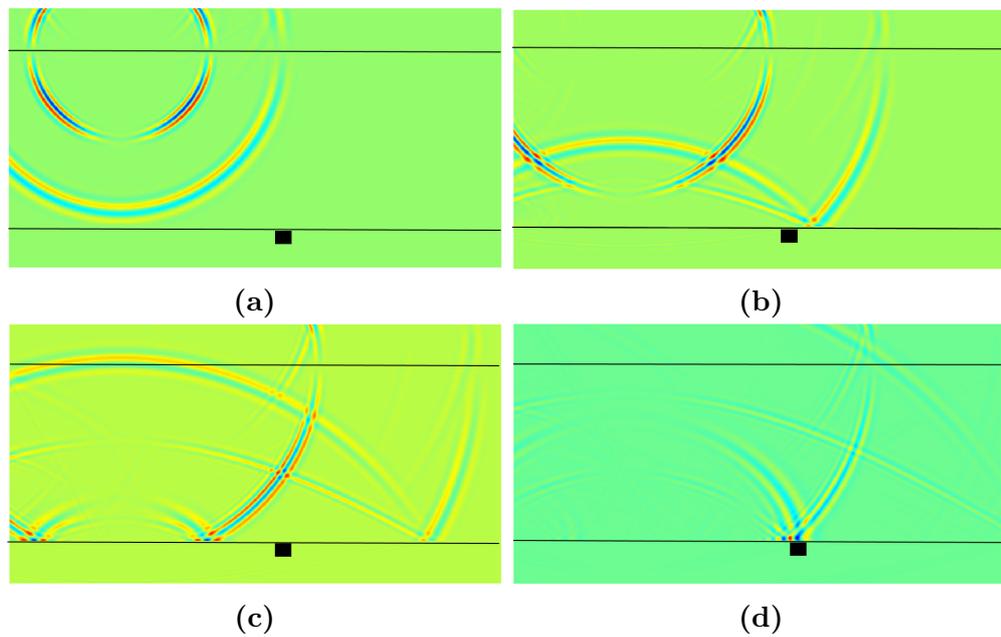


Figure 6.8: Plots of the in-plane displacement at four points in time ($8.5\mu s$, $12.5\mu s$, $16.5\mu s$, $20\mu s$) which highlights the shear wave being generated by the transmitter. The black square denotes the location of receiver 21.

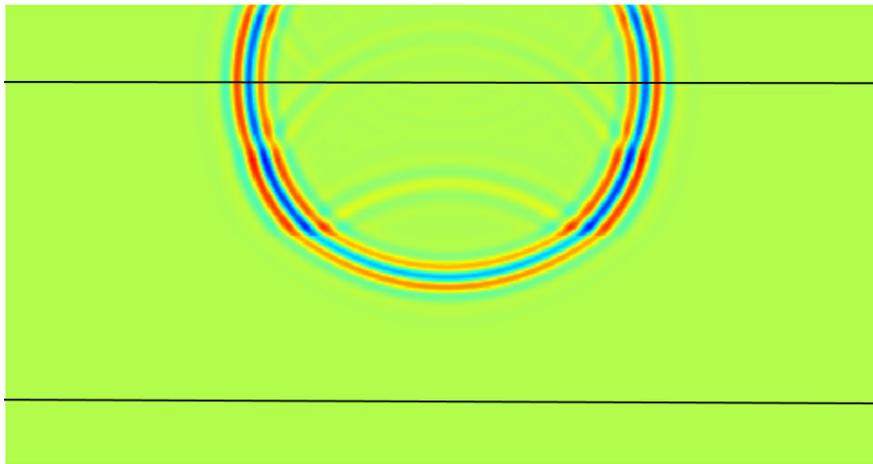


Figure 6.9: Plot of the acoustic pressure wave being generated by a point source.

To examine the similarity of the observed data over the full array, the log based objective function was calculated for every pair of transmit-receive elements and is plotted as a heatmap (Figure 6.11). Recall from Chapter 5 that this log based

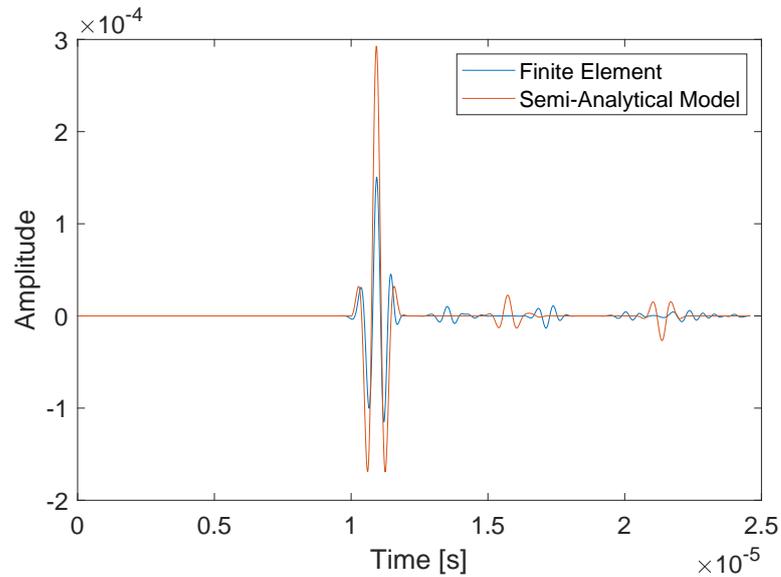


Figure 6.10: An A-scan comparison between a finite element simulation (blue line) and the equivalent A-scan from semi-analytical model (orange line) from transmitter 1 to receiver 21. The transmitters are now point sources.

calculation is used in the inversion algorithm to give greater weight to the smaller amplitudes in the A-scans. Figure 6.11 shows that the objective function values are now higher along the main diagonal. Taking the log of the signal means that the smaller signals (say due to shear waves), are given a far greater importance. These additional features in the data arising from the enhanced physics of the finite element simulation were not present in the previous chapter.

The last step before the tomographic inversion is to look at how the log objective function varies as the model is perturbed. It is important that the objective function is sensitive to changes in the material to be reconstructed. It was found

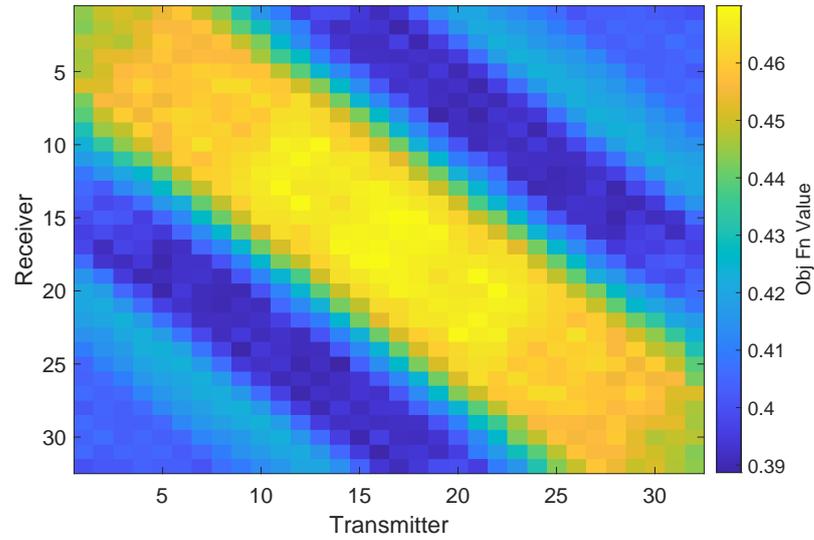


Figure 6.11: Log based objective function heatmap between the A-scans generated from the semi-analytical model and the finite element data using a point source transmitting element for the layered model in Figure 5.1.

that the objective function is insensitive to quite large perturbations. Despite these discouraging results, in the next section, an inversion using this objective function will be attempted.

6.2 Reconstruction of a Layered Medium from Simulated Data using the Log Based Objective Function

In this section, the inversion of the layered media (Figure 5.1) is attempted using the log based objective function (as used in Chapter 5) on data from a finite element simulation. The rj-MCMC was run for a total of 600,000 samples, where the first 400,000 were discarded (burn-in period) and the remaining models were sampled at an interval $\kappa = 100$. The initial model used to initialise the algorithm was composed of eight Voronoi cells (in a rectangular grid pattern as before), each assigned a wave speed of 5600 m/s. The algorithm had the following prior uniform distributions: the number of cells must lie in the range of 5 to 250 cells; the wave speed for a Voronoi cell must lie in the range 5600 ± 700 m/s; and the noise parameter must lie in the range 0.001 and 0.5. The standard deviation for the proposal distributions σ were: on *wave speed* 600m/s, and 150m/s for the delayed rejection; on *cell position* 30%, and 3% for the delayed rejection; on the wave speed for a new cell generated by a *birth* 100m/s; on the *noise* parameter 0.03. These are summarised in Table 6.1

Table 6.1: Proposal distributions in the rj-MCMC framework for the case where finite element simulated data is being inverted.

Perturbation	Range	σ	σ^{dr}
Wave speed	5600 ± 700 m/s	600 m/s	150m/s
Seed position	–	30%	3%
Birth	–	100 m/s	–
Noise	0.001 - 0.5	0.03	–

With these priors, the algorithm had an overall acceptance rate of 47.4%. The acceptance rates were: 88.1% for the wave speed; 47.7% for the position; 29.8% for the birth and death of a cell; and 29.7% for the noise. These values are very high with an overall acceptance rate of 23-44% is desired.

The average objective function value (over all transmit-receiver pairs), as the algorithm iterates through the samples, is shown in Figure 6.12, which shows a sharp initial decrease and then steady objective function value of 0.35. This is lower than the calculated objective function value of 0.475 between the modelled data of the known Voronoi diagram and the observed data. The number of Voronoi cells, as the algorithm iterates through the samples, is shown in Figure 6.13 (a). This shows a steady number of cells around 200 cells. This is also shown in Figure 6.13 (b) which shows an approximate normal distribution with a mean of around 200 cells for the last 200,000 samples. This is a very large number, most likely that the algorithm is trying to account for the additional waves present within the full wavefield finite element simulated data but absent in the semi-analytic ray-tracing model.

Examining the posterior distributions of the wave speed at each point in the domain, the median material map was reconstructed, (see Figure 6.14), and the standard deviation map (see Figure 6.15). These both clearly show that the algorithm failed to reconstruct the material map. It is believed that by taking the log of the A-scans this emphasises small amplitude waves within the objective function calculation. Since the finite element simulated data contains much more of the physics (for example, wave refraction, and shear waves) it has a richer set of

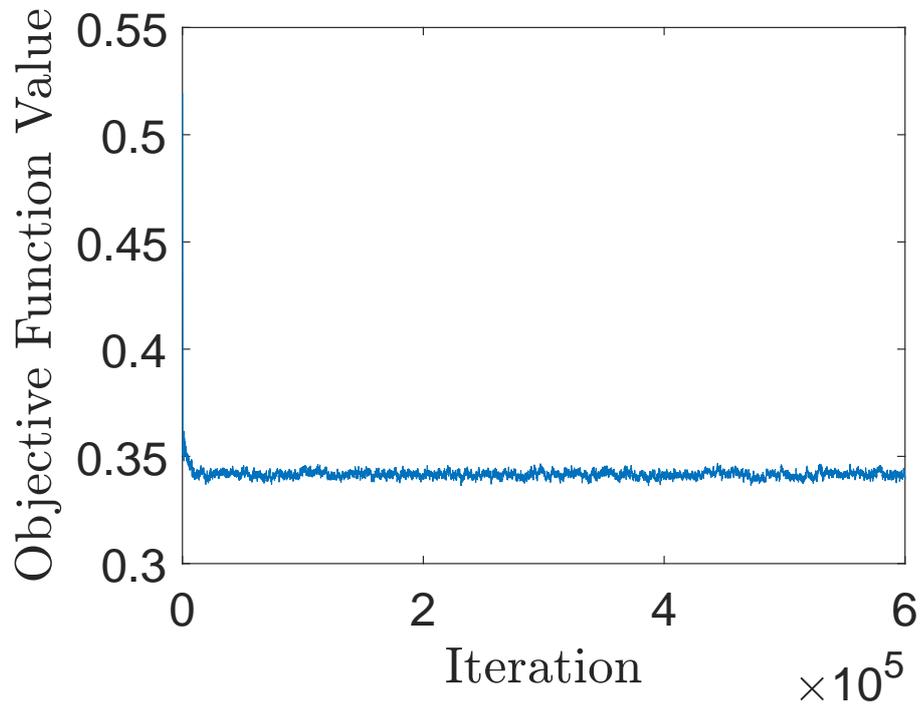


Figure 6.12: The average objective function value (Equation (2.1)) as the rj-MCMC algorithm iterates through each sample.

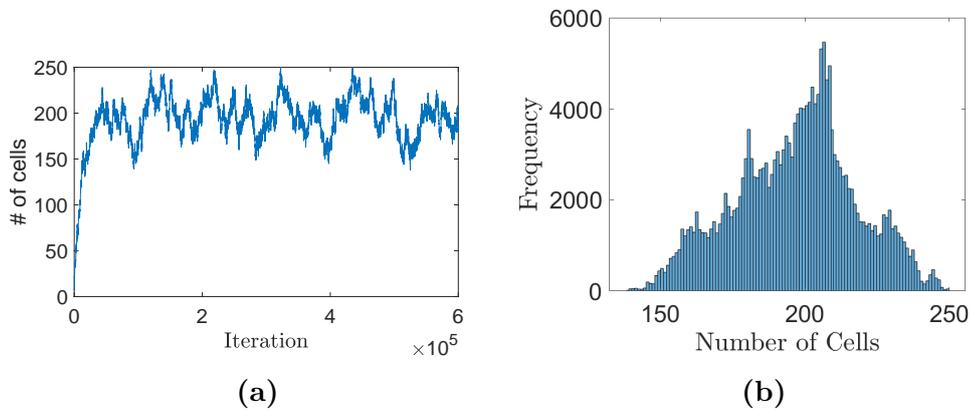


Figure 6.13: Plot (a) is the number of cells at each iteration of the rj-MCMC algorithm and plot (b) is the histogram of the number of cells of the last 200,000 samples.

small amplitude waves and hence has a poor correlation with the semi-analytical, ray-tracing modelled data. This leads to a larger objective function value and thus makes it hard to accurately reconstruct the medium. In future work, a new analytical model that contains this additional physics may be required. For now, to overcome this issue, a new objective function will be needed that is more robust. Before one is developed a flaw reconstruction is attempted and quantified.

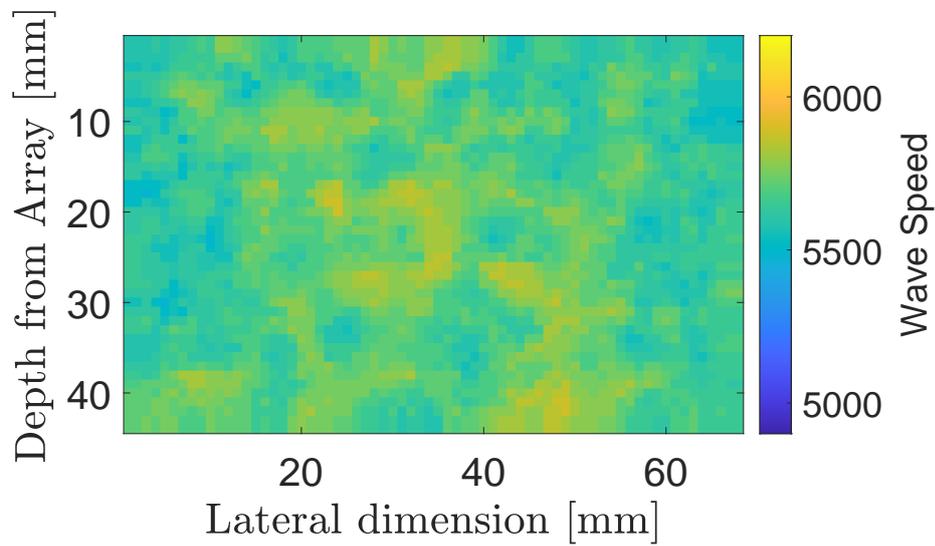


Figure 6.14: Material map obtained from the median of the posterior distribution on wave speed across the domain.

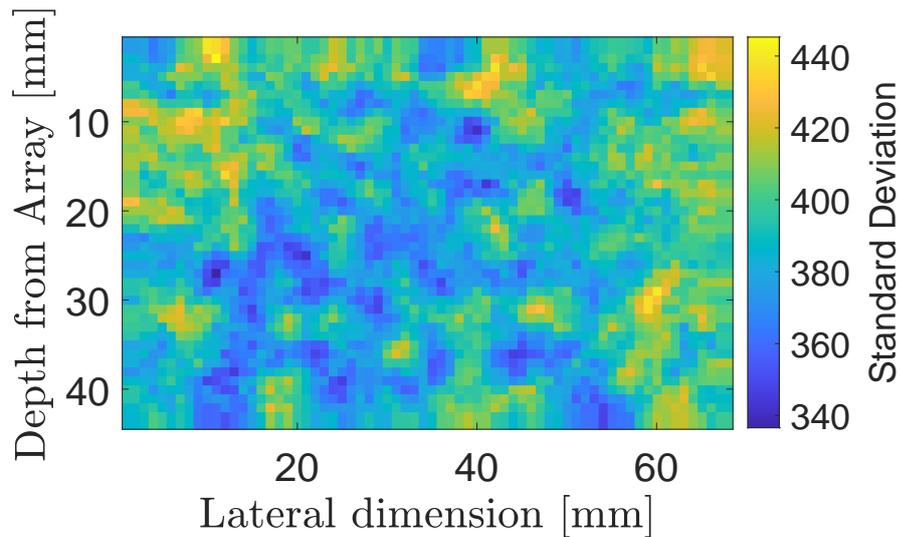


Figure 6.15: Standard deviation of the posterior distribution on wave speed across the domain.

6.2.1 TFM+ from the Reconstructed Material Map from Simulated Data

Using the median of the posterior distribution, shown in Figure 6.14, as the reconstructed material map within the TFM+ algorithm gives the flaw reconstruction in Figure 6.16 recall there are three embedded defects - see Figure 5.1. From this, three flaws can be detected, however, the third flaw could go unnoticed as it appears to have a similar magnitude as the surrounding noise. The signal-to-noise ratio for the third flaw is 15.9 dB, calculated by the maximum intensity of a rectangle around the flaw and a mean value of a rectangle that is placed at the same depth but centred at width 15mm. This is due to the additional noise that is present in this image compared to the standard TFM image (see Figure 5.2) and the TFM+ image with the known map (see Figure 5.3) which exhibits SNRs

of 17.5 dB and 23.6 dB respectively. The error in the flaw location is 0.1mm. From the ROC curves (see Figure 6.17) it can be observed that the performance is slightly better than the TFM (based on a homogeneous medium assumption) and the TFM+ (with the known map); the latter is slightly better at lower FPRs. This is quantified by the area under the curve (AUC) value of 0.901 for TFM+ with the known map, 0.875 for TFM+ with the reconstructed map, and 0.830 for TFM with a homogeneous map. However, this value uses the full reconstruction meaning there are two more flaws that can create some false positives. Calculating the ROC curves for just the fourth layer focuses attention on the third flaw, where the TFM performs better than TFM+ with the reconstructed map. All these results are summarised in Table 6.2.

