


# A transdimensional Bayesian approach to ultrasonic travel-time tomography for non-destructive testing

K M M Tant<sup>1</sup> , E Galetti<sup>2</sup>, A J Mulholland<sup>1</sup>, A Curtis<sup>2</sup> and A Gachagan<sup>3</sup>

<sup>1</sup> Department of Mathematics and Statistics, University of Strathclyde, Glasgow, G1 1XH, United Kingdom

<sup>2</sup> School of Geosciences, University of Edinburgh, Edinburgh, United Kingdom

<sup>3</sup> Centre for Ultrasonic Engineering, University of Strathclyde, Glasgow, United Kingdom

E-mail: [katy.tant@strath.ac.uk](mailto:katy.tant@strath.ac.uk)

Received 11 October 2017, revised 22 May 2018

Accepted for publication 5 June 2018

Published 29 June 2018



CrossMark

## Abstract

Traditional imaging algorithms within the ultrasonic non-destructive testing community typically assume that the material being inspected is primarily homogeneous, with heterogeneities only at sub-wavelength scales. When the medium is of a more generally heterogeneous nature, this assumption can contribute to the poor detection, sizing and characterisation of defects. Prior knowledge of the varying wave speeds within the component would allow more accurate imaging of defects, leading to better decisions about how to treat the damaged component. This work endeavours to reconstruct the inhomogeneous wave speed maps of random media from simulated ultrasonic phased array data. This is achieved via application of the reversible-jump Markov chain Monte Carlo method: a sampling-based approach within a Bayesian framework. The inverted maps are used in conjunction with an imaging algorithm to correct for deviations in the wave speed, and the reconstructed flaw images are then used to quantitatively assess the success of this methodology. Using full matrix capture data arising from a finite element simulation of a phased array inspection of a heterogeneous component, a six-fold improvement in flaw location is achieved by taking into account the reconstructed wave speed map which exploits almost no *a priori* knowledge



Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/3.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

of the material's internal structure. Receiver operating characteristic curves are then calculated to demonstrate the enhanced probability of detection achieved when the material speed map is accounted for.

Keywords: inverse problems, Bayesian inversion, wave propagation, ultrasonics, non-destructive testing, imaging, tomography

(Some figures may appear in colour only in the online journal)

## 1. Introduction

The oil and gas, nuclear, power and aerospace industries are only a subset of the sectors dependent on the routine maintenance of safety-critical structures [1]. Failure to detect structural weaknesses in components integral to the work being carried out by these industries can be catastrophic. Ultrasonic non-destructive testing is a technique which involves the transmission of mechanical waves through the component under inspection [2]. As in medical ultrasound, these waves can be passed through the component and subsequently collected without disturbing the internal composition of the medium [3]. The resulting datasets can then be used to create images of internal component features. However, inspections can become challenging when the material exhibits an inhomogeneous response to the passing wave energy [4, 5]. Due to the spatial variation of material properties, ultrasonic wavepaths are distorted and their expected arrival times (on which current imaging algorithms are based) cannot be accurately modelled by straight rays travelling at a constant velocity. When commonly used imaging algorithms (which assume a constant wave speed throughout the inspection domain) are applied to these ultrasonic datasets, the resulting images typically display poorly characterised and mislocated flaws [6, 7]. Ultrasonic wave propagation through inhomogeneous media has previously been studied using models [8–10] and simulations within finite element packages [5, 11] and it has been shown that some prior knowledge of how the material properties vary spatially can be used to correct for the deviation in wave speed and path, producing improved reconstructions or images of any internal defects [6, 7, 12].

Previous studies have extracted microstructural maps experimentally which allowed insight into the variation in structure and crystallography of materials and how this can effect wave propagation [6]. However, acquiring these maps usually requires a destructive aspect where cross sections of the material are analysed via macrographs or interrogated by techniques which require sophisticated equipment (such as electron backscatter diffraction (EBSD) [6] or spatially resolved acoustic spectroscopy (SRAS) [13]). Each component under inspection will have some dependence on its manufacturing conditions so it is not prudent to project the map of one component onto another as they will vary from case to case. So, although these destructive measurements are important for understanding complex materials, the techniques used to gather them cannot practically be transferred to an *in situ* environment. Modelling the material map using a small number of parameters presents an attractive, non-destructive alternative. In the weld specific case, MINA (modelling of anisotropy based on notebook of arcwelding) uses information from the welding procedure such as the dimensions of the weld pool, number, inclination and order of weld passes, and the remelting rates [14]. This allows the consideration of the physical polycrystal growth and a reasonably accurate prediction of the wave's passage throughout the weld can thus be made [15–17].