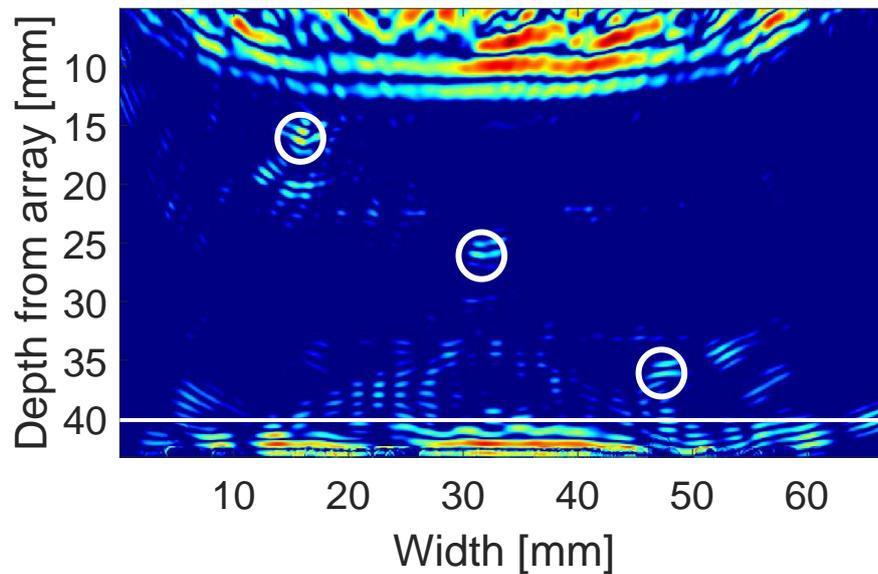


Figure 6.16: A TFM+ image of the 3 inclusions, using the material map obtained from the median of the posterior distribution as shown in Figure 6.14.

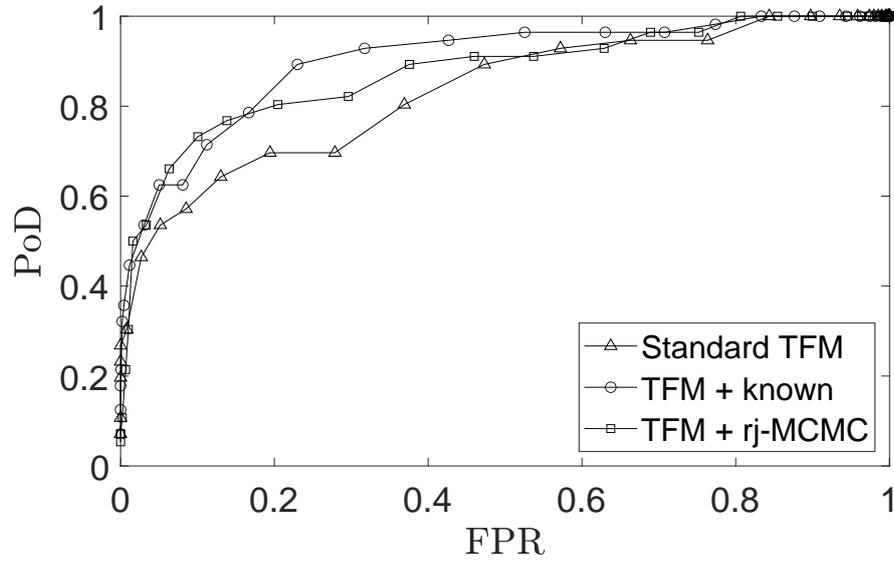


Figure 6.17: ROC curves from Figure 5.6 and the ROC curve generated from the TFM+ image Figure 6.16 calculated over a domain that contains all three inclusion.

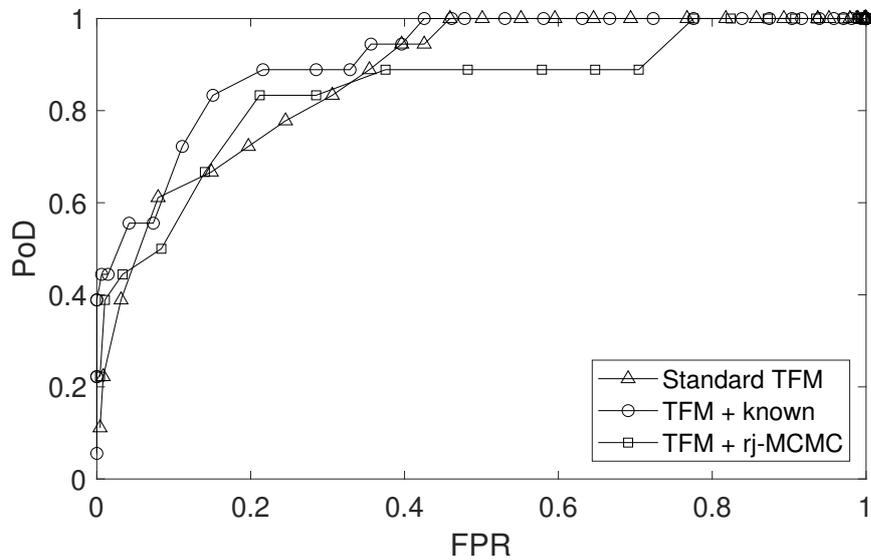


Figure 6.18: ROC curves from Figure 5.6 and the ROC curve generated from the TFM+ image Figure 6.16 calculated over a domain that only contains the fourth layer.

Table 6.2: Comparison of the quality metrics between the TFM image based on a homogeneous map, the TFM+ image based on the known map, and the TFM+ image based on the Full Waveform Inversion (FWI) reconstructed map (Figure 6.14).

Method	SNR	Flaw Position Error	AUC	AUC (4^{th} Layer)
TFM	17.5	1.315	0.830	0.876
TFM+	23.6	1.300	0.901	0.915
TFM+ FWI	15.9	0.100	0.875	0.845

6.3 Discussion

In this chapter, the A-scans generated from the semi-analytical model were compared to the A-scans generated from a finite element plane wave simulation. From these it was found that the modelled waves and simulated waves had different dimensions. Adjustments to the model was made such that it was of a pressure wave and better reflected what is being produced within the finite element simulation. This new semi-analytical model was then compared to a spherical wave within the finite element simulation. From these comparisons, adjustments to the semi-analytical model were implemented to account for the attenuation due to spherical decay, and the finite element simulation was adjusted to limit the presence of shear waves. An inversion was then attempted using the simulated data from finite element simulation as the observed data.

Although there was some improvement in the flaw detection, the log based objective function failed to create a material map approximation that contained any features of a layered medium. The main problem is that the small amplitudes arising from the additional physics contained in the finite element simulated A-scans which are then given significant weighting in the log of the A-scan is calculated.

In order to progress it would appear that a new objective function has to be created. This should take into account the waves that arrive at the receiver after the initial wave arrives but focus on the subset of the arriving waves that the ray-tracing, semi-analytical model predicts.

Chapter 7

Tomographic Reconstruction using a Time Windowed Objective Function

From the previous chapters, it is evident that the initial peak in the received A-scans are dominating the objective function Equation (2.1). This is not ideal as critical information on the structure of the layered medium (Figure 5.1) is contained in the smaller peaks that follow. It is therefore sensible to explore other objective functions. For the exploration of a new objective function the same scenario of a medium with four layers (Figure 5.1) will be investigated as in Chapter 5 and Chapter 6. The simplicity of the model allows investigation at every step of the rj-MCMC and facilitates the comparison of the log based PCC and any new proposed objective functions.

To explore a new objective function, the natural place to start is with examination of a simple L_1 or L_2 -norm, as it is very common in the literature and has produced positive results [88, 120, 121]. The advantage is the ease at which this can be implemented and the low computational cost. Unfortunately, the L_1 -norm objective function has similar issues as the PCC objective function. Firstly, it is sensitive to a small timeshift and prone to cycle skipping; the phenomena where peaks and troughs in a wave are matched but out of phase by one period. This is demonstrated in Figure 7.1 which plots the L_1 -norm of the difference between two identical Ricker wavelets $f(s)$ and $g(s - t)$ (for a fixed (s)) as they transition in and out of phase. Secondly, the initial wave dominates the objective function value. This is simply due to the relative amplitudes of the waves that follow the initial wave which is generally the largest. This is highlighted in Figure 7.2 which plots (a) an A-scan example and (b) the L_1 -norm as a function of the length of time of the A-scan taken. It can be observed here that the initial wave accounts for over 50% of the objective function value even though the information about the layered medium is contained within the later arriving waves. Therefore to reconstruct the layered structure, there needs to be a greater importance put on the later arriving waves.

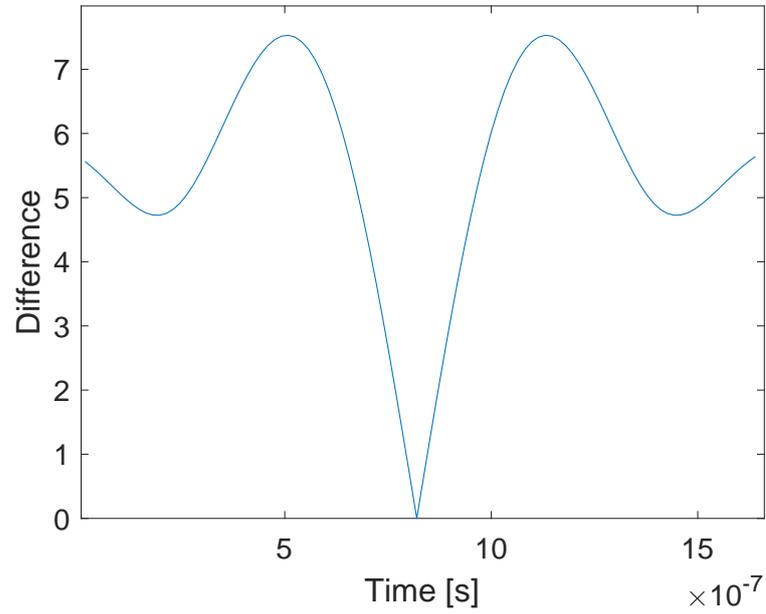


Figure 7.1: The L_1 -norm of the difference between two Ricker wavelet functions $f(s)$ and $g(s - t)$ (for a fixed s), moving in and out of phase.

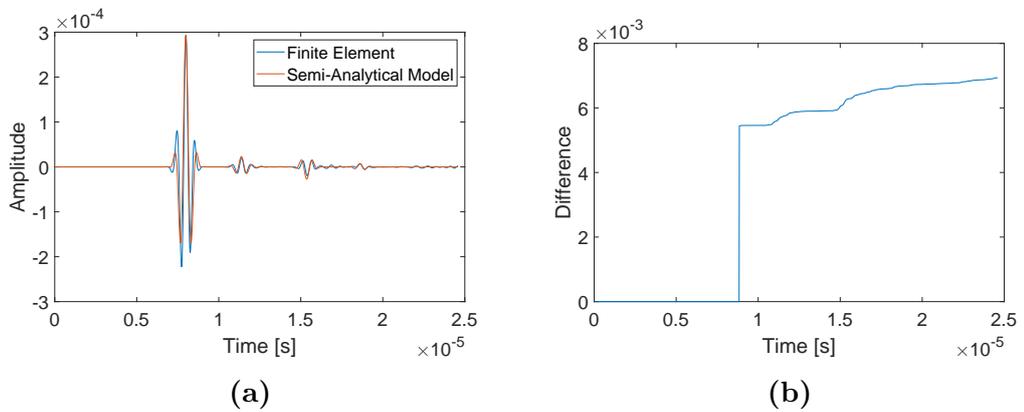


Figure 7.2: Plot (a) A-scan travelling through the layered medium and plot (b) the L_1 -norm as a function of the length of time taken of the A-scans from (a).

7.1 The Hilbert Transform

The first step to addressing the challenges associated with the objective function is to use the envelope of the received wave by taking the Hilbert transform $H(f(t))$, [122]. The Hilbert transform $H(f(t))$ of a real signal is defined as

$$H(f(t)) = \frac{1}{\pi} C \int_{-\infty}^{+\infty} \frac{f(\tau)}{t - \tau} d\tau. \quad (7.1)$$

where C stands for the Cauchy principal value. The Hilbert transform produces a real part (\Re) and an imaginary part (\Im). The envelope of the wave is then obtained by

$$E(t) = \sqrt{\Re(H(f(t)))^2 + \Im(H(f(t)))^2}, \quad (7.2)$$

where H denotes the Hilbert transform (see Figure 7.3).

To check for the suitability of this objective function, two wavelets were moved in and out of phase $f(s)$ and $g(s - t)$ (for a fixed s). This produced a smooth convex output which means that the envelope would stop cycle skipping (see Figure 7.4). The envelope for a full A-scan was then obtained (see Figure 7.5).

This still has the issue of the initial peak dominating the objective function value due to its amplitude being far greater than the other peaks present in an A-scan. To overcome this a time windowed approach was implemented. This takes a window around the regions of interest (first, second, and third peaks) and then normalise the window by the maximum value of the observed data in

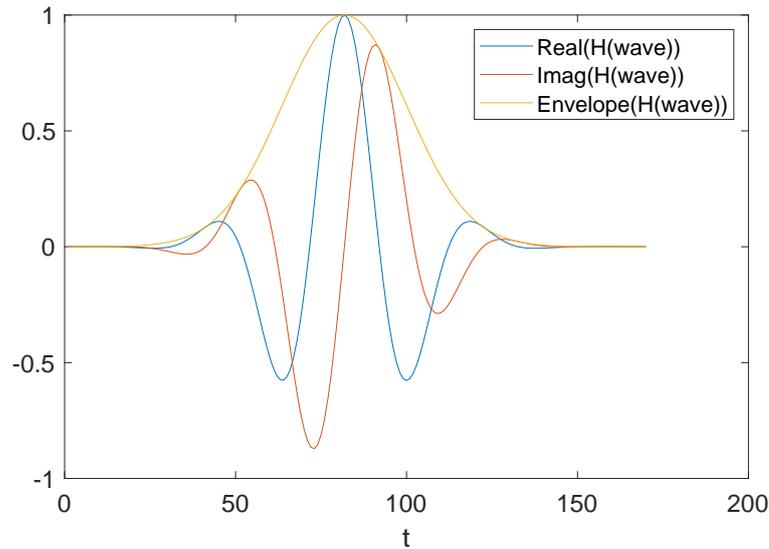


Figure 7.3: Real and Imaginary component of the Hilbert transform (and the envelope) of a Ricker wavelet.

this window. This means that each peak has an equal weighting on the objective function value. The number of windows chosen was three for the three main peaks that can be observed in a typical A-scan.

Starting with two A-scans (one from the semi-analytical model (forward model) and the other from finite element simulation (observed data)) travelling through the layered medium (see Figure 7.6), the Hilbert transform envelope for these are obtained (see Figure 7.7). The three most dominant peaks are found and a window around each one is taken and normalised by the maximum value of the observed (simulated) data within each window (see Figure 7.8). The window size is chosen to be the same size as the input waveform that is used within the inversion, for this case the Ricker wavelet.

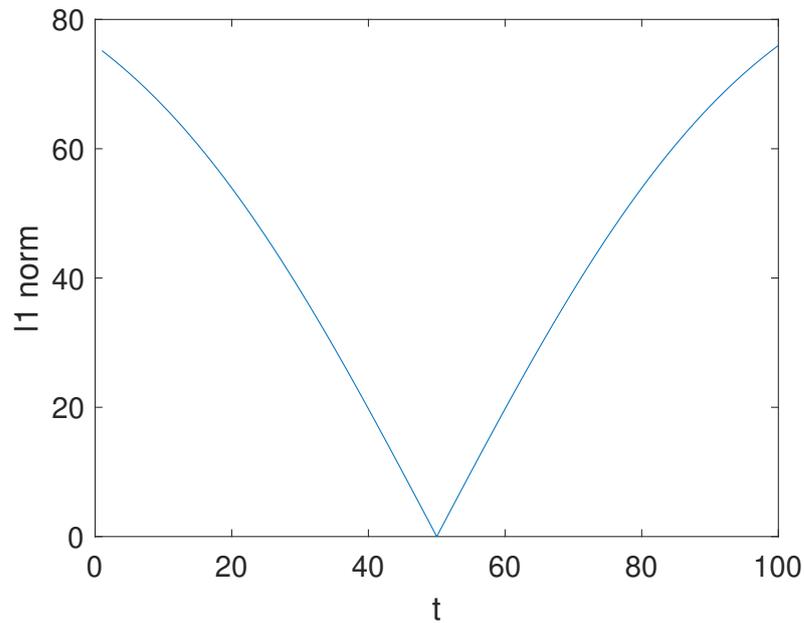


Figure 7.4: The L_1 -norm for the difference between two envelopes of a Ricker wavelet function $f(s)$ and $g(s - t)$ (for a fixed (s)) moving into and then out of phase.

To check how sensitive the proposed objective function (L_1 -norm of the Hilbert transform envelope and time windowing of the three tallest peaks) is, two identical A-scans generated from the semi-analytical model were compared. In one of the A-scans the second largest peak's amplitude was increased by 10%. This led to an objective function value increase of 4.88 from zero (see Figure 7.9). The objective function matrix heatmap generated by varying the seed location and wave speeds of a Voronoi with 8 cells (see Figure 7.10) shows a good sensitivity and suggests this could be a promising objective function.

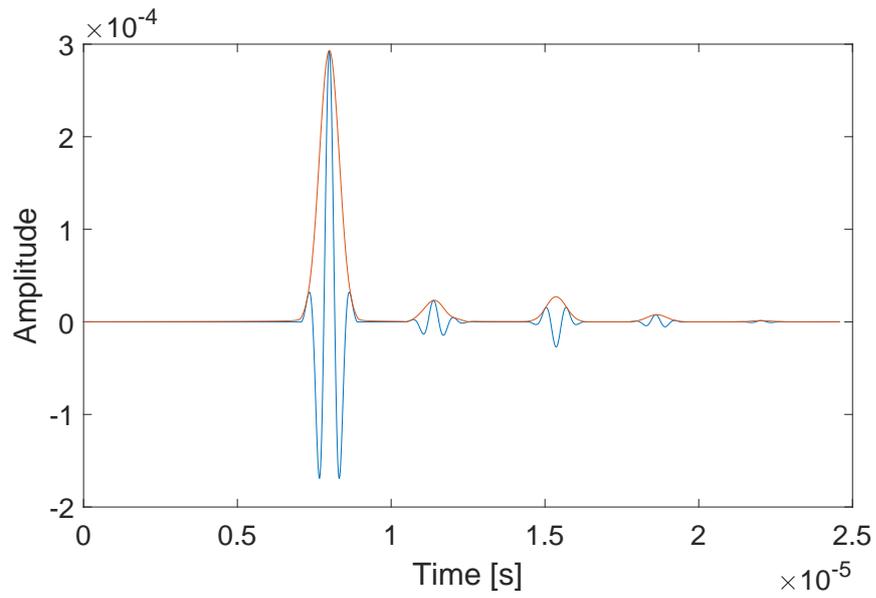


Figure 7.5: An A-scan of a wave travelling through the layered medium (blue line) with its associated envelope (orange line).

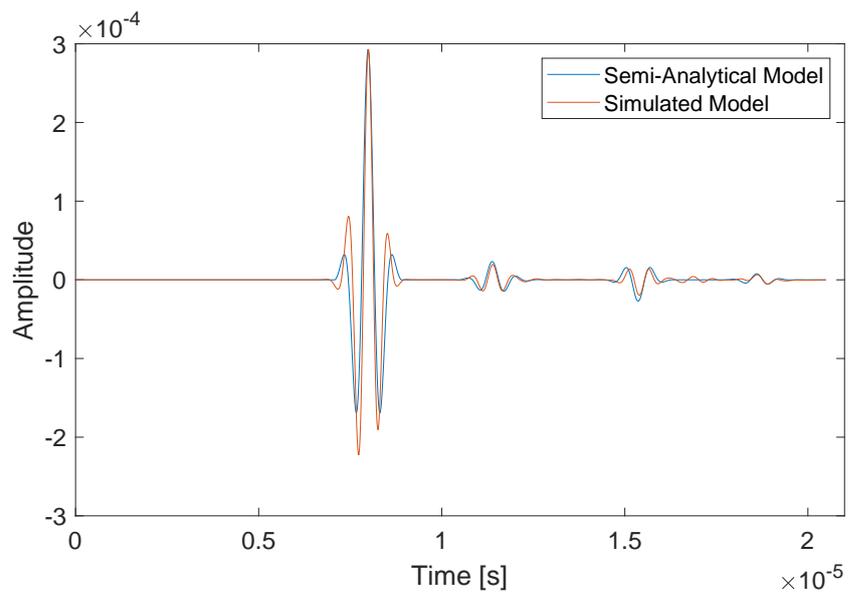


Figure 7.6: A-scans from a layered medium generated by the semi-analytical model (blue line) and the equivalent finite element simulation (orange line).

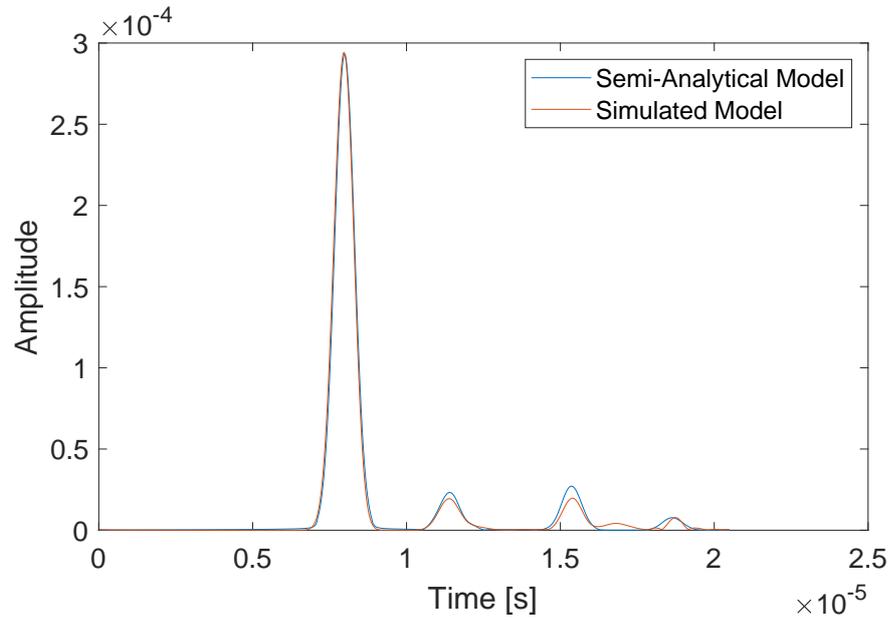


Figure 7.7: Envelopes of the A-scans generated by the semi-analytical model (blue line) and the equivalent finite element simulation (orange line).

7.2 Implementing the Hilbert Transform Envelope

To calculate the Hilbert transform the code from the supplementary materials in [151] was adapted. The Hilbert transform envelope of a vector \underline{x} (the discretised form of an A-scan) was created by

1. Calculating the discrete FFT of \underline{x} denoted by \hat{x} .

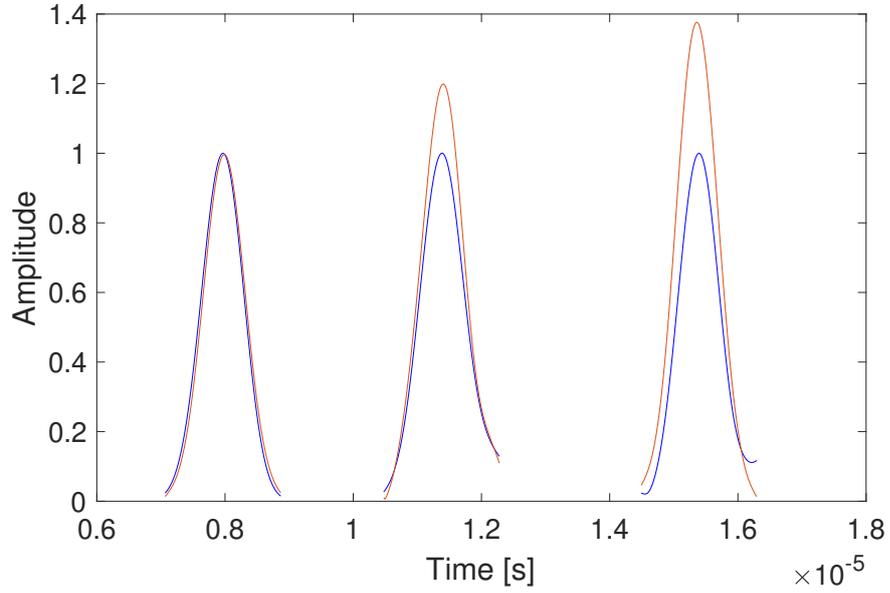


Figure 7.8: Normalised windows of the envelopes of the A-scans generated by the semi-analytical model (blue line) and the equivalent finite element simulation (orange line)

2. Create a vector \underline{h} where

$$h_i = \begin{cases} 1 & \text{when } i = 1 \text{ or } (n/2 + 1) \\ 2 & \text{when } i = 2, \dots, n/2 \\ 0 & \text{when } i = (n/2 + 2), \dots, n \end{cases}$$

3. Calculate $\{\hat{x}_i h_i : i = 1, \dots, n\}$, the element-wise product of $\hat{\underline{x}}$ and \underline{h} .
4. Calculate the inverse discrete FFT of $\hat{\underline{x}}\underline{h}$ denoted by \underline{xh} .
5. Calculate the envelope by taking the square root of the element wise product $\Re(\underline{xh})^2$ plus the element wise product $\Im(\underline{xh})^2$ (see Equation (7.2)).

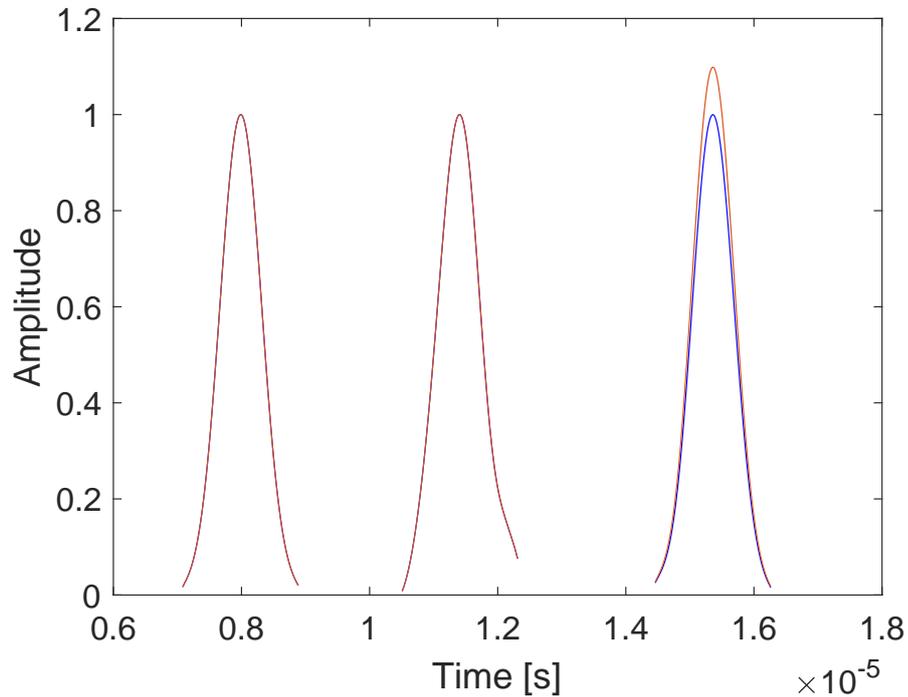


Figure 7.9: The windows of the envelopes of three largest peaks of an A-scan of the 4 layered medium unedited (orange line), and 10% increase is the second largest peak (blue line). Objective function value of 4.88.

One downside of the code was that the number of elements of the vector \underline{xh} needs to be 2^n for some natural number n . To make sure the third peak is included, choosing $n = 11$ ($2^n = 2048$ sample points) provided a sufficiently long enough A-scan.

An objective function matrix heatmap was obtained between the observed data generated by the semi-analytical model implemented in MATLAB (inverse crime) and data generated by the semi-analytical model as implemented in the tomography algorithm in FORTRAN (see, Figure 7.11). It can be observed that there are a few outliers. With these included the mean value is 6.34, however,

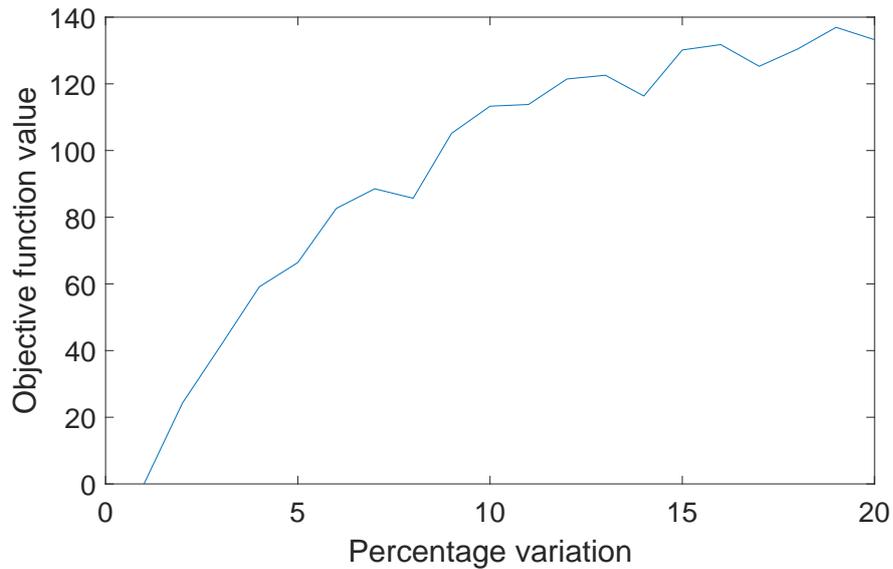


Figure 7.10: The Hilbert transform time windowed objective function value as the seed location and wave speeds are varied from 0-25% from their original value.

when they are excluded it is 5.76. They are therefore affecting the overall objective function value for this particular model but not by a large amount. This difference between the two approaches is simply down to the discretisation of the Voronoi diagram in the tomography algorithm leading to a wave speed mismatch at Voronoi cell boundaries. One way to minimise this would be to increase resolution (number of discrete blocks) that represents the models. However, the goal is not to find this known map. It is to find a set of Voronoi diagrams that form an ensemble to represent the known map. Therefore, these small errors should cancel out over many samples.

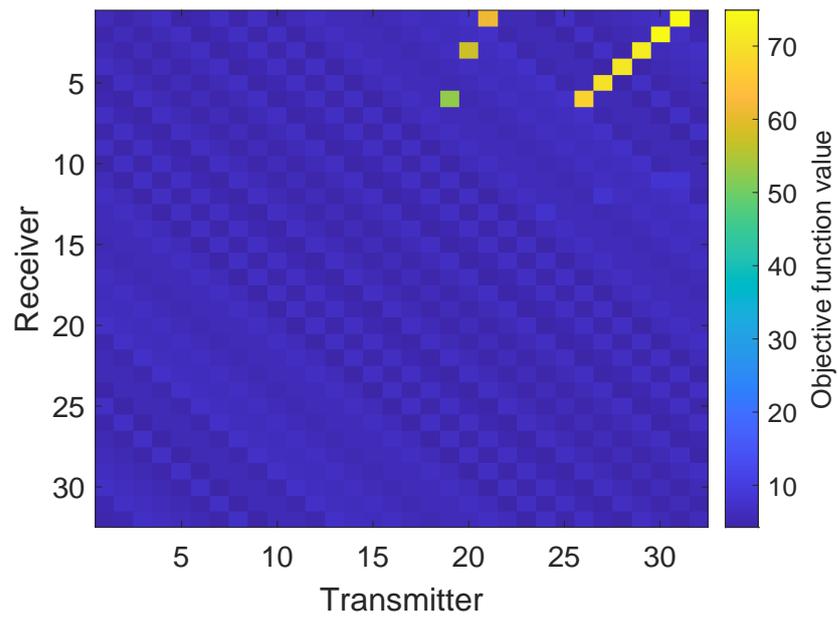


Figure 7.11: Objective function matrix heatmap between the observed data generated in MATLAB and the semi-analytical model data generated within the rj-MCMC.

7.3 Tomographic Reconstruction of a Layered Medium using a Hilbert Transform

In this section, a tomographic reconstruction of the layered medium (Figure 5.1) is attempted using observed data generated from the semi-analytical model as implemented in MATLAB (similar to the approach in Chapter 5) and the Hilbert transform time-windowed (HTTW) objective function.

7.3.1 Reconstruction over 300,000 samples

The rj-MCMC was run for a total of 300,000 iterations with a burn-in of 200,000. The remaining models were sampled at an interval of $\kappa = 100$. A homogeneous initial model composed of 8 cells with each Voronoi cell having a wave speed of 5600 m/s was used.

The algorithm had the following prior (uniform) distributions: the number of cells must lie in the range 5 to 300 cells; the wave speed for a Voronoi cell must lie in the range 5600 ± 700 m/s; and the noise parameter must lie in the range 1 and 120. The standard deviation for the proposal distributions (σ) were: on *wave speed* 200 m/s, and 60 m/s for the delayed rejection; on *cell position* 15%, and 3% for the delayed rejection; the wave speed for a new cell generated by a *birth* perturbation 50 m/s; on the *noise* parameter 6. These are summarised in Table 7.1.

Table 7.1: Prior and Proposal distributions set for the simulation using the semi-analytical model as the observed data.

Perturbation	Range	σ	σ^{dr}
Wave speed	5600 ± 700 m/s	200 m/s	60m/s
Seed position	–	15%	3%
Birth	–	50 m/s	–
Noise	1-120	6	–

The standard deviation, for the proposal distributions were, chosen such that the overall acceptance rate would lie between 23% and 44%. The limits of the noise parameter were chosen by setting the lower limit to be less than the objective function value between the FMC generated with the known map and the observed data and the upper limit being greater than the objective function value between the FMC generated for homogeneous map and the observed data. Note, that the noise value should be commensurate with the objective function value, hence, why there is a much larger range than in previous chapters. With these priors the algorithm had an overall acceptance rate of 41.0%. The acceptance rates were: 82.5% for the wave speed; 42.6% for the position; 18.7% for the birth; 18.6% for the death; and 28.5% for the noise.

The average objective function value, as the algorithm iterates through the samples, is shown in Figure 7.12. The initial objective function value is 108.0 and it converges to approximately 60 relatively quickly. The number of Voronoi cells as the algorithm iterates through the samples is shown in Figure 7.13 (a)

which converges to around 100 cells. The histogram for the last 100,000 iterations (omitting the burn-in period), is shown in (b) which shows a skewed distribution with a tail on the right.

From the moments and characteristics of the posterior distribution on wave speed at each point in the domain, material maps were constructed. Figure 7.14, shows (a) the mean, (b) the median, (c) the MAP, and (d) the standard deviation map of the posterior probability distributions. Two distinct layers can be observed in the bottom half of the material maps. This is supported by the high levels of deviation in (d) indicating layer interfaces. These two layers have wave speeds of approximately 5100 m/s and 5900 m/s respectively which is close to the correct wave speeds of 5010 m. and 5900 m/s respectively. In the top region, there is no clear layer interface however the wave speeds vary from 5700 - 6000m/s which is close the correct wave speeds of 5700 m/s and 6100 m/s. The layer interfaces in these reconstructions are not as well defined as in the reconstructions in Figure 5.27. However, the wave speeds are closer to the known wave speeds. The differences in wave speed from the known map is highlighted in Figure 7.15 which is a vertical cross-section of the medium at the halfway point (32mm). It shows that both the median and mean have a relative constant wave speed of 5800 m/s for the first 20mm and then shows the estimated wave speed, approximating the correct wave speed of 5010m/s for 10mm then correctly a wave speed of 5900m/s for 10mm. It also shows the MAP does relatively well in reconstructing the wave speeds. It shows a wave speed of 5700m/s for 1mm, a little less than 6100 m/s for 10mm, about 5010 m/s for 10mm and then a little more than 5900m/s for

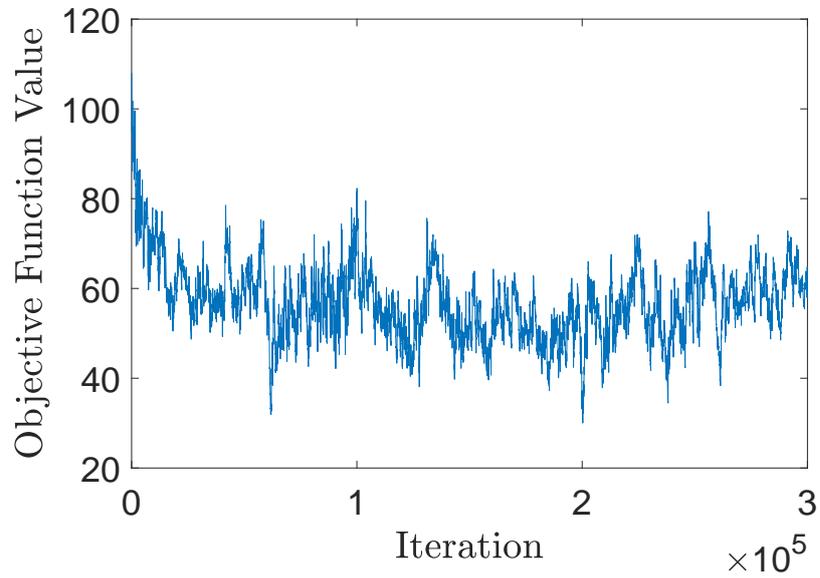


Figure 7.12: The average objective function value as the algorithm iterates through each Voronoi diagram sample.