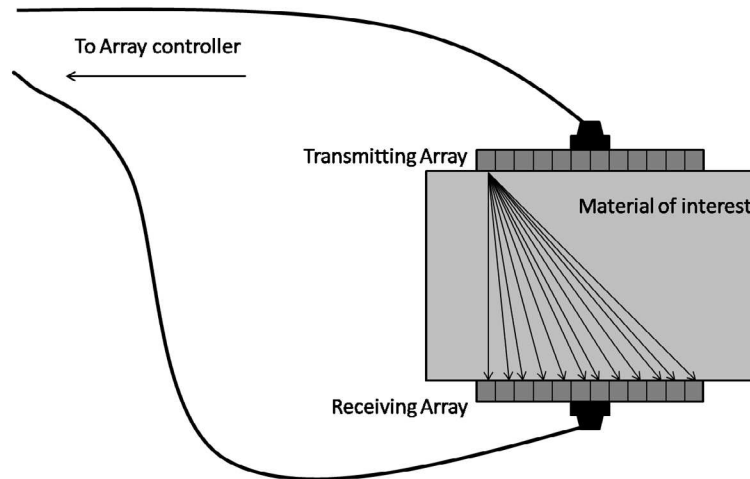
Tomographic inversion presents another practicable approach where the aim is to recover some information on an object's internal properties from measurements taken on its surface [18]. Tomographic imaging is commonplace in diagnostic medicine and seismology, and

within the NDT community much effort has been expended on developing ultrasonic guided wave tomography to detect and locate corrosion in pipes [19–23].

Although guided wave tomography has proved successful, it is restricted to the examination of plate-like structures. For the inspection of the third dimension of an object (depth), we might consider bulk waves (shear waves could also be considered but will not be studied in this work). A tomographic approach has previously been taken in [7] to map the anisotropic distribution of grains within austenitic welds using a Markov chain Monte Carlo (MCMC) approach, where an initial estimate of the weld map is taken from a model [24] and the ultrasonic wave propagation through the weld is predicted using Dijkstra’s shortest path algorithm. A weld map is generated by taking the mean of the postconvergence ensemble of samples and this is then used in conjunction with the total focussing method (TFM) [25] to produce more accurate images of a flaw. In this paper, we focus on heterogeneous media in which the constituent materials are locally isotropic (as opposed to the locally anisotropic regions studied in [7]) and the reversible jump Markov chain Monte Carlo (rj-MCMC) method [26], a sampling-based ensemble inference approach within a Bayesian framework, has been selected as the inversion scheme. This method has already been used successfully in seismology to map the spatially varying wave speed within the Earth’s crust [27–29] where the spatial domain was partitioned using a Voronoi tessellation [30]. This paper applies this methodology for the first time within an NDT setting, mapping the wave speed over a rectangular metal component where the sources and receivers are spaced at regular intervals (as they are in linear phased arrays [1]). We introduce the multistencils fast marching method (MSFM) [31] as the forward model, to obtain more precise realisations of the travel-time field through the heterogeneous structure than those obtained using the standard fast marching method. This forward model is successfully adapted to handle data arising from pulse-echo inspections (that is where a single array is used for both transmission and reception of the waves simultaneously). Once we have inverted for the material map, it is used in conjunction with the TFM (an algorithm which we will refer to as TFM+) to correct for deviations in the expected arrival times of each wave caused by the heterogeneous nature of the material and the subsequent refraction of the wave, thus producing better focussed images of any embedded defects. The success of the method is analysed via the accuracy of the resulting flaw reconstructions and how they compare to those arising from application of the standard TFM with a constant wave speed assumption. Analysis of these images can of course be affected by the subjective choice of threshold at which the images are plotted. To address this element of subjectivity and ensure that the comparisons are made fairly, we present an objective methodology for producing receiver operating characteristic (ROC) curves [32] from each image, and demonstrate the enhanced probability of detection achieved when the material map is accounted for within the imaging algorithm.