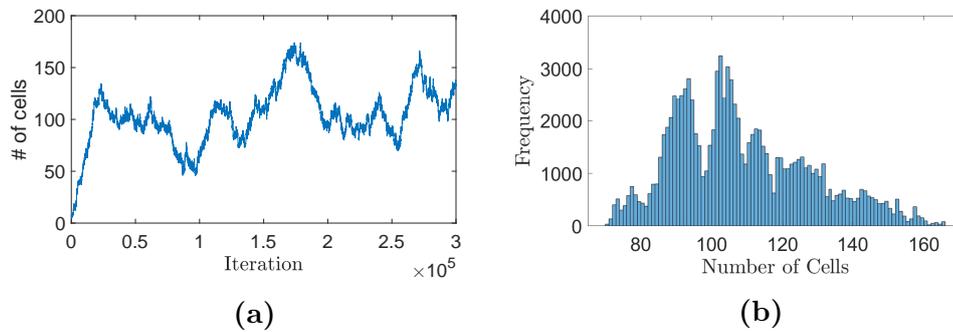


Figure 7.13: Plot (a) is the number of cells at each iteration of the rj-MCMC algorithm and plot (b) is the distribution of the number of cells from the last 100,000 samples.

10mm. However, this is one particular cross-section that gives a good wave speed profile. This does not seem to be consistent when observing other cross-sections at other x values. For example the cross section $x = 60$ would give a poor wave speed profile from Figure 7.14. The cross-section in the centre arises from the area with the most ray coverage and is often the best.

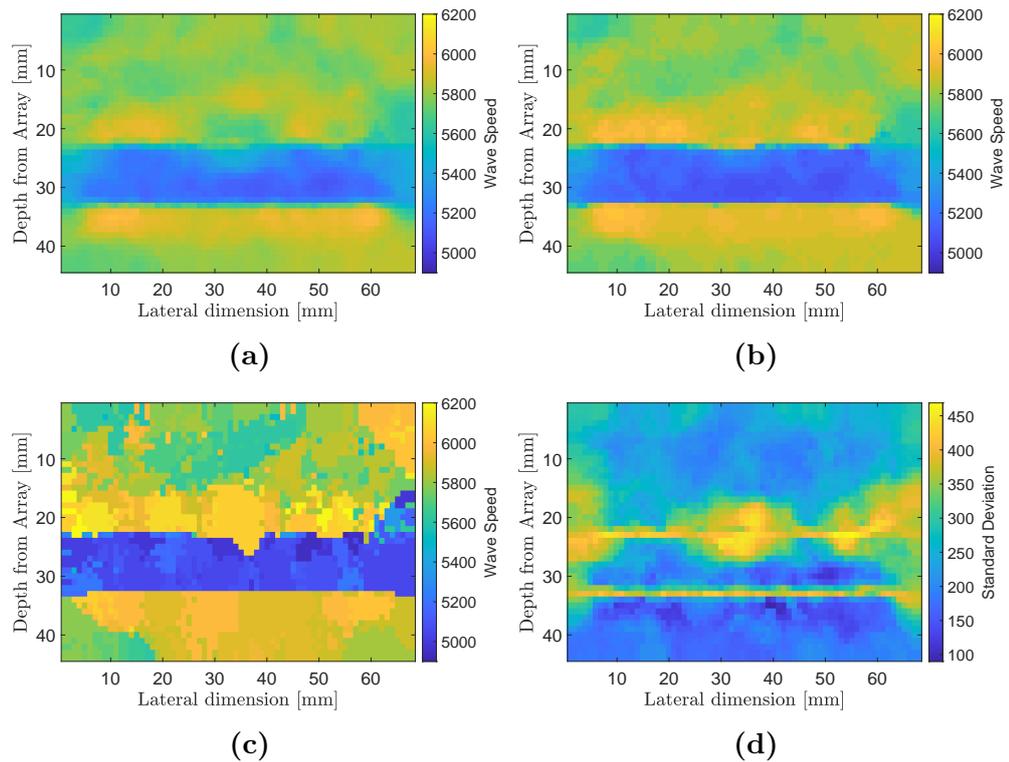


Figure 7.14: Material map reconstructions of the 4 layered medium. Plot (a) is the mean of the posterior distribution on wave speed at each point in the domain whilst plots (b)-(d) show the median, MAP, and the standard deviation respectively. Both the forward model and the observed data arise from the semi-analytical model.

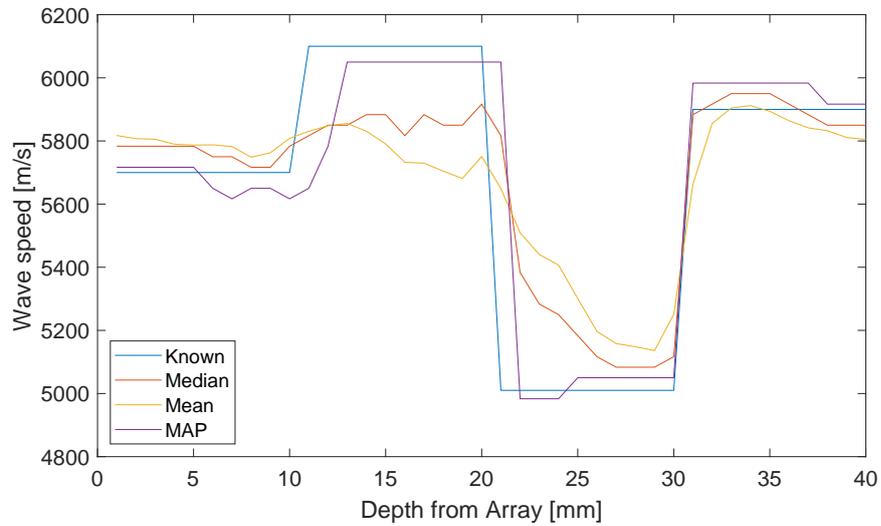


Figure 7.15: Wave speed profile of the moments of the posterior distributions (Figure 7.14) along the vertical cross-section at $x = 32\text{mm}$.

The variation of the wave speed profile along the lateral dimension of the medium is shown in Figure 7.16. Where plot (a) shows the mean, (b) the median, and (c) the MAP. This shows that the lower two layers are most consistent along the lateral dimension of the medium. The top two layers do not exhibit the same consistency. However, in the MAP profile (c) the average wave speed along the lateral dimension sits about 5800 m/s for the first 10 mm then starts to increase to about 6050 m/s. These values are closest to the known model. The interquartile range is also very small for the second layer, however, it should be noted that there is a large range when looking at the outliers.

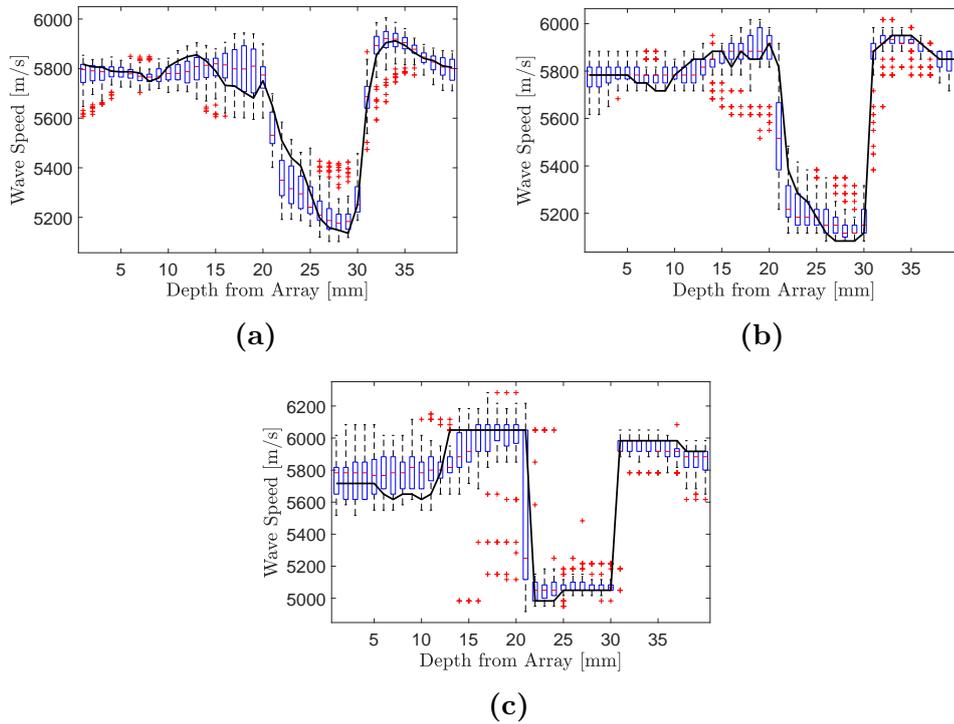


Figure 7.16: Wave speed profiles at $x = 32\text{mm}$ and box plots of the wave speed variation at every mm depth of the medium. Plot(a) is the median, plot (b) is the mean, and plot (c) is the MAP, of the posterior distributions from Figure 7.14.

The distance, in the location of the layer interfaces, from the reconstruction to the known map is shown in Figure 7.17. This shows that there is no interface identified for the first layer, the second layer is approximately correct but this is difficult to confirm due to the large region (centred at $(32\text{mm}, 25\text{mm})$) of a high standard deviation, and the third interface is off by about 1mm.

Now examining the posterior distribution of wave speeds at particular points in the tomographic reconstruction (Figure 7.18 and Figure 7.19 points (a) - (f)). These points are distributed such that there is one in each layer and two on the

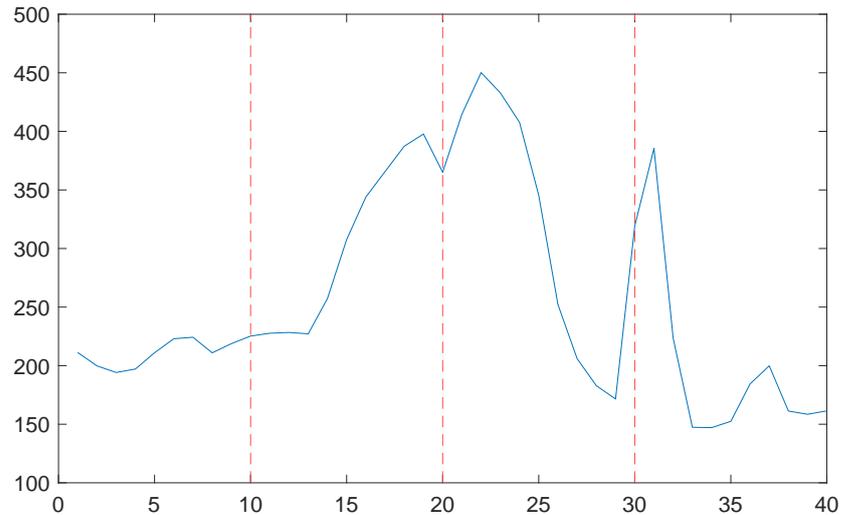


Figure 7.17: A vertical profile of the standard deviation (blue) - Figure 7.14 (d) - at $x = 32\text{mm}$ and the known layer interface (red dashed).

well reconstructed interfaces (on regions of high standard deviation) (see Figure 7.20, (a) - (f)). It is expected that there will be normal distributions within the layers and bimodal distributions at the interfaces. However the interface region is particularly small (just one or two pixels) and can be hard to precisely detect. It can be observed that: point (a) has a normal distribution at around 5900 m/s however there is a small tail to the left; point (b) has a bimodal distribution with peaks at 5100 m/s and 6000 m/s which indicates it is on the interface and thus exhibits values close to the known wave speeds of 5010m/s and 5900m/s on either side of the interface; point (c) has a skewed normal distribution centred at 5100-5200 m/s and a tail to the right; point (d) has a bimodal distribution with peaks at 5200 m/s and 5900 m/s indicating an interface; point (e) has a skewed distribution which has a dominant peak at 6100 m/s which is the wave

speed that was expected (however, the skewed distribution indicates that it has not fully converged in this region; point (f) has a normal distribution centred around 5800 m/s which is close to the correct wave speed but the large standard deviation indicates that the algorithm has not fully converged. Overall these distributions show that the model has not converged in the top half of the material map but indicates that it is starting to converge and should be run for longer (this is done in Subsection 7.3.3). In the second layer (20mm-30mm range) it is close to converging but not quite due to the high standard deviation of the distributions. The distribution in the last layer suggests that the algorithm has converged in this region.

One possible explanation lies in how the semi-analytical model builds the A-scan (see Chapter 2 for details). It is done with a backwards recursive relation starting at the backwall and recording the largest reflections layer by layer until it reaches the front wall. This means that the bottom layer has the potential to have the most internal reflections and therefore the greatest impact on the A-scan. Whereas the first layer has no internal reflections recorded. This could explain why this is the region with the greatest error in all the reconstructions.

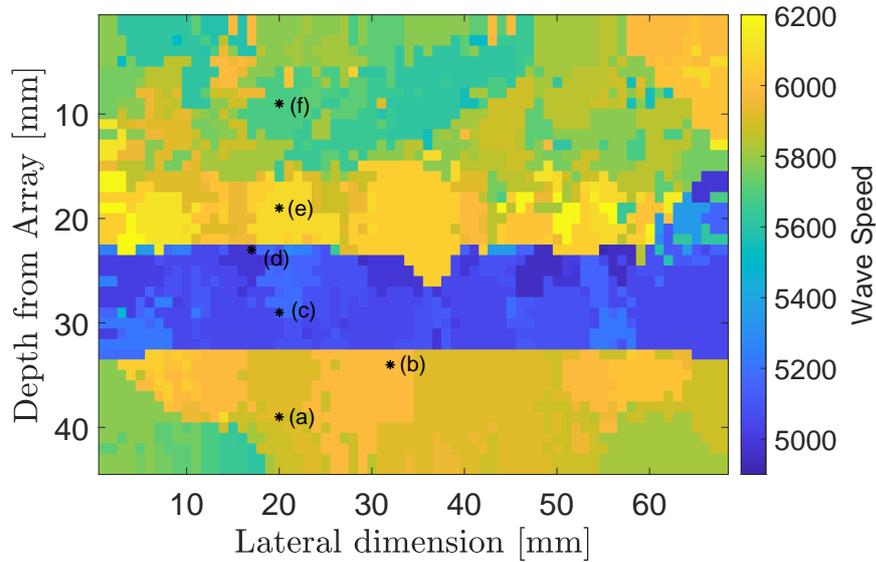


Figure 7.18: MAP of the posterior distribution of the tomographic reconstruction. Points where the wave speed distributions (Figure 7.20) have been analysed are shown.

7.3.2 TFM+ from the Reconstructed Map using the Hilbert Transform Time Windowed Objective Function

To further test this tomographic reconstruction, it is used to image the flaws embedded in the layered medium used in conjunction with the TFM+ algorithm. Using the same FMC data as in Section 5.1, the MAP of the posterior distribution map is used with the TFM+ algorithm, and the resulting flaw reconstruction is displayed in Figure 7.21. The last flaw is clearly detected which indicates an improvement from the standard material reconstruction using the TFM (Figure 5.2). The SNR is 22.3dB and the flaw position error is 0.2mm. These results are compared to the TFM and TFM+ with the known map are summarised in

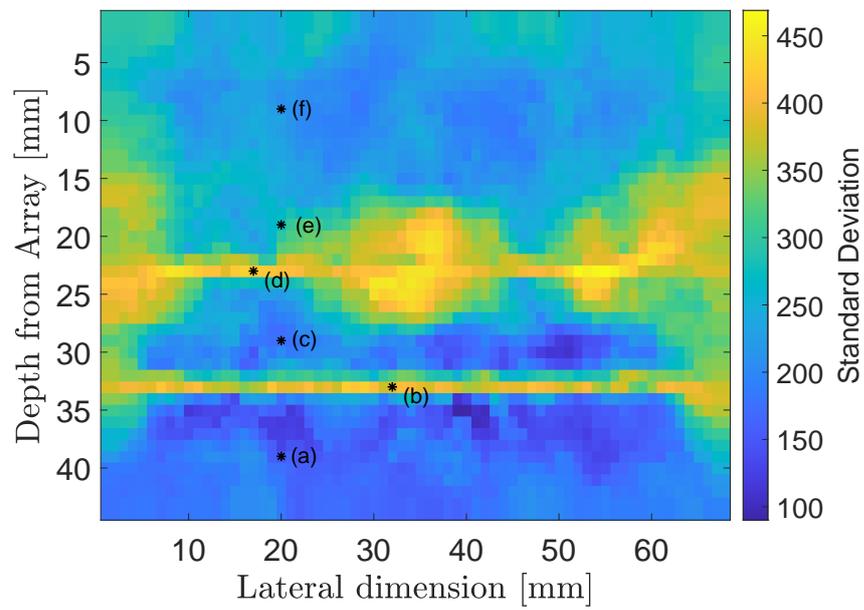


Figure 7.19: Standard deviation of the posterior distribution of the tomographic reconstruction. The points where the wave speed distributions (Figure 7.20) are analysed are also shown.

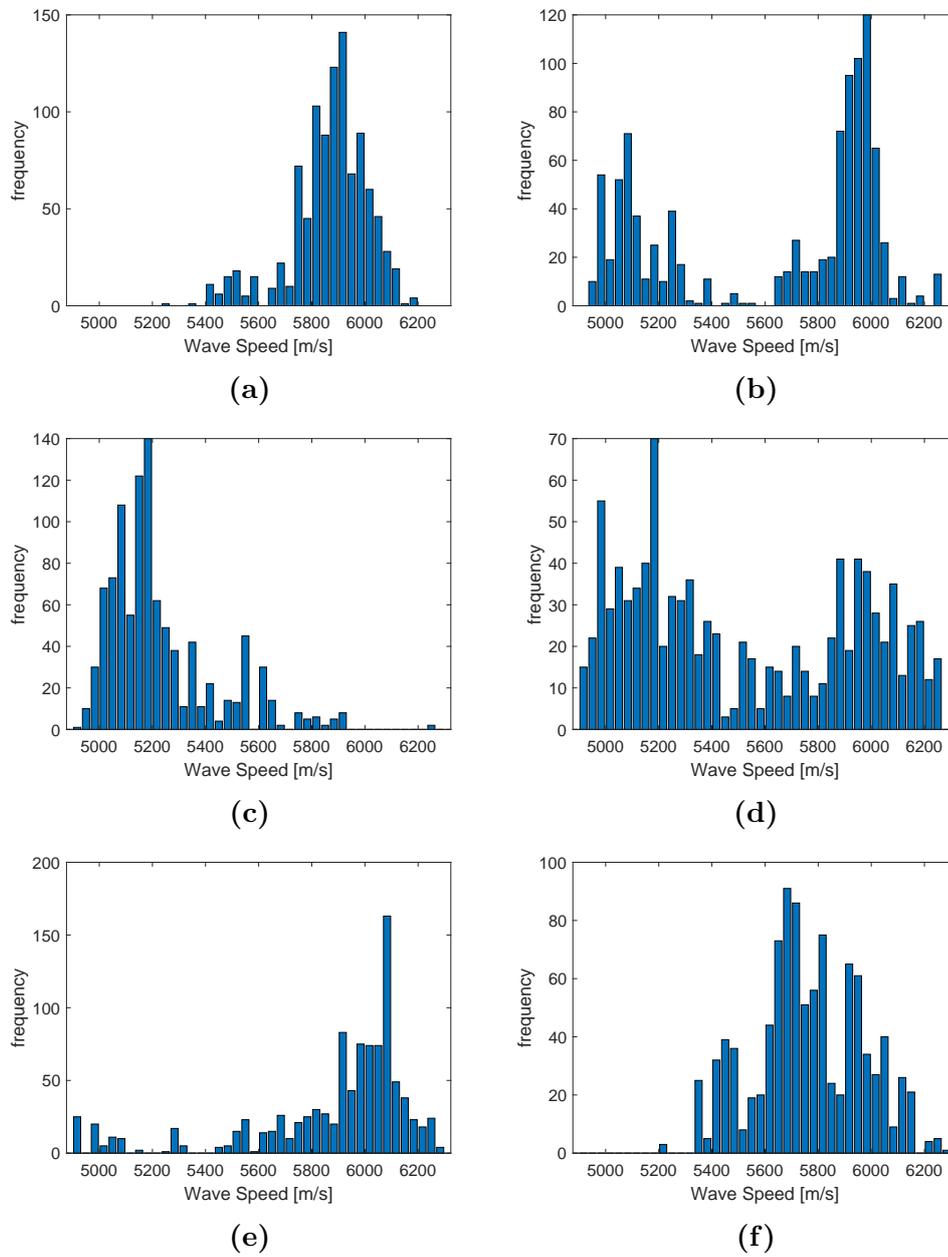


Figure 7.20: Wave speed distributions for selected points in the domain (as shown in Figure 7.19).

Table 7.2. From the ROC curves (see Figure 7.22) it can be observed that the TFM+ with the MAP, with an AUC=0.86197, is better than the standard TFM but not as good as the TFM+ with the known map, which is to be expected.

These results can be compared with the results in Subsection 5.5.1 which used the Pearson Correlation Coefficient objective function based on the log of the A-scans. It can be observed that there is a greater flaw location error here, however, the SNR and AUC are not as good. These comparisons are summarised in Table 7.3

Table 7.2: Comparison of quality metrics SNR, Flaw Error location, and ROC AUC for the lower flaw in Figure 5.2, Figure 7.21, and Figure 5.3.

	TFM	TFM+FWI	TFM+
SNR	17.5	22.3	23.6
Flaw Error (mm)	1.315	0.200	1.300
AUC	0.830	0.862	0.901

Table 7.3: Comparison of the quality metrics SNR, Flaw Error location, and ROC AUC, between the log PCC and the Hilbert transform time windowed (HTTW) objective functions.

	Log PCC	HTTW
SNR	28.4	22.3
Flaw Error (mm)	1.315	0.2
AUC	0.871	0.862

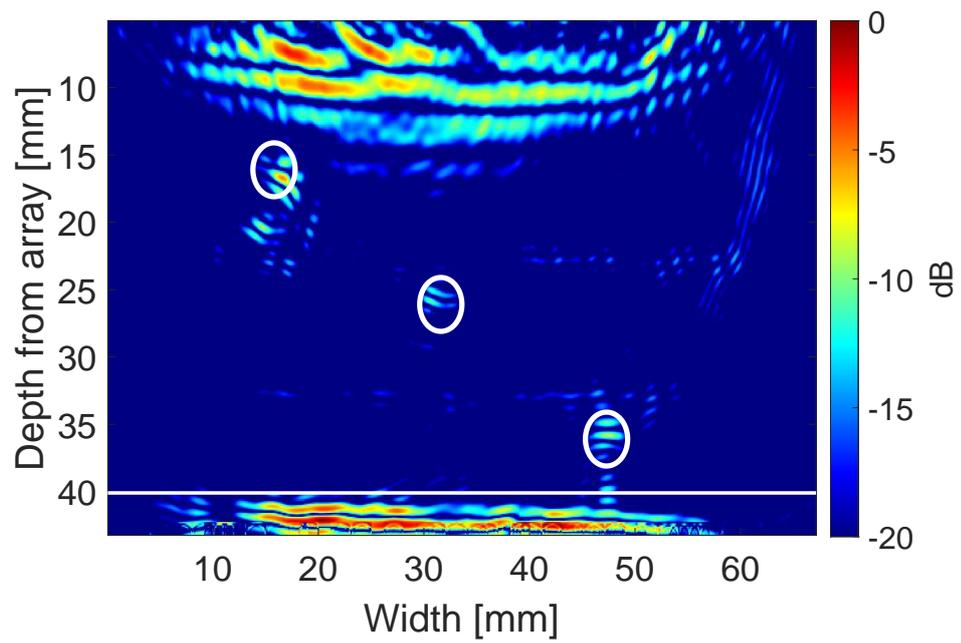


Figure 7.21: A TFM+ image using the MAP of the posterior distribution map, Figure 7.14(c). The white circles indicate the locations of the inclusions.

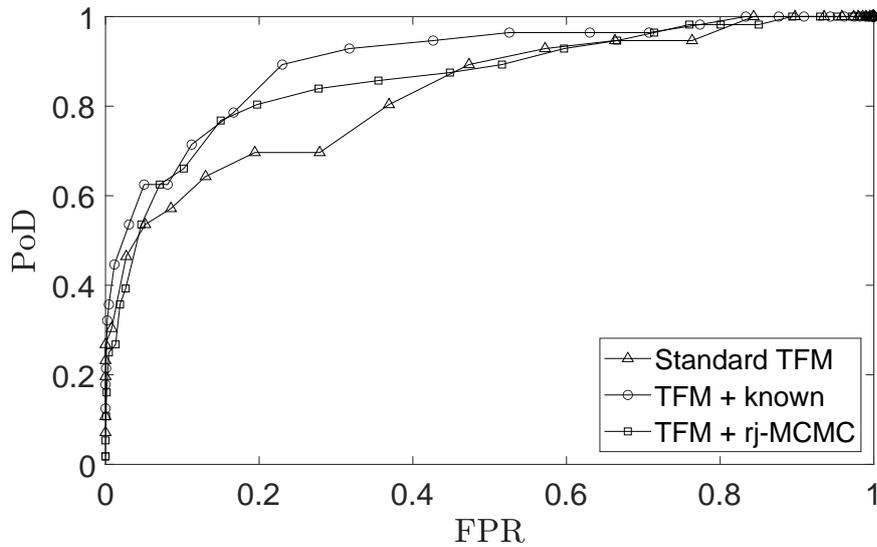


Figure 7.22: ROC Curves of the TFM (homogeneous material map), TFM+ (known heterogeneous material map) and TFM+ FWI (tomographic reconstruction of heterogeneous map shown in (Figure 7.21) using the binary map (Figure 5.5) as ground truth.

7.3.3 Reconstruction over 600,000 samples

To examine the convergence of the tomographic algorithm, the rj-MCMC was run for a further of 300,000 iterations (giving a total of 600,000 iterations) with a burn-in of 400,000. The remaining models were sampled at an interval of $\kappa = 100$. The initial model was the homogeneous model as used before. The prior and proposal distribution parameters were the same as before and can be found in Table 7.1. The simulation was seeded such that the first 300,000 iterations are the same as in Subsection 7.3.1.

The average objective function value, as the algorithm iterates through the samples is shown in Figure 7.23, where it can be observed that at around 400,000th iteration the objective function value is more or less level with less variation. However, this is most likely due to the higher number of cells (shown in Figure 7.24). With a higher number of cells each perturbation will have a lesser effect on the objective function.

From the median of the resulting posterior distribution, Figure 7.25, it can be observed that there is a large region of varying wave speed at 18mm in depth followed by 4 distinct layers of approximately 5mm each followed another layer of varying wave speed. This material wave speed map is subjectively worse than in Figure 7.14.

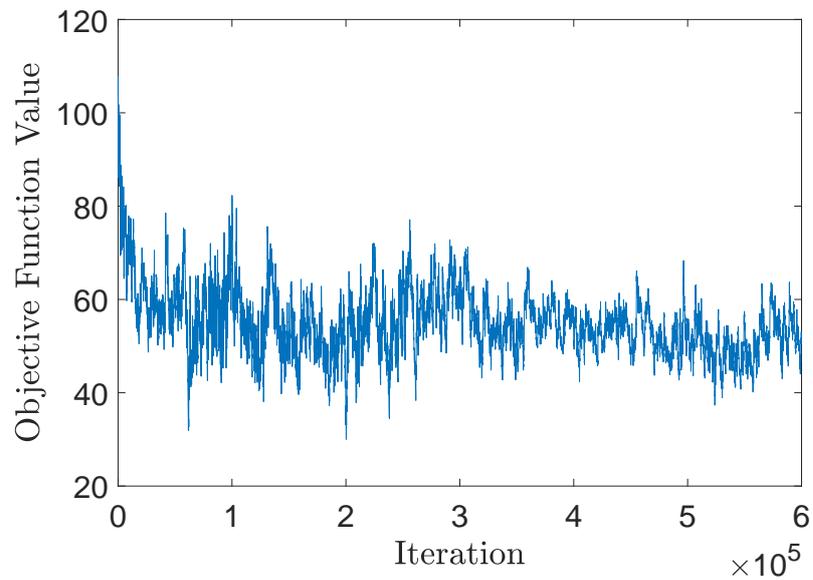


Figure 7.23: The average HTTW objective function value as the algorithm iterates through each Voronoi sample.

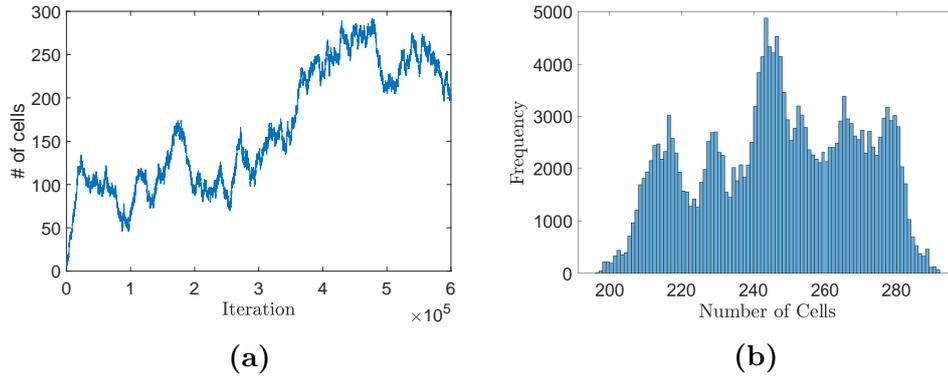


Figure 7.24: Plot (a) is the number of cells at each iteration of the rj-MCMC algorithm and plot (b) is the distribution of the number of cells from the last 200,000 samples.

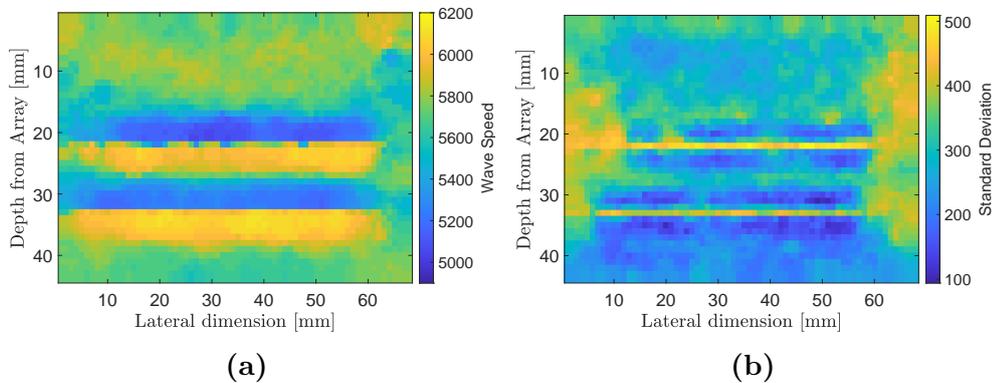


Figure 7.25: Material map reconstructions of the 4 layered medium. Plot (a) is the median of the posterior distribution on wave speed at each point in the domain whilst plot (d) shows the standard deviation respectively.

Looking at the acceptance rates of the parameters, Figure 7.26, it can be observed that they all remain steady after the first 100,000 runs. This behaviour is dictated by the Metropolis-Hastings criteria (Equation (2.11)). If the combination of these terms remains constant from model M to model M' then the probability of acceptance will remain constant. For example, in Equation (2.14), which describes the prior ratio for a perturbation caused by a birth of a cell, when

the number of cells (L) gets large the relative difference becomes small. If the objective function doesn't vary much (as is suggested in Figure 7.23) then the likelihood in Equation (2.16) will be a constant.

There are several possible ways to improve this reconstruction. One is to improve the objective function. Another is to improve the forward model used; since the semi-analytical model is used to generate the observed data it could be that it is not capturing enough information of the layered medium to allow for an accurate reconstruction. As stated before, there is an imbalance between the reflections each layer can produce. With this in mind, a material reconstruction based on finite element simulated data is attempted next.

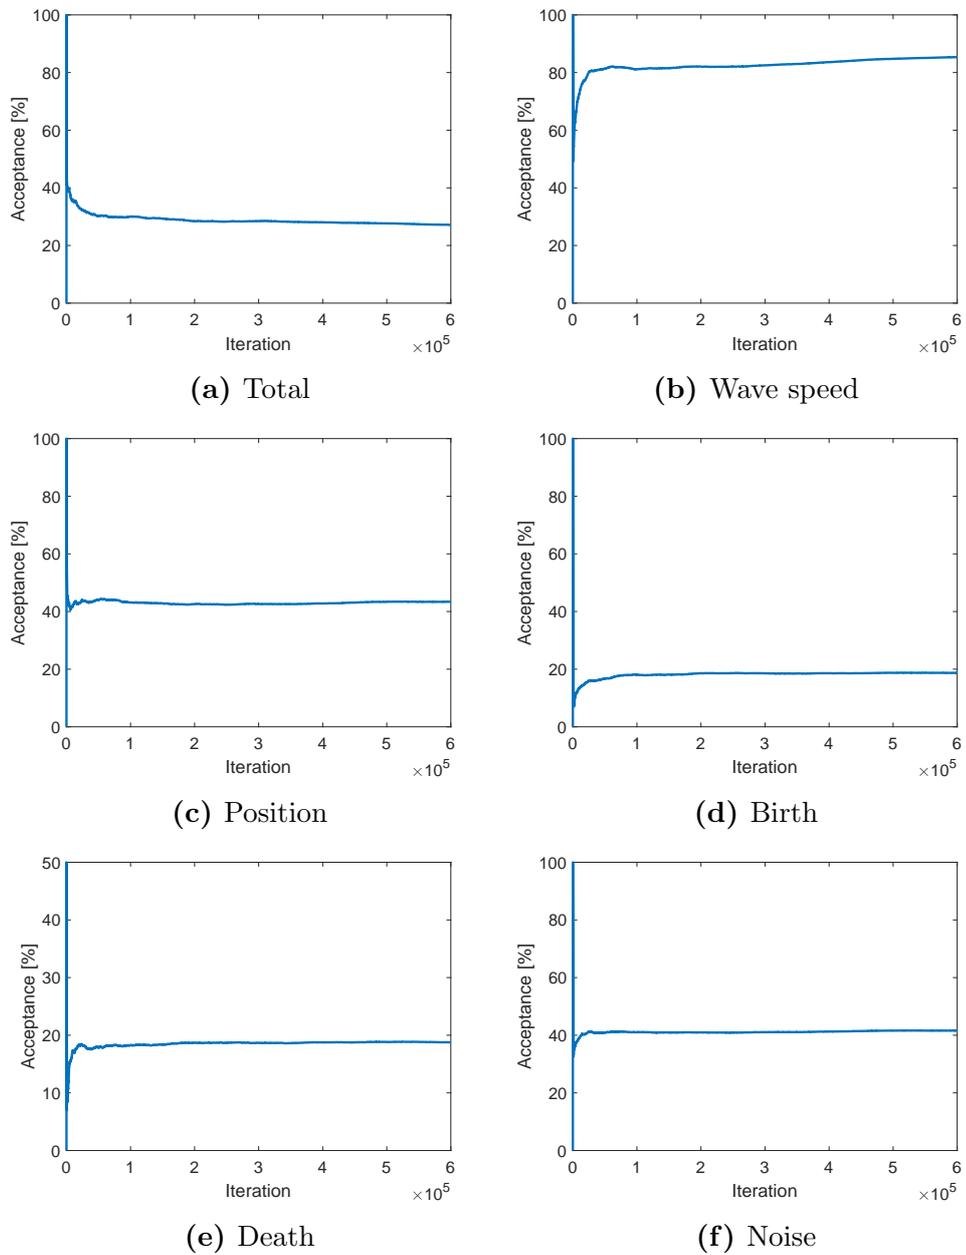


Figure 7.26: The acceptance rates of the different parameters as the algorithm iterates through 600,000 samples.

7.4 Tomographic Reconstruction of a Layered Medium using the Hilbert Transform Time Windowed Objective Function on Finite Element Simulated Data

In the previous section, the tomographic reconstruction was based on observed data generated from the semi-analytical model. In this section, it will be run using finite element simulated data generated from OnScale. The goal is to be able to reconstruct a material wave speed map that contains more layered features than Figure 6.14 using the Hilbert transform time windowed objective function.