## 2. Data and processing

The production of ultrasonic phased arrays, which are capable of simultaneously transmitting and receiving ultrasound signals, has surged in recent years [1]. These multi-element transducers allow for greater coverage (hence faster inspection times) and provide the possibility of performing inspections with ultrasonic beams of various angles and focal lengths, giving rise to a richer set of data. These dual-purpose arrays provide us with two primary experimental set-up options. The most common is the pulse-echo inspection, where a single phased array is employed to simultaneously transmit and receive signals. This set-up has the obvious advantage of requiring only single-sided access to a component. An alternative set-up is the pitch-catch arrangement, where two phased arrays are employed, one to transmit and one to receive, and are placed at two carefully considered locations on the component’s surface [33] so that




**Figure 1.** Ultrasonic phased array through-transmission experimental set-up.

the receiving array can record the energy input to the system by the transmitting array. For the work presented in this paper, we begin by studying the case known as through-transmission, where two arrays are placed directly opposite each other, on either side of the component (see figure 1). This set up simplifies the extraction of the time of flight (ToF) between each pair of elements and thus reduces the amount of uncertainty present in the input data for our inversion algorithm. We then progress to examine the pulse-echo case which is more relevant to industrial applications. In both cases we have  $N$  transmitters and  $N$  receivers, giving rise to  $N^2$  A-scans (the time series data), which are stored in an  $N \times N \times T$  matrix (where  $T$  is the number of time steps recorded by the receiving array), which in the NDT literature is termed the full matrix capture (FMC) [25].

### 2.1. The observed data

The FMC is a rich dataset which contains information on the wave path between every pair of transmit and receive array elements. However, for our current implementation of the algorithm, we need only the first time of arrival of the wave between each transmitter and receiver. In datasets arising from simulated through-transmission inspections, this is easily extracted as it is measured as the first point in time at which the receiving element detects any signal (that is the first point when  $A_{t_x, r_x}(t) > \varepsilon$ , where  $A$  is our time domain signal transmitted at element  $t_x$  and received at element  $r_x$ ,  $\varepsilon$  is some chosen threshold, and  $0 \leq t \leq T$  is the time period over which the signal is collected). For application to data arising from pulse-echo inspections, the corresponding ToF is the shortest time taken for the longitudinal wave to travel from the transmitting element to the receiving element, via the back wall of the component. In that case, scattering from facets of the microstructure can interfere with the detection of the first time of arrival and an element of uncertainty is thus introduced. However, cross correlation of the input signal with the received signal allows us to obtain sufficiently good estimates of the ToF in such cases [34]. Once the arrival times have been obtained, they are stored in a ToF matrix  $T_0$ , where each element  $t_{t_x, r_x}$  represents the time taken for the wave to travel from transmitting element  $t_x$  to receiving element  $r_x$ . In a homogeneous, isotropic medium, the time taken is dependent only on distance and so we obtain banded matrices as shown in figure 2(a).



**Figure 2.** Time of flight (ToF) matrices for (a) a homogeneous, isotropic medium and (b) a heterogeneous medium with locally isotropic regions.

However, if the material properties vary throughout the component as they would in heterogeneous media, these bands are distorted and we obtain a matrix as displayed in figure 2(b).

### 3. Ultrasonic travel-time tomography

#### 3.1. Material parametrisations using Voronoi diagrams

To minimise the degrees of freedom within our inverse problem, we reconstruct a lower resolution map of the material's heterogeneous structure than that afforded by destructive, experimental measurements. To this end, we use Voronoi diagrams to randomly partition our domain into cells. These diagrams have already been used successfully as the basis for FE simulations of waves propagating through polycrystalline materials [35] and to parametrise tomographic imaging problems [36]. Given an arbitrary set of *seeds*  $S$ , the Voronoi diagram consists of a set of non-overlapping convex regions where any point within a *cell* is closer to the seed of that cell than any other seed. Note that although Voronoi diagrams can partition  $n$ -dimensional spaces, we are only interested in the two dimensional case (our linear array of sensors allow us to inspect a two dimensional slice of the material) and so a single seed  $s_i$  is given by two dimensional Cartesian co-ordinates. To parametrise a heterogeneous material, we introduce a third parameter to each cell: its locally isotropic wave speed  $v_i$ . Note that we assume that the density throughout the component's domain is constant and so this variance in wave speed causes a mismatch in mechanical impedance between neighbouring cells. Thus we have a material model  $m(S, V, M)$  ( $V = \{v_i : i = 1, \dots, M\}$  is the set of assigned wave speeds) with  $3M + 1$  unknowns (since  $M$ , the number of cells, is itself an unknown), and  $N^2$  equations which describe the known time of arrival between every transmit/receive pair of elements.