The rj-MCMC was run for a total of 600,000 samples, with a burn-in of 400,000, and the remaining samples were sampled at an interval $\kappa = 100$. A homogeneous initial model was used with each Voronoi cell assigned a wave speed of 5600 m/s. The algorithm had the following limits on the prior distribution: the number of cells must lie in the range of 5 to 300 cells; the wave speed for a Voronoi cell must lie in the range 5600 ± 700 m/s; and the noise parameter must lie in the range of 1 and 120. The standard deviation for the proposal distributions (σ) were: on *wave speed* 400 m/s and 120 m/s for the delayed rejection; on *cell position* 15% and 3% for the delayed rejection; on the wave speed for a *birth* perturbation 50 m/s; and on the *noise* 6. These are summarised in Table 7.4.

Table 7.4: Proposal distributions set for the rj-MCMC on simulated data using the new objective function.

Perturbations	Range	σ	σ^{dr}
Wave speed	5600 ± 700 m/s	400 m/s	1200m/s
Seed position	–	15%	3%
Birth	–	50 m/s	–
Noise	1 - 120	10	–

With these proposals, the algorithm had an overall acceptance rate of 39.7%. The acceptance rate for each parameter was: 77.1% for the wave speed; 48.0% for the position; 17.0% for the birth and 17.0% death of a cell; and 26.9% for the noise.

The average objective function value, as the algorithm iterates through the samples, is shown in Figure 7.27, which converges to around 98. The number of Voronoi cells for the last 200,000 iterations is shown in Figure 7.28, which shows a bimodal distribution with a peak around 90 cells and the other around 120 cells. This indicates that the model has not converged and the algorithm is still exploring the model space.

From the moments of the posterior distribution on the wave speed at each point in the domain, material maps were constructed. Figure 7.29 shows, (a) the mean, (b) the median, (c) the MAP, and (d) the standard deviation of the posterior probability distributions. One distinct layer at the bottom of the domain can be observed with a depth approximately of 10 mm and a wave speed of 6000-6200 m/s. A faint second layer and a heterogeneous region above this with a high standard deviation can be observed. This could be in part due to the

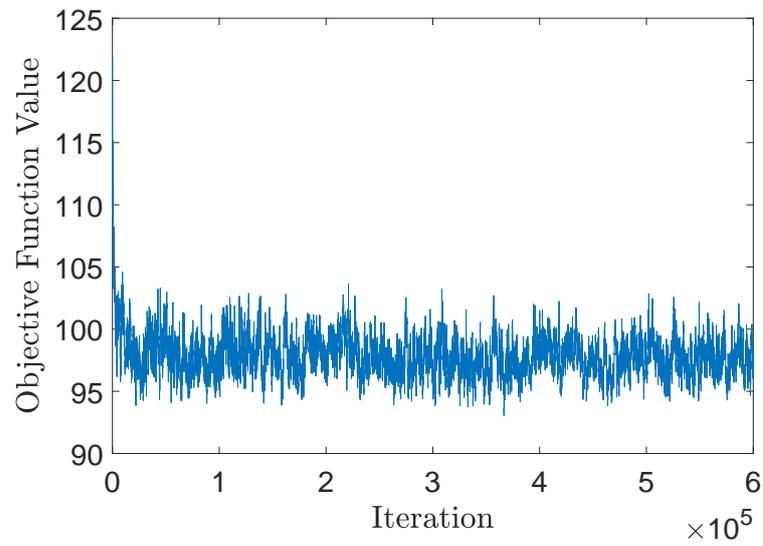


Figure 7.27: The average objective function value as the algorithm iterates through each Voronoi sample for the rj-MCMC where the observed data is generated in a FE package.

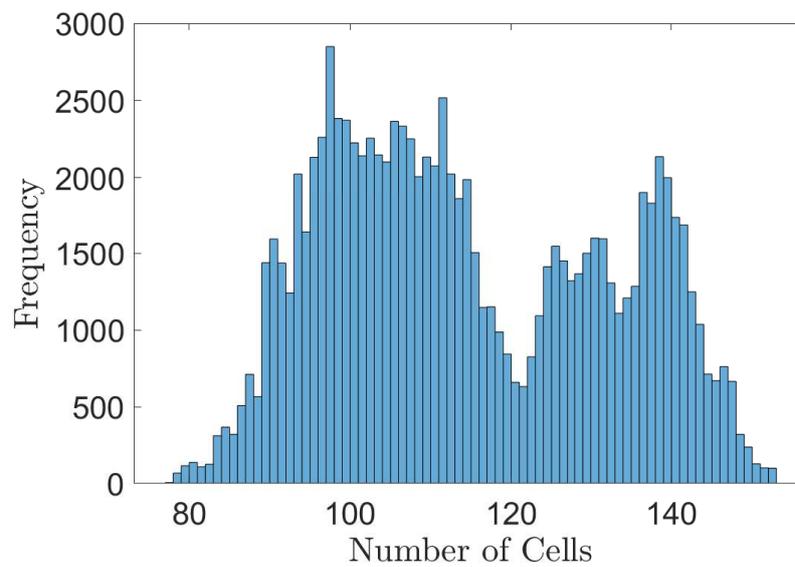


Figure 7.28: The distribution of the number of cells from the last 200,000 samples.

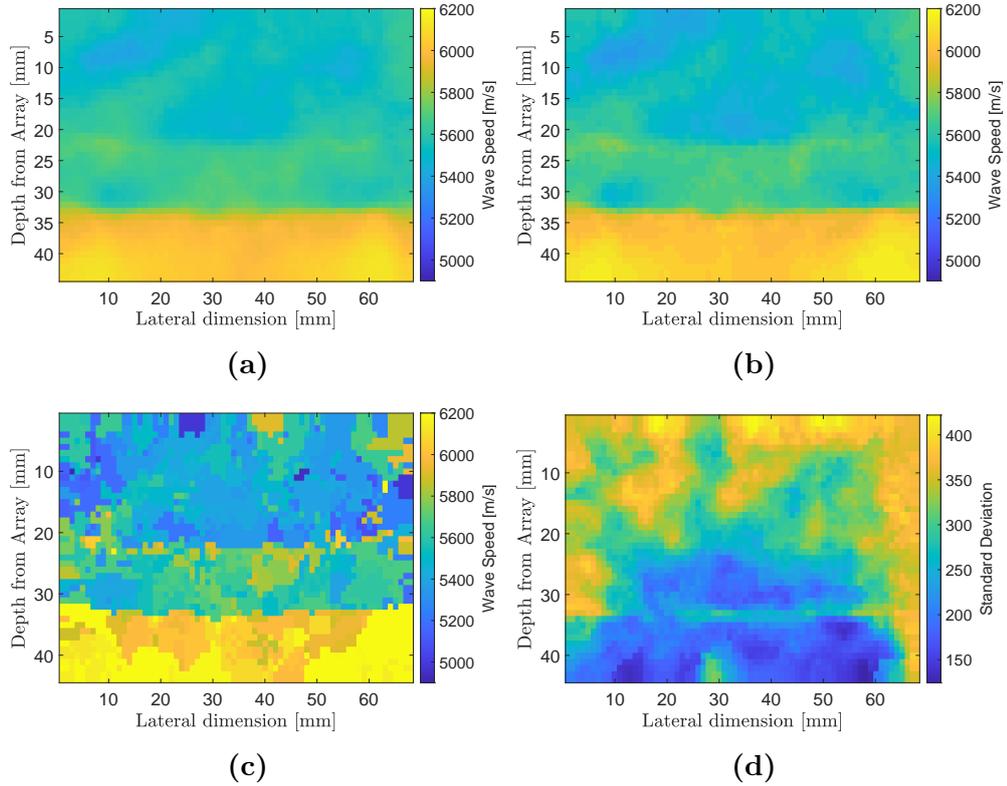


Figure 7.29: Material map reconstructions of the 4 layered medium for the simulation using the full 1024 rays. Plot (a) is the mean of the posterior distribution on wave speed at each point in the domain whilst plots (b)-(d) show the median, MAP, and the standard deviation respectively.

semi-analytical model’s potential to contain more internal reflections from the bottom layer as discussed in Section 5.5 and Section 7.3. Another reason, is the difference in arrival times for the subsequent wave packets in the A-scans that have a diagonal raypath. This is shown in Figure 7.30 and Figure 7.31, where it can be observed that the second wave packet in the FE A-scan arrives sooner than predicted by the semi-analytical model. The reason for this is shown in Figure 6.6, which shows two alternative raypaths, FE simulation and semi-analytical model,

for a diagonal ray with a reflection in the last layer. The semi-analytical model (solid line) is one dimensional, however, the FE simulation is two dimensional and therefore has a raypath which has a more acute angle to the vertical and undergoes refraction. This allows for a shorter distance and therefore a faster time of arrival.

7.4.1 TFM+ from the Reconstructed Material Maps from Simulated data

To quantify the quality of the reconstruction with respect to enhanced flaw imaging, the median of the posterior distribution was used to correct the delay laws in the TFM+ algorithm and the resulting flaw reconstructions are displayed in

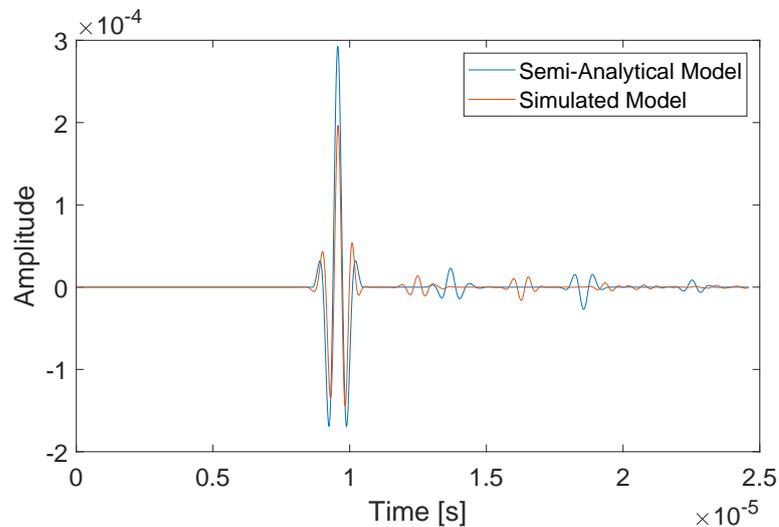


Figure 7.30: A-scans from a layered medium generated by the semi-analytical model (blue line) and the equivalent finite element simulation (orange line) for a diagonal raypath (transmitting on element 1 and receiving on element 15).

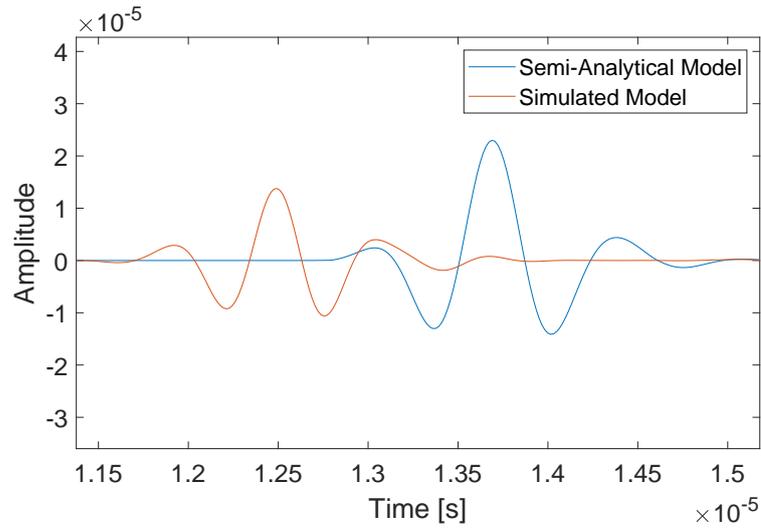


Figure 7.31: The second reflection from the A-scans (Figure 7.30) from a layered medium generated by the semi-analytical model (blue line) and the equivalent finite element simulation (orange line) for a diagonal raypath (transmitting on element 1 and receiving on element 15).

Figure 7.32. It can be observed that there is a flaw in the bottom layer contained within the white circle indicating the known location of the inclusion. This flaw has an SNR value 14.7 dB and a flaw location error of 1.005 mm. From the ROC curves (Figure 7.33) it can be observed that the performance is marginally better than the standard TFM with an ROC AUC value of 0.855.

Table 7.5: Comparison of quality metric SNR, Flaw Error and AUC between the standard TFM, TFM+ FWI, TFM+ known, for the simulation using simulated data from a FE package.

	TFM	TFM+FWI	TFM+
SNR	17.5	14.7	23.6
Flaw Error [mm]	1.315	1.005	1.300
AUC	0.8304	0.8550	0.9012

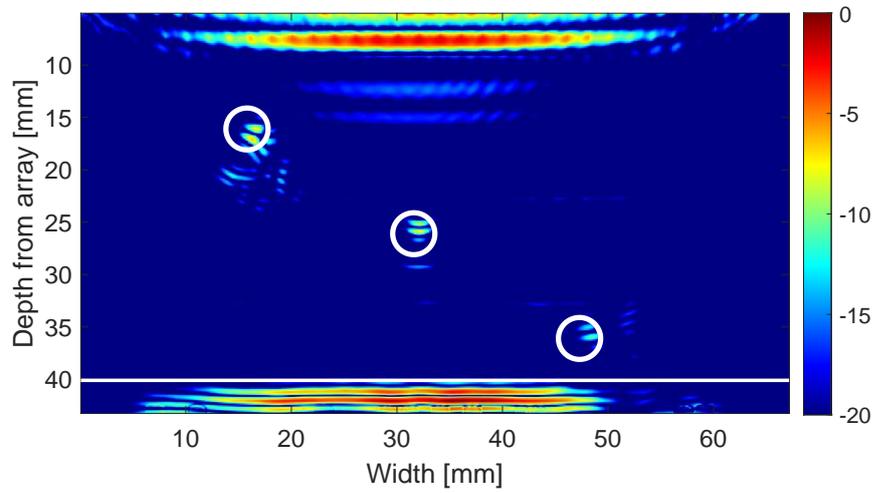


Figure 7.32: A TFM+ image from the material map given by the MAP of the posterior distribution as shown in Figure 7.29(c)

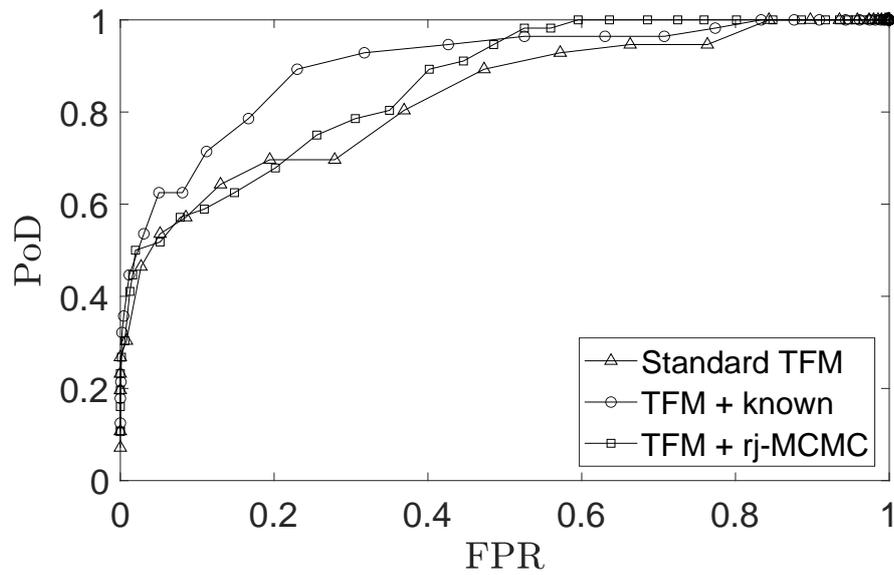


Figure 7.33: ROC curves from Figure 5.6 including the ROC curve from the TFM+ generated from the median of the posterior distribution from Figure 7.29(c).

Comparing these values to the SNR, flaw location errors, and AUC values from the case in Chapter 6 which used the PCC on the log of the A-scans as the objective function; it could be argued Figure 7.29 provides a better material reconstruction due to indications of layers. However, when looking at the SNR values, flaw location error, and AUC values the material reconstruction Chapter 6 performs better in all three metrics. The values can be found in Table 7.6.

7.5 Discussion

In this chapter, a new objective function was developed that calculated the envelope of a wave using a Hilbert transform, windowed regions of interest and then calculated the L_1 -norm between these windows. This Hilbert transform time windowed (HTTW) objective function was then implemented into the tomographic inversion scheme. The first section used observed data generated by the forward model (so called inverse crime), where the layer interfaces in the material reconstruction were not as defined compared to these generated using the PCC log based objective function. However, the wave speeds of the layers were more ac-

Table 7.6: Comparison of quality metric SNR, Flaw Error and AUC between the standard TFM+ images arising from using PCC on the log of A-scans (Log PCC) and l1-norm on the windowed envelopes of A-scans (l1 Windows) as objective functions from simulated data from a FE package.

	Log PCC	HTTW
SNR	15.9	14.7
Flaw Error [mm]	0.1	1.005
AUC	0.875	0.855

curate. The next potential research would be to look into the robustness of this HTTPW objective function via different levels of proportional noise and additive noise similar to Section 5.6.

When the material reconstruction was used with the TFM+ algorithm there was an improvement in the imaging of the lowest (most difficult) flaw. However, it did not perform any better than the imaged flaw from Section 5.5, which had a better flaw location error and AUC value.

The HTTPW objective function was then used to reconstruct the layered medium from finite element simulated data. This provided a reconstruction which showed indication of layers but had not fully converged. This material reconstruction was visually an improvement from Chapter 6, however, the quality metrics arising indicated, from the subsequent application of the TFM, that it performed worse.

Chapter 8

Conclusion and Future Work

This thesis has developed a framework for using ultrasonic array data to reconstruct material maps and to use these to improve the imaging of flaws.

8.1 Conclusion

This thesis began by giving a review of the current literature on Non-destructive testing (NDT), introducing the different methods of flaw detection and characterisation, and tomographic reconstruction. There was an introduction to inverse problems and different tools for inversion such as optimisation techniques, Bayesian inverse problems, and full waveform inversion. Following this, Chapter 2 gave a detailed summary of the methods and tools used to create the full waveform inversion (FWI) presented in this thesis. This included a discussion on Voronoi tessellations, the different types of observed data that are used in the to-

mographic inversion, and the underlying principles of the Bayesian optimisation method used in the *rj*-MCMC method. A mathematical model is then developed to describe the propagation of ultrasonic waves travelling through layered media in one-dimension. Using this, Chapter 3 develops, a bespoke one-dimensional semi-analytical model is developed for the forward solver in the inversion algorithm that follows.

In Chapter 4, the first approach to the tomographic inversion was conducted on a synthetically generated random material. The material wave speed reconstruction was used within the TFM+ imaging algorithm to reconstruct a flaw present within the medium. Metrics were introduced to quantify the relative improvement in the flaw reconstruction when compared to the benchmark TFM imaging algorithm (an industry standard imaging algorithm). This approach was then compared to that obtained using a time of flight (ToF) tomographic inversion method. It was shown that FWI had comparable results to the ToF through transmission (TT) results with a flaw location error of 1.2mm and an area under the curve (AUC) of 0.7 for the ROC curves generated.

Chapter 5 looked at the case of a layered medium with the observed data being generated from the semi-analytical forward model. The objective function was adapted by taking the logarithm of the A-scans before calculating the objective function value. With this new objective function, and using only a subset of the data available, a tomographic reconstruction was obtained that contained layered features. These layers had an incorrect thickness of 6mm for one and 14mm for another. It gave an AUC value of 0.84, which was comparable to the

standard TFM. Using the full data set the tomographic reconstruction contained clearly defined layers of around 10mm each, however, the wave speeds of the layers were not correct for all the layers with one layer being off by 700 m/s. The TFM+ image from the tomographic reconstruction gave an AUC value of 0.87, an improvement on the standard TFM. It was then shown that this FWI method was robust to proportional noise of 20% but not additive noise.

In Chapter 6 the FWI tomographic methodology was challenged with data that was generated from a finite element simulation. This brought in additional physics and a greater level of complexity. The FWI method was unable to reconstruct the layered medium. However, this reconstruction still gave an AUC value of 0.875 which is an improvement on the standard TFM.

Next, in Chapter 7, a new objective function was developed. One that took the envelope of the wave, windowed regions of interest, and then calculated the L_1 -norm between the forward model and the observed data. This Hilbert transform time windowed objective function was not dominated by the initial wave and was robust to additive noise. Tomographic reconstructions with the observed data generated from the semi-analytical model and finite element simulated data were then attempted. With the semi-analytical data the tomographic reconstruction managed to have regions of the correct wave speed but did not have clearly defined layered interfaces. The TFM+ image gave an AUC value of 0.86 which is an improvement on the standard TFM. With the finite element simulation one layer was successfully reconstructed, with an AUC value of 0.855.

8.2 Future Work

There is a large scope of future work for this research. There are various methods that could improve the tomographic reconstruction of the layered media. However, the goal of this thesis was to create an inversion method of a general isotropic heterogeneous material. The next step could be to focus on the inversion of more complex heterogeneous material, such as a medium with a grid layout. Another would be to adapt the full waveform inversion for locally anisotropic materials.

The next step could be to undertake tomographic inversion from experimental data from an ultrasonic phased array. To achieve this, improvements to the forward model and/or objective function could be implemented. There are various methods of improving the forward model used. The first would be a one-dimensional forward model that generates information over the whole medium uniformly thus allowing for accurate material reconstruction close to the transmitting array. Another would be to develop a two-dimensional model. With a two-dimensional model it could be possible for additional physics to be included such a refraction, shear waves, and mode conversion. An alternative would be using a finite element solver as the forward model however this would require more computational power. Another approach would be to develop the forward model from the elastic wave equations instead of the acoustic wave equations. A new objective function could also be developed, one option is the Wasserstein metric.

The inversion algorithm could be adapted and one method that could be implemented is parallel tempering as presented in [101]. Another would be to add in elements of simulated annealing to the evolution of the noise parameter in the rj-MCMC algorithm to improve convergence.

This thesis has laid out a framework for tomographic reconstruction of heterogeneous materials using limited aperture ultrasonic array data. This platform will hopefully pave the way for a fully functioning tomography capability in the years to come.

References

- [1] R. Halmshaw, *Non-destructive testing*. London: Edward Arnold, 2nd ed., 1991.
- [2] G. Honeymoon, *Non-destructive testing*. London: Institute of Metals, 1989.
- [3] K. G. Boving, *NDE Handbook: Non-destructive examination methods for condition monitoring*. London ; Boston: Butterworths, 1989.
- [4] T. W. Liao and J. Ni, “An automated radiographic NDT system for weld inspection: part I weld extraction,” *NDT&E International*, vol. 29, pp. 157–162, 1996.
- [5] D. Popescu, F. D. Anania, C. E. Cotet, and C. G. Armza, “Fully-automated liquid penetrant inspection line simulation model for increasing productivity,” *International Journal of Simulation Modelling*, vol. 12, pp. 82–93, 2013.

REFERENCES

- [6] J. Garcia-Martn, J. Gomez-Gil, and E. Vazquez-Sanchez, “Non-destructive techniques based on eddy current testing,” *Sensors*, vol. 11, no. 3, pp. 2525–2565, 2011.
- [7] B. W. Drinkwater and P. D. Wilcox, “Ultrasonic arrays for non-destructive evaluation: A review,” *NDT & E International*, vol. 39, no. 7, pp. 525 – 541, 2006.
- [8] Y.-H. Chung and N. Moori, “An overview of a non-medical x-ray safety program at a large university,” *Health Physics*, vol. 108, pp. S37–S41, 2015.
- [9] C. Holmes, B. W. Drinkwater, and P. D. Wilcox, “Post-processing of the full matrix of ultrasonic transmit/receive array data for non-destructive evaluation,” *NDT & E International*, vol. 38, no. 8, pp. 701 – 711, 2005.
- [10] S. Kočiš and Z. Figura, *Ultrasonic Measurements and Technologies. Sensor physics and technology*; 4, London: Chapman & Hall, 1996.
- [11] H. Kuttruff, *Ultrasonics: Fundamentals and Applications*. London ; New York : New York, NY, USA: Elsevier Applied Science, 1991.
- [12] K. Mizutani, K. Nishizaki, K. Nagai, and K. Harakawa, “Measurement of temperature distribution in space using ultrasound computerized tomography,” *Japanese Journal of Applied Physics*, vol. 36, no. Part 1, No. 5B, pp. 3176–3177, 1997.

REFERENCES

- [13] T. Mizoguchi, K. Tamura, J. Mamou, J. A. Ketterling, K. Yoshida, and T. Yamaguchi, “Comprehensive backscattering characteristics analysis for quantitative ultrasound with an annular array: a basic study on homogeneous scattering phantom,” *Japanese Journal of Applied Physics*, vol. 58, p. SGGE08, 2019.
- [14] N. B. Smith and A. Webb, *Introduction to Medical Imaging: Physics, Engineering and Clinical Applications*. Cambridge: Cambridge University Press, 2010.
- [15] N. Ozmen, R. Dapp, M. Zapf, H. Gemmeke, N. V. Ruitter, and K. W. A. van Dongen, “Comparing different ultrasound imaging methods for breast cancer detection,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 62, no. 4, pp. 637–646, 2015.
- [16] L. A. Miller and A. Surlykke, “How some insects detect and avoid being eaten by bats: Tactics and countertactics of prey and predator: Evolutionarily speaking, insects have responded to selective pressure from bats with new evasive mechanisms, and these very responses in turn put pressure on bats to improve their tactics,” *BioScience*, vol. 51, no. 7, pp. 570–581, 2001.

REFERENCES

- [17] M. Cosentino, F. Guarato, J. Tougaard, D. Nairn, J. C. Jackson, and J. F. C. Windmill, “Porpoise click classifier (PorCC): A high-accuracy classifier to study harbour porpoises (*Phocoena phocoena*) in the wild,” *The Journal of the Acoustical Society of America*, vol. 145, no. 6, pp. 3427–3434, 2019.
- [18] D. Voss, “March 1880: The Curie brothers discover piezoelectricity,” *APS News*, vol. 23, no. 3, 2014.
- [19] R. T. Beyer, *Sounds of our times: 200 years of acoustics*. New York: Springer-Verlag, 1999.
- [20] M. Nicolson and J. E. E. Fleming, *Imaging and Imagining the Fetus*. Baltimore, Maryland: The John Hopkins University Press, 2013.
- [21] W. J. McGonnagle, *Non-Destructive Testing*. New York: McGraw-Hill, 1961.
- [22] S. Chatillon, G. Cattiaux, M. Serre, and O. Roy, “Ultrasonic non-destructive testing of pieces of complex geometry with a flexible phased array transducer,” *Ultrasonics*, vol. 38, no. 1, pp. 131 – 134, 2000.
- [23] C. Nageswaran, C. Carpentier, and Y. Y. Tse, “Microstructural quantification, modelling and array ultrasonics to improve the inspection of austenitic welds,” *Insight*, vol. 51, no. 12, pp. 660 – 666, 2009.

REFERENCES

- [24] C. Li, D. Pain, P. D. Wilcox, and B. W. Drinkwater, “Imaging composite material using ultrasonic arrays,” *NDT & E International*, vol. 53, pp. 8 – 17, 2013.
- [25] P. R. L. Alves, J. S. C. Filho, and M. G. Farinhas, “A new approach to determine tensile stress states from the parameters of longitudinal waves,” *Applied Mathematical Modelling*, vol. 83, pp. 189 – 201, 2020.
- [26] C. A. Guertler, R. J. Okamoto, J. A. Ireland, C. P. Pacia, J. R. Garbow, H. Chen, and P. V. Bayly, “Estimation of anisotropic material properties of soft tissue by mri of ultrasound-induced shear waves,” *Journal of Biomechanical Engineering*, vol. 142, no. 3, pp. 0310011–03100117, 2020.
- [27] L. Xiang, D. Greenshields, S. Dixon, and R. S. Edwards, “Phased electromagnetic acoustic transducer array for Rayleigh wave surface defect detection,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 67, pp. 1403–1411, 2020.
- [28] Z. Tang, N. Munir, T. Lee, Y. Yeom, and S. Song, “Lamb wave flaw classification in al plates using time reversal and deep neural networks,” *Journal of the Korean Physical Society*, vol. 75, pp. 978 – 984, 2019.
- [29] A. Mandi, S. Kundu, P. Chandra Pal, and P. Pati, “An analytic study on the dispersion of love wave propagation in double layers lying over inhomogeneous half-space,” *Journal of Solid Mechanics*, vol. 11, no. 3, pp. 570–580, 2019.

REFERENCES

- [30] T. Stratoudaki, M. Clark, and P. D. Wilcox, “Laser induced ultrasonic phased array using full matrix capture data acquisition and total focusing method,” vol. 24, no. 19, pp. 21921–21938, 2016.
- [31] D. Pieris, T. Stratoudaki, Y. Javadi, P. Lukacs, S. Catchpole-Smith, P. D. Wilcox, A. Clare, and M. Clark, “Laser induced phased arrays (LIPA) to detect nested features in additively manufactured components,” *Materials & Design*, vol. 187, p. 108412, 2020.
- [32] C. B. Scruby and L. E. Drain, *Laser ultrasonics : techniques and applications*. Bristol, England ; Philadelphia: A. Hilger, 1990.
- [33] G. Kaiser, *A Friendly Guide to Wavelets*. Boston: Birkhäus Boston, 1994.
- [34] H. Ozaktas, D. Mendlovic, M. Kutay, and Z. Zalevsky, *The Fractional Fourier Transform: With Applications in Optics and Signal Processing*. Wiley Series in Pure and Applied Optics, Chichester ; New York: Wiley, 2001.
- [35] “Data presentation.” <https://www.nde-ed.org/EducationResources/CommunityCollege/Ultrasonics/EquipmentTrans/DataPres.htm>, (accessed 19/03/20). NDT Resource Center.
- [36] M. Klemm, I. J. Craddock, J. A. Leendertz, A. Preece, and R. Benjamin, “Improved delay-and-sum beamforming algorithm for breast cancer detection,” *International Journal of Antennas and Propagation*, pp. 1–9, 2008.

REFERENCES

- [37] S. Li and S. Fomel, “Kirchhoff migration using eikonal-based computation of traveltimes source derivatives,” *Geophysics*, vol. 78, no. 4, pp. 211–219, 2013.
- [38] C. Fan, M. Caleap, M. Pan, and B. W. Drinkwater, “A comparison between ultrasonic array beamforming and super resolution imaging algorithms for non-destructive evaluation,” *Ultrasonics*, vol. 54, no. 7, pp. 1842–1850, 2014.
- [39] M. Jobst and G. D. Connolly, “Demonstration of the application of the total focusing method to the inspection of steel welds,” in *10th European Conference on Non-Destructive Testing*, 2010.
- [40] C. Wang, J. Mao, T. Ieng, Z. Y. Zhuang, and X. M. Wang, “Efficient acceleration for total focus method based on advanced parallel computing in FPGA,” *International Journal of Acoustics and Vibration*, vol. 22, no. 4, pp. 536–540, 2017.
- [41] K. M. M. Tant, E. Galetti, A. J. Mulholland, A. Curtis, and A. Gachagan, “A transdimensional bayesian approach to ultrasonic travel-time tomography for non-destructive testing,” *Inverse Problems*, vol. 34, no. 9, p. 095002, 2018.
- [42] M. V. Felice, A. Velichko, and P. D. Wilcox, “Accurate depth measurement of small surface-breaking cracks using an ultrasonic array post-processing technique,” *NDT & E International*, vol. 68, pp. 105 – 112, 2014.

REFERENCES

- [43] C. Zhang, P. Huthwaite, and M. Lowe, “Eliminating backwall effects in the phased array imaging of near backwall defects,” *The Journal of the Acoustical Society of America*, vol. 144, no. 2, pp. 1075–1088, 2018.
- [44] K. Mayer, R. Marklein, K. Langenberg, and T. Kreutter, “Three-dimensional imaging system based on fourier transform synthetic aperture focusing technique,” *Ultrasonics*, vol. 28, no. 4, pp. 241 – 255, 1990.
- [45] L. J. Busse, “Three-dimensional imaging using a frequency-domain synthetic aperture focusing technique,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 39, no. 2, pp. 174–179, 1992.
- [46] H. Jin, E. Wu, Y. Han, K. Yang, and J. Chen, “Frequency domain synthetic aperture focusing technique for variable-diameter cylindrical components,” *The Journal of the Acoustical Society of America*, vol. 142, no. 3, pp. 1554–1562, 2017.
- [47] T. Olofsson and T. Stepinski, “Phase shift migration for imaging layered materials and objects immersed in water,” in *2009 IEEE International Ultrasonics Symposium*, pp. 673–676, 2009.
- [48] J. Jensen, M. B. Stuart, and J. A. Jensen, “Increased frame rate for plane wave imaging without loss of image quality,” in *2015 IEEE International Ultrasonics Symposium (IUS)*, pp. 1–4, 2015.

REFERENCES

- [49] L. Sandrin, M. Tanter, S. Catheline, and M. Fink, “Shear modulus imaging with 2-D transient elastography,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 49, no. 4, pp. 426–435, 2002.
- [50] O. Couture, M. Fink, and M. Tanter, “Ultrasound contrast plane wave imaging,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 59, no. 12, pp. 2676–2683, 2012.
- [51] J. Jensen, M. B. Stuart, and J. A. Jensen, “Optimized plane wave imaging for fast and high-quality ultrasound imaging,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 63, no. 11, pp. 1922–1934, 2016.
- [52] M. Fink, D. Cassereau, A. Derode, C. Prada, P. Roux, M. Tanter, J.-L. Thomas, and F. Wu, “Time-reversed acoustics,” *Reports on Progress in Physics*, vol. 63, no. 12, pp. 1933–1995, 2000.
- [53] M. Fink and C. Prada, “Acoustic time-reversal mirrors,” *Inverse Problems*, vol. 17, no. 1, pp. 1–38, 2001.
- [54] F. K. Gruber, E. A. Marengo, and A. J. Devaney, “Time-reversal imaging with multiple signal classification considering multiple scattering between the targets,” *The Journal of the Acoustical Society of America*, vol. 115, no. 6, pp. 3042–3047, 2004.

REFERENCES

- [55] Y. Guo, P. Monk, D. Colton, and D. I. Hastie, “The linear sampling method for sparse small aperture data,” *Applicable Analysis*, vol. 95, pp. 1599 – 1615, 2015.
- [56] R. Potthast, “A study on orthogonality sampling,” *Inverse Problems*, vol. 26, no. 7, p. 074015, 2010.
- [57] K. Ito, B. Jin, and J. Zou, “A direct sampling method to an inverse medium scattering problem,” *Inverse Problems*, vol. 28, no. 2, p. 025003, 2012.
- [58] F. J. Humphreys, “Characterisation of fine-scale microstructures by electron backscatter diffraction (EBSD),” *Scripta Materialia*, vol. 51, no. 8, pp. 771 – 776, 2004.
- [59] C. Ma, C. Dong, D. Elsworth, Q. Wang, Z. Huang, H. Liu, G. Luan, B. Yu, D. Li, H. Yin, M. Feng, and Y. He, “Insights from electron backscatter diffraction into the origin of fibrous calcite veins in organic-rich shale from lower Es3 to upper Es4, Jiyang Depression, China,” *Marine and Petroleum Geology*, vol. 113, p. 104131, 2020.
- [60] T. Tanaka, N. Maruyama, N. Nakamura, and A. J. Wilkinson, “Tetragonality of Fe-C martensite a pattern matching electron backscatter diffraction analysis compared to X-ray diffraction,” *Acta Materialia*, vol. 195, pp. 728–738, 2020.

REFERENCES

- [61] A. Matsushita, N. Fujimori, Y. Tsuchida, N. Ohtani, D. Dojima, K. Koide, T. Kaneko, and S. Shikata, “Evaluation of diamond mosaic wafer crystallinity by electron backscatter diffraction,” *Diamond and Related Materials*, vol. 101, p. 107558, 2020.
- [62] X. Zhang, K. Wang, Q. Zhou, J. Kong, Y. Peng, J. Ding, C. Diao, D. Yang, Y. Huang, T. Zhang, and S. W. Williams, “Element partitioning and electron backscatter diffraction analysis from feeding wire to as-deposited microstructure of wire and arc additive manufacturing with super duplex stainless steel,” *Materials Science and Engineering*, vol. 773, p. 138856, 2020.
- [63] J. S. Weaver, V. Livescu, and N. A. Mara, “A comparison of adiabatic shear bands in wrought and additively manufactured 316L stainless steel using nanoindentation and electron backscatter diffraction,” *Journal of Materials Science*, vol. 55, pp. 1738–1752, 2020.
- [64] S. D. Sharples, M. Clark, and M. G. Somekh, “Spatially resolved acoustic spectroscopy for fast non-contact imaging of material microstructure,” *Optics Express*, vol. 14, pp. 10435–10440, 2006.
- [65] R. J. Smith, W. Li, J. Coulson, M. Clark, M. G. Somekh, and S. D. Sharples, “Spatially resolved acoustic spectroscopy for rapid imaging of material microstructure and grain orientation,” *Measurement Science and Technology*, vol. 25, no. 5, p. 055902, 2014.