#### 3.2. The forward model

Now that the material's geometry has been parametrised, we require an efficient forward model which outputs the ToF matrix  $T_m$  for any particular instance of the material model  $m$ . We assume that the first time of arrival represents the arrival time of the longitudinal wave (as transverse waves travel more slowly) and so mode conversion at cell interfaces is ignored

in our forward model. In this paper, we compare two forward models. The first assumes that the wave travels in a straight line between the transmitter and receiver, neglecting the effects of refraction (in the pulse-echo case, we assume that the wave travels in a straight line from the transmitter to the midpoint between the transmitter and receiver on the back wall of the component, before reflecting and travelling in a straight line back to the receiver). Given a material model  $m$ , the distance a wave travels through each distinct region lying on this path is known and, when coupled with the known wave speeds in each region, gives rise to the ToF between the two points. Although this represents very little of the true physics, this forward model is fast and surprisingly effective for locally isotropic geometries when inspected in the through-transmission set-up.

In the second forward model, we consider the effects of raybending by implementing the fast marching method (FMM) [37] (unlike Dijkstra's algorithm, which is used in [7], the fast marching method diminishes the grid bias and converges to the underlying geodesic distance when the grid step size tends to zero [38]). Let  $\tau(x_i, y_j, t_x)$  denote the shortest time for a wave to travel from the transmitter  $t_x \in \partial I$ , on the boundary of the discretised representation of our image domain  $I = x \times y$ , to the point  $(x_i, y_j) \in I$  (in our setting,  $x$  and  $y$  are typically vectors of points spaced at 1 mm intervals along the horizontal and vertical axes of our component). The travel-time field  $\tau$  can be obtained by solving the Eikonal equation,  $|\nabla\tau| = 1/\hat{m}$  using an upwind finite difference scheme [39–41] where  $\hat{m}$  is a recti-linear grid of wave speeds obtained by discretising the material model  $m(S, V, M)$ . By calculating the travel-time field for each transmitter, the shortest travel-time between each transmitter  $t_x$  and receiver  $r_x$  can be calculated and the matrix  $T_m$  constructed. It is well known that the fast marching method is not accurate along diagonal trajectories on coarsely discretised domains as it considers only nearest neighbours of each node and, as the wave front propagates through the rectangular grid, errors accumulate along the directions which lie between the axes vectors. To combat this, we adopt the MSFM [31] which operates by additionally rotating the initial four-point stencil by  $45^\circ$ , allowing neighbouring nodes along the diagonal to contribute to the shortest time calculations at each point. Note that here we use only a first order finite difference scheme to approximate the directional derivative as it has been proven to be unconditionally stable [39]. However, a mixed order scheme which is nominally second order accurate (it reverts to a first order approximation when the travel-times for the second order approximation are unavailable) can also be implemented within the multistencil framework [31].

To demonstrate the advantage of employing the MSFM, figure 3 depicts two travel-time fields arising from solving the Eikonal equation through a heterogeneous material model  $\hat{m}$  (the discretised representation of the map shown in figure 10(a)) using (a) the standard fast marching method and (b) MSFM. In both cases the source is located at the top of the domain, 24 mm along the horizontal axis. The travel-times are calculated over a coarse grid with  $1 \text{ mm}^2$  cells to allow efficient computation over the entire domain (which has an area of  $64 \text{ mm} \times 40 \text{ mm}$ ). The arrival times at a set of points placed at a depth of 40 mm, at 2 mm intervals along the horizontal axis, are plotted for both cases in plot (c) (the dashed line arises from the FMM and the dotted line arises from the MSFM algorithm). To examine the accuracy with which the travel-times are calculated using these methods, a full finite element simulation of the wave propagation through this material model was also run. The domain, measuring  $64 \text{ mm} \times 40 \text{ mm}$ , was meshed with square elements with side lengths of approximately  $260 \mu\text{m}$ . A 1.5 MHz sinusoidal pulse was used to excite a single source located 24 mm along the  $x$  axis at a height of  $y = 40 \text{ mm}$ . An array of 32 receivers were distributed at 2 mm intervals along the bottom of the geometry ( $y = 0 \text{ mm}$ ). The solid line in figure 3(c) depicts estimates of the first arrival times extracted from this full finite element simulation of the wave propagation through the medium (this data itself is subject to minor errors dependent on the threshold at