REFERENCES

- [66] R. Patel, W. Li, R. J. Smith, S. D. Sharples, and M. Clark, "Orientation of macro-sized polysilicon grains on wafers using spatially resolved acoustic spectroscopy," *Scripta Materialia*, pp. 67–70, 2017.
- [67] R. Patel, M. Hirsch, P. Dryburgh, D. M. Pieris, S. Achamfuo-Yeboah, R. J. Smith, R. Light, S. D. Sharples, A. Clare, and M. Clark, "Imaging material texture of as-deposited selective laser melted parts using spatially resolved acoustic spectroscopy," *Applied Science*, p. 1991, 2018.
- [68] P. Dryburgh, R. Patel, D. M. Pieris, M. Hirsch, W. Li, S. D. Sharples, R. J. Smith, A. Clare, and M. Clark, "Spatially resolved acoustic spectroscopy for texture imaging in powder bed fusion nickel superalloys," *AIP Conference Proceedings*, p. 020004, 2019.
- [69] J. Moysan, A. Apfel, G. Corneloup, and B. Chassignole, "Modelling the grain orientation of austenitic stainless steel multipass welds to improve ultrasonic assessment of structural integrity," *International Journal of Pressure Vessels and Piping*, vol. 80, no. 2, pp. 77 – 85, 2003.
- [70] A. Mark, Z. Fan, F. Azough, M. Lowe, and P. Withers, "Investigation of the elastic/crystallographic anisotropy of welds for improved ultrasonic inspections," *Materials Characterization*, vol. 98, pp. 47 – 53, 2014.

REFERENCES

- [71] Z. Fan, A. F. Mark, M. J. S. Lowe, and P. J. Withers, “Nonintrusive estimation of anisotropic stiffness maps of heterogeneous steel welds for the improvement of ultrasonic array inspection,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 62, no. 8, pp. 1530–1543, 2015.
- [72] C. Gueudre, L. L. Marrec, J. Moysan, and B. Chassignole, “Direct model optimisation for data inversion. application to ultrasonic characterisation of heterogeneous welds,” *NDT & E International*, vol. 42, no. 1, pp. 47 – 55, 2009.
- [73] J. Zhang, A. Hunter, B. W. Drinkwater, and P. D. Wilcox, “Monte Carlo inversion of ultrasonic array data to map anisotropic weld properties,” *IEEE TUFFC*, vol. 59, no. 11, pp. 2487–2497, 2012.
- [74] K. M. M. Tant, E. Galetti, A. Mulholland, A. Curtis, and A. Gachagan, “Effective grain orientation mapping of complex and locally anisotropic media for improved imaging in ultrasonic non-destructive testing,” *Inverse Problems in Science and Engineering*, vol. 28, no. 12, pp. 1694–1718, 2020.
- [75] A. L. Bukhgeim, *Introduction to the theory of inverse problems*. Inverse and ill-posed problems series, Utrecht: VSP, 2000.
- [76] F. Natterer, *The mathematics of computerized tomography*. SIAM, 2001.

REFERENCES

- [77] S. I. Kabanikhin, *Inverse and ill-posed problems: theory and applications*, vol. 55 of *Inverse and ill-posed problems series; v. 55*. Berlin: Walter De Gruyter, 2011.
- [78] A. Fichtner, *Full seismic waveform modelling and inversion*. Advances in geophysical and environmental mechanics and mathematics, Berlin ; Heidelberg: Springer, 1st ed. 2011.. ed., 2011.
- [79] J. Virieux, A. Asnaashari, R. Brossier, L. Métivier, A. Ribodetti, and W. Zhou, *An introduction to full waveform inversion*. Society of Exploration Geophysicists, 2017.
- [80] L. Lucy, “Astronomical inverse problems.,” in *Reviews in Modern Astronomy*, vol. 7, pp. 31–50, 1994.
- [81] S. N. Yanushkevich, A. Stoica, V. P. Shmerko, and D. V. Popel, *Biometric inverse problems*. CRC Press, 2018.
- [82] K. Wang, T. Matthews, F. Anis, C. Li, N. Duric, and M. A. Anastasio, “Waveform inversion with source encoding for breast sound speed reconstruction in ultrasound computed tomography,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 62, no. 3, pp. 475–493, 2015.
- [83] N. Nakata, L. Gualtieri, and A. Fichtner, *Seismic ambient noise*. Cambridge University Press, 2019.

REFERENCES

- [84] S. Mao, M. Campillo, R. D. van der Hilt, F. Brengular, L. Stehly, and G. Hillers, “High temporal resolution monitoring of small variation in crustal strain by dense seismic arrays,” *Geophysical Research Letter*, vol. 46, pp. 128–137, 2018.
- [85] E. Beretta, C. Cavaterra, M. C. Cerutti, A. Manzoni, and L. Ratti, “An inverse problem for a semilinear parabolic equation arising from cardiac electrophysiology,” *Inverse Problems*, vol. 33, p. 105008, 2017.
- [86] F. Cakoni, I. de Teresa, and P. Monk, “Nondestructive testing of delaminated interfaces between two materials using electromagnetic interrogation,” *Inverse Problems*, vol. 34, no. 6, p. 065005, 2018.
- [87] M. Almekkawy, A. Carevic, A. Abdou, J. He, G. Lee, and J. Barlow, “Regularization in ultrasound tomography using projection-based regularized total least squares,” *Inverse Problems in Science and Engineering*, vol. 28, no. 4, pp. 556–579, 2020.
- [88] X. Zhang and A. Curtis, “Variational full-waveform inversion,” *Geophysical Journal International*, vol. 222, no. 1, pp. 406–411, 2020.
- [89] Y. Eren, I. Kucukdemiral, and I. Ustoglu, “Introduction to optimization,” in *Optimization in Renewable Energy Systems: Recent Perspectives*, pp. 27–74, Elsevier Inc., 2017.

REFERENCES

- [90] U. M. Diwekar, *Introduction to applied optimization / [internet resource]*. Springer optimization and its applications; v. 22, New York: Springer, 2nd ed., 2008.
- [91] D. B. Shmoys and C. Swamy, “Stochastic optimization is (almost) as easy as deterministic optimization,” in *45th Annual IEEE Symposium on Foundations of Computer Science*, pp. 228–237, 2004.
- [92] M.-H. Lin, J.-F. Tsai, and C.-S. Yu, “A review of deterministic optimization methods in engineering and management,” *Mathematical Problems in Engineering*, vol. 2012, 2012.
- [93] M. Aleardi and A. Mazzotti, “1D elastic full-waveform inversion and uncertainty estimation by means of a hybrid genetic algorithm-Gibbs sampler approach: FWI and uncertainty estimation,” *Geophysical Prospecting*, vol. 65, pp. 64–85, 2016.
- [94] S. N. Sivanandam, *Introduction to genetic algorithms / [internet resource]*. Berlin; New York: Springer, 1st ed. 2008.. ed., 2007.
- [95] D. Aleksendri and P. Carlone, “4 - soft computing techniques,” in *Soft Computing in the Design and Manufacturing of Composite Materials* (D. Aleksendri and P. Carlone, eds.), pp. 39 – 60, Oxford: Woodhead Publishing, 2015.

REFERENCES

- [96] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, *Markov chain Monte Carlo in practice*. Interdisciplinary statistics, London: Chapman & Hall, 1996.
- [97] S. Brooks, *Handbook of Markov Chain Monte Carlo*. London; New York: Chapman & Hall, 2011.
- [98] X. Zhang and A. Curtis, “Seismic tomography using variational inference methods,” *Journal of Geophysical Research: Solid Earth*, vol. 125, no. 4, p. e2019JB018589, 2020.
- [99] X. Zhang, F. Hansteen, A. Curtis, and S. de Ridder, “1-D, 2-D, and 3-D Monte Carlo ambient noise tomography using a dense passive seismic array installed on the North Sea seabed,” *Journal of Geophysical Research: Solid Earth*, vol. 125, no. 2, p. e2019JB018552, 2020.
- [100] T. Bodin and M. Sambridge, “Seismic tomography with the reversible jump algorithm,” *Geophysical Journal International*, vol. 178, no. 3, pp. 1411–1436, 2009.
- [101] E. Galetti and A. Curtis, “Transdimensional electrical resistivity tomography,” *Journal of Geophysical Research: Solid Earth*, vol. 123, no. 8, pp. 6347–6377, 2018.

REFERENCES

- [102] M. Warner, A. Ratcliffe, T. Nangoo, J. Morgan, A. Umpleby, N. Shah, V. Vinje, I. tekl, L. Guasch, C. Win, G. Conroy, and A. Bertrand, “Anisotropic 3D full-waveform inversion,” *Geophysics*, vol. 78, no. 2, pp. 59–80, 2013.
- [103] R. M. Shipp and S. C. Singh, “Two-dimensional full wavefield inversion of wide-aperture marine seismic streamer data,” *Geophysical Journal International*, vol. 151, no. 2, pp. 325–344, 2002.
- [104] J. Virieux and S. Operto, “An overview of full-waveform inversion in exploration geophysics,” *Geophysics*, vol. 74, no. 6, pp. WCC1–WCC26, 2009.
- [105] D. Vigh and E. W. Starr, “3D prestack plane-wave, full-waveform inversion,” *Geophysics*, vol. 73, no. 5, pp. 135–144, 2008.
- [106] Z. Xue, J. Sun, S. Fomel, and T. Zhu, “Accelerating full-waveform inversion with attenuation compensation,” *Geophysics*, vol. 83, no. 1, pp. A13–A20, 2017.
- [107] D. Borisov, R. Modrak, F. Gao, and J. Tromp, “3D elastic full-waveform inversion of surface waves in the presence of irregular topography using an envelope-based misfit function,” *Geophysics*, vol. 83, no. 1, pp. R1–R11, 2017.

REFERENCES

- [108] A. Klotzsche, H. Vereecken, and J. van der Kruk, “Review of crosshole ground-penetrating radar full-waveform inversion of experimental data: Recent developments, challenges, and pitfalls,” *Geophysics*, vol. 84, no. 6, pp. H13–H28, 2019.
- [109] S. Busch, J. van der Kruk, J. Bikowski, and H. Vereecken, “Quantitative conductivity and permittivity estimation using full-waveform inversion of on-ground GPR data,” *Geophysics*, vol. 77, no. 6, pp. H79–H91, 2012.
- [110] S. Busch, J. van der Kruk, and H. Vereecken, “Improved characterization of fine-texture soils using on-ground GPR full-waveform inversion,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 52, no. 7, pp. 3947–3958, 2014.
- [111] A. Kalogeropoulos, J. [van der Kruk], J. Hugenschmidt, J. Bikowski, and E. Brhwiler, “Full-waveform gpr inversion to assess chloride gradients in concrete,” *NDT & E International*, vol. 57, pp. 74 – 84, 2013.
- [112] O. C. Aguda, L. Guasch, P. Huthwaite, and M. Lowe, “3D imaging of the breast using full-waveform inversion,” *Proceedings of the International Workshop on Medical Ultrasound Tomography*, 2017.
- [113] L. Guasch, O. C. Agudo, M.-X. Tang, P. Nachev, and M. Warner, “Full-waveform inversion imaging of the human brain,” *NPJ digital medicine*, vol. 3, no. 1, pp. 1–12, 2020.

REFERENCES

- [114] L. T. Nguyen and R. T. Modrak, “Ultrasonic wavefield inversion and migration in complex heterogeneous structures: 2D numerical imaging and nondestructive testing experiments,” *Ultrasonics*, vol. 82, pp. 357 – 370, 2018.
- [115] J. Rao, M. Ratassepp, and Z. Fan, “Guided wave tomography based on full waveform inversion,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 63, no. 5, pp. 737–745, 2016.
- [116] J. Rao, M. Ratassepp, and Z. Fan, “Investigation of the reconstruction accuracy of guided wave tomography using full waveform inversion,” *Journal of Sound and Vibration*, vol. 400, pp. 317 – 328, 2017.
- [117] P. C. Gibson, “The combinatorics of scattering in layered media,” *SIAM Journal on Applied Mathematics*, vol. 74, no. 4, pp. 919–938, 2014.
- [118] P. C. Gibson, “The refined impedance transform for 1D acoustic reflection data,” *Inverse Problems*, vol. 34, no. 7, p. 075013, 2018.
- [119] P. C. Gibson, “Acoustic imaging of layered media,” *Journal of Computational Physics*, vol. 372, pp. 524–545, 2018.
- [120] L. Métivier, R. Brossier, Q. Mérigot, E. Oudet, and J. Virieux, “Measuring the misfit between seismograms using an optimal transport distance: application to full waveform inversion,” *Geophysical Journal International*, vol. 205, no. 1, pp. 345–377, 2016.

REFERENCES

- [121] Y. Yang, B. Engquist, J. Sun, and B. F. Hamfeldt, “Application of optimal transport and the quadratic Wasserstein metric to full-waveform inversion,” *Geophysics*, vol. 83, no. 1, pp. 43–62, 2018.
- [122] E. Bozdag, J. Trampert, and J. Tromp, “Misfit functions for full waveform inversion based on instantaneous phase and envelope measurements,” *Geophysical Journal International*, vol. 185, no. 2, pp. 845–870, 2011.
- [123] M. Sambridge and N. Rawlinson, “Seismic tomography with irregular meshes,” *Washington DC American Geophysical Union Geophysical Monograph Series*, vol. 157, pp. 49–65, 2005.
- [124] J. Wiskin, D. Borup, E. Iuanow, J. Klock, and M. Lenox, “3-D nonlinear acoustic inverse scattering: Algorithm and quantitative results,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, pp. 1161–1174, 2017.
- [125] P. A. Ullrich, P. H. Lauritzen, and C. Jablonowski, “Geometrically exact conservative remapping (GECORE): Regular latitude-longitude and cubed-sphere grids,” *Monthly Weather Review*, vol. 137, no. 6, pp. 1721–1741, 2009.
- [126] A. Staniforth and J. Thuburn, “Horizontal grids for global weather and climate prediction models: A review,” *Quarterly Journal of the Royal Meteorological Society*, vol. 138, no. 662, pp. 1–26, 2012.

REFERENCES

- [127] T. Bodin, M. Sambridge, and K. Gallagher, “A self-parametrizing partition model approach to tomographic inverse problems,” *Inverse Problems*, vol. 25, no. 5, p. 055009, 2009.
- [128] M. Sambridge and O. Gudmundsson, “Tomographic systems of equations with irregular cells,” *Journal of Geophysical Research: Solid Earth*, vol. 103, no. 1, pp. 773–781, 1998.
- [129] M. Sambridge, J. Braun, and H. McQueen, “Geophysical parametrization and interpolation of irregular data using natural neighbours,” *Geophysical Journal International*, vol. 122, no. 3, pp. 837–857, 1995.
- [130] D. Avis and B. K. Bhattacharya, “Algorithms for computing d-dimensional Voronoi diagrams and their duals,” *Advances in computing research*, vol. 1, pp. 159–180, 1983.
- [131] T. Bodin, M. Sambridge, N. Rawlinson, and P. Arroucau, “Transdimensional tomography with unknown data noise,” *Geophysical Journal International*, vol. 189, no. 3, pp. 1536–1556, 2012.
- [132] X. Zhang, A. Curtis, E. Galetti, and S. deRidder, “3-D monte carlo surface wave tomography,” *Geophysical Journal International*, vol. 215, no. 3, pp. 1644–1658, 2018.
- [133] A. V. Pamel, P. Huthwaite, C. R. Brett, and M. J. Lowe, “Numerical simulations of ultrasonic array imaging of highly scattering materials,” *NDT & E International*, vol. 81, pp. 9 – 19, 2016.

REFERENCES

- [134] A. Van Pamel, G. Sha, M. J. S. Lowe, and S. I. Rokhlin, “Numerical and analytic modelling of elastodynamic scattering within polycrystalline materials,” *The Journal of the Acoustical Society of America*, vol. 143, no. 4, pp. 2394–2408, 2018.
- [135] Y. Liu, A. Van Pamel, P. B. Nagy, and P. Cawley, “Investigation of ultrasonic backscatter using three-dimensional finite element simulations,” *The Journal of the Acoustical Society of America*, vol. 145, no. 3, pp. 1584–1595, 2019.
- [136] M. Bock, A. K. Tyagi, J.-U. Kreft, and W. Alt, “Generalized Voronoi tessellation as a model of two-dimensional cell tissue dynamics,” *Bulletin of Mathematical Biology*, vol. 72, no. 7, pp. 1696–1731, 2010.
- [137] S. Johnson, *The Ghost Map: The Story of London’s Most Terrifying Epidemic—and How It Changed Science, Cities, and the Modern World*. Penguin Publishing Group, 2006.
- [138] G. Aletti and D. Saada, “Survival analysis in Johnson-Mehl tessellation,” *Statistical Inference for Stochastic Processes*, vol. 11, pp. 55–76, 2008.
- [139] B. Bollobás and O. Riordan, “Percolation on random Johnson-Mehl tessellations and related models,” *Probability Theory and Related Fields*, vol. 140, pp. 319–343, 2008.

REFERENCES

- [140] J. Belhadj, T. Romary, A. Gesret, M. Noble, and B. Figliuzzi, “New parameterisation for Bayesian seismic tomography,” *Inverse Problems*, vol. 34, no. 6, p. 065007, 2018.
- [141] *OnScale (formally PZFlex)*. 770 Marshall Street, Redwood City, CA 94063 USA, 2019.
- [142] J. L. Rodgers and W. A. Nicewander, “Thirteen ways to look at the correlation coefficient,” *The American Statistician*, vol. 42, no. 1, pp. 59 – 66, 1988.
- [143] M. K. Cowles and B. P. Carlin, “Markov chain Monte Carlo convergence diagnostics: a comparative review,” *Journal of the American Statistical Association*, vol. 91, no. 434, pp. 883–904, 1996.
- [144] L. Tierney and A. Mira, “Some adaptive Monte Carlo methods for Bayesian inference,” *Statistics in Medicine*, vol. 18, 1999.
- [145] J. P. Green and D. I. Hastie, “Reversible jump MCMC,” *Genetics*, vol. 155, pp. 1391–1403, 2009.
- [146] M. S. Hassouna and A. A. Farag, “Multistencils fast marching methods: A highly accurate solution to the eikonal equation on cartesian domains,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 9, pp. 1563–1574, 2007.
- [147] C. E. Metz, “Basic principles of ROC analysis,” in *Seminars in Nuclear Medicine*, vol. 8, pp. 283–298, 1978.

REFERENCES

- [148] T. Fawcett, “An introduction to ROC analysis,” *Pattern Recognition Letters*, vol. 27, no. 8, pp. 861–874, 2006.
- [149] R. Jarvis, P. Cawley, and P. B. Nagy, “Performance evaluation of a magnetic field measurement NDE technique using a model assisted Probability of Detection framework,” *NDT & E International*, vol. 91, pp. 61 – 70, 2017.
- [150] J. P. Fouque, J. Garnier, G. Papanicolaou, and K. Sølna, *Wave propagation and time reversal in randomly layered media*. Springer, 2007.
- [151] B. N. P. Agarwal and S. Srivastava, “FORTRAN codes to implement enhanced local wave number technique to determine the depth and location and shape of the causative source using magnetic anomaly,” *Computers & Geosciences*, vol. 34, no. 12, pp. 1843 – 1849, 2008.