**Figure 3.** Travel-time fields through a heterogeneous material model  $\hat{m}$  calculated by (a) the standard fast marching method and (b) the multistencils fast marching method (MSFM). The transmitting ultrasonic array element  $t_x$  is located at the top of the component at a horizontal position of 24 mm. Plot (c) shows the calculated arrival times at the receiver locations (placed at 2 mm intervals along the horizontal axis at a depth of 40 mm) obtained by a finite element simulation of the wave propagation (solid line), the fast marching method (dashed line) and the MSFM (dotted line).

which the first time-of-arrival is estimated). Note that the FMM achieves a reasonable fit when the wave is propagating vertically (around receiver index 12, directly opposite the source) but fails to capture the behaviour of the data elsewhere. The MSFM is more accurate across the full set of receivers and this improved accuracy incurs a small additional computational expense (circa 7%). We see in section 3.3 that this additional expense is offset by the fact that we do not need to consider ray length when parametrising the data uncertainty.

Note that in the pulse-echo case, the shortest time for a wave to travel between the transmitter and receiver arises from the path that the surface wave takes. However, we want to consider the time of arrival taken for the longitudinal wave to travel from our transmitter at  $t_x$  to the back wall of the component at a depth  $y = d$ , back to the receiver at  $r_x$ , and so the grid on which we implement the MSFM must be modified. This is achieved by reflecting the material geometry across the horizontal axis at  $d$ , effectively doubling the size of the domain and creating a virtual receiving array at  $y = 2d$  (see figure 4). The wave is then propagated through the extended material map from the original transmitter  $t_x$  to the mirrored, virtual receiver  $r_x^*$ .

### 3.3. Uncertainty parametrisation

An important consideration to make when employing a transdimensional inversion scheme is the parametrisation of the data uncertainty. This uncertainty arises from data measurement errors and from simplification of the physical phenomena by the forward model [42]. Figure 3(c) showed that by using a more sophisticated forward model, the residuals between the modelled data and observed data (in this case arising from a finite element simulation—solid line) can be reduced but not eliminated. Regardless of the model’s enhanced ability to capture the physics of the problem, any remaining differences will be magnified when the observed data arises from experiment: as data noise becomes more prevalent, extracting the first times of arrival becomes subject to increased error. In this work, we treat the uncertainty as a single unknown, which aggregates the uncertainty inflicted on the system by both data noise and model deficiencies. Much work was done in [29] to examine different approaches to data noise parametrisation, and it was shown that parametrising the travel-time uncertainty as a function of ray length is an effective strategy. It can be seen from figure 3(c) that when the standard FMM is employed (dashed line), this assumption is valid: the points directly under the source (around receiver index 12) will presumably arise from ray paths with a



**Figure 4.** The MSFM requires modification to cope with pulse-echo inspection geometries. The material map is mirrored over the horizontal axis at  $y = d$  and a virtual receiver  $r_x^*$  is created at  $y = 2d$ .

shorter length than those lying near the vertical edges of the domain, and accordingly, these exhibit comparably smaller errors. However, to implement this parametrisation of the data noise exactly, the ray must be traced from the receiver to the source through the calculated travel-time field, creating significant additional computational expense. Fortunately, it can be observed that when the MSFM is implemented, the travel-time uncertainty appears to be independent of ray length (figure 3(c) dotted line) and so this approach is less relevant and the ray tracing step can be neglected. Instead, we will allow the algorithm to infer the uncertainty level present in the system by drawing values from a uniform distribution over some predefined range. The level of uncertainty attributed to the dataset has a direct impact on the complexity of the solution. By restricting the range too much, the algorithm will overfit the data, increasing the complexity of the model in order to minimise the data misfit. Permitting larger data uncertainties will allow the algorithm to fit the data with lower dimensional models. In this paper, we will allow the standard deviation of the noise parameter,  $\sigma_n$  to explore a range of values in the interval between  $0.01 \mu\text{s}$  and  $1 \mu\text{s}$  (the average travel-times are around  $10 \mu\text{s}$ ) and in each case the standard deviation in the uncertainty converges to a distribution with mean value in this range.

### 3.4. A probabilistic framework

The rj-MCMC method produces a posterior distribution for transdimensional spaces (that is when the number of degrees of freedom of the material model is not fixed). Its basis in a Bayesian framework necessitates that all information is written in probabilistic terms. The posterior probability density function is given by Bayes' rule,  $p(m|T_0) \propto p(T_0|m)p(m)$ , where  $p(m)$  is the *a priori* probability of the material model  $m$  and  $p(T_0|m)$  is the likelihood that the observed ToF data  $T_0$  arises from that model. Naturally, the likelihood must somehow



account for the misfit between the observed data and the data arising from a given material model. In this work we use the least squares misfit function  $\phi = \|(T_m - T_0)/\sigma_n\|^2$ , where  $\sigma_n^2$  is the variance of the uncertainty parameter. The best fitting material model  $m$  will be the model where this misfit is minimised. This is equivalent to maximising the probability of a Gaussian likelihood function and so can be written as  $p(T_0|m) \propto \exp(-\phi/2)$ . To calculate the posterior probability density function we also require information on the prior,  $p(m)$ . In this work, we choose the prior probability density functions for each model parameter to be a uniform distribution, as used in [27–29]. However, other weakly informative priors such as Jeffrey’s prior [43] or the Gamma distribution could be implemented and will be studied by the authors in future applications of the method. Given that we are interested in reconstructing random media and we have no prior knowledge of the covariance between the model parameters, we choose the partitioning of the spatial domain to be independent of the regional wave speed assignment and system noise level. The full *a priori* probability density function can then be written as a product of the probability density functions of the individual model parameters. Firstly, we allow the prior on the number of Voronoi cells used to parameterise the material’s underlying structure,  $p(M)$ , to be defined by a discrete uniform distribution given by

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

where  $\mathcal{M} = [M_{\min}, M_{\min} + 1, \dots, M_{\max} - 1, M_{\max}]$ ,  $\Delta M = M_{\max} - M_{\min} + 1$  and the bounds  $M_{\min}, M_{\max} \in \mathbb{Z}$  are chosen to reflect the wavelengths present in our system (our ultrasonic wave is most distorted by regions of size commensurate with the wavelength and so this determines the minimal resolution we require and can thus be used to define  $M_{\max}$ ). The wave speeds,  $v_i \in V$ , can be altered to reflect any prior information on the constituent materials of the component under inspection. In the case of the media studied in this paper, we assume the wave speeds are uniformly distributed

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

and due to the independence of the wave speed of one cell from another we have

$$p(V|M) = \prod_{i=1}^M p(v_i).$$

Here,  $v_i$  is measured in  $\text{m s}^{-1}$  and  $\Delta v = v_{\max} - v_{\min} + 1$ , where the bounds  $v_{\min}$  and  $v_{\max}$  on the range of  $v_i$  can be chosen by selecting sensible bounds on the speed of sound through the materials of interest. This can of course be further restricted if reliable statistics on the distribution of materials (and therefore speeds) throughout the component under inspection are available. However, it is important to note that the choice of these limits directly influences the efficiency and ability of the rj-MCMC algorithm to reconstruct a viable approximation of the material map: providing a narrow prior on the wave speed could inhibit the algorithm’s ability to properly sample the model space, subsequently biasing the solution. Conversely, if the prior is prescribed as a very broad uniform distribution, the model space could become prohibitively large and solutions which fit the observed data well but produce material maps which exhibit properties that are unphysical, may be considered. Assuming that the seed positioning has a uniform probability distribution, and that our  $M$  seeds must have  $M$  distinct locations (that is the seeds cannot lie on top of each other), we have

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

where  $|I_*|$  is the cardinality of our computational domain where the discretisation is determined by the numerical precision with which the seed co-ordinates are assigned. Finally, to define the prior on the level of uncertainty present in the system, we have

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

where  $\Delta\sigma_n = \sigma_n^{\max} - \sigma_n^{\min} + 1$  and  $\sigma_n^{\min}$  and  $\sigma_n^{\max}$  are chosen as discussed in section 3.3. Thus, the probability of a given model  $m$  is

$$p(m) = p(M) p(V|M) p(S|M) p(\sigma_n) = \frac{M!(|I_*| - M)!}{\Delta M (\Delta v)^M |I_*|! \Delta\sigma_n},$$

provided that the model parameters lie within their predefined ranges, and is equal to 0 otherwise. Since typically  $|I_*| \gg 2M$  and  $\Delta v > 1$ , it is clear that models with fewer cells will be assigned a higher probability: this is in line with the parsimonious properties of Bayesian inference.

### 3.5. The rj-MCMC method

The nature of the rj-MCMC method can be inferred from its name: reversible-jump refers to the process of allowing dimensional jumps (that is, changes in the number of Voronoi cells  $M$ ) in our material model space, which can later be reversed [26]; Markov chain dictates that we possess the memoryless characteristic so that each perturbation is solely dependent on the current model and not on its predecessors; and finally Monte Carlo stipulates that the process is iterated many times. To begin, the initial number of Voronoi cells,  $M$ , the system noise level  $\sigma_n$ , the Voronoi seeds  $s_i$  and corresponding wave speed values  $v_i$  are drawn from their uniform prior distributions. This sparsity of prior information differentiates this work from the weld-specific algorithms, in for example [7, 14]. Using either of the two forward models discussed in section 3.2, a first time of arrival matrix  $T_m$  is calculated. This is compared with the first time of arrival matrix extracted from the observed dataset,  $T_0$ , and the posterior value  $p(m|T_0)$  for the initial model is calculated. The material model is then perturbed to create a new model  $m'$ . Since the problem is ill-posed and thus sensitive to small changes, the model parameters are perturbed individually to isolate their effects. In [54] it is shown that large steps through parameter space can change the complexity of the solution towards which the rj-MCMC algorithm converges, so our proposal strategy avoids this potential problem. Each model perturbation is subject to proposal distributions which are the conditional probabilities of proposing a state  $m'$  given  $m$ . In this work, the model can be perturbed in one of five ways:

- (1) A wave speed change where the value of  $v_i$  is changed in cell  $i$ , subject to

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

where  $v'_i$  is the proposed wave speed in cell  $i$ ,  $X \sim \mathcal{N}(0, 1)$  is a random variable drawn from the standard normal distribution with mean 0 and variance 1, and  $\sigma_v$  is the standard deviation of the proposal distribution for wave speed.

- (2) A system noise change where the value of  $\sigma_n$  is changed subject to

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

where  $\sigma'_n$  is the proposed standard deviation on the noise parameter, and  $\sigma_n^{\text{prop}}$  is the standard deviation of the proposal distribution for a noise perturbation.

- (3) A cell move where the coordinates of a seed  $s_i = (x_i, y_i)$  are changed subject to

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

and

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

to obtain the perturbed Voronoi seed location  $s'_i = (x'_i, y'_i)$ , where  $\sigma_c$  is the standard deviation of the proposal distribution for a cell move.

- (4) A cell birth where an additional seed at a randomly selected location  $s_{n+1}$ , is added to the set  $S$  (recall  $S$  is the set of 2D cartesian coordinates which define the seed locations), where  $s_{n+1}$  is drawn from the prior distribution on  $S$ . The wave speed assigned to this new cell is given by

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

where  $v_{n+1}^*$  is the wave speed in the previous model at the point  $s_{n+1}$  and  $\sigma_b$  is the standard deviation on the wave speed proposal for a birth perturbation.

- (5) A cell death where a randomly selected seed at location  $s_i \in S$  is removed from the set  $S$ .

Note that steps 4 and 5 result in a dimensional jump in the model: if our model  $m$  has  $M$  cells, a birth would result in the perturbed model  $m'$  having  $M + 1$  cells, whilst a death would leave only  $M - 1$  cells. Once a perturbation has been made, the posterior  $p(m'|T_0)$  is calculated. The Metropolis–Hastings criterion

$$p(\text{accept}) = \min \left( 1, \frac{p(m'|T_0)}{p(m|T_0)} \right),$$

is used to determine whether the perturbation should be accepted. Once accepted, the model  $m'$  replaces the material model  $m$  and the process begins again. If  $m'$  is rejected, the model is discarded and the original material model  $m$  is repeated as the sample from the current iteration and is perturbed again in the following iteration.

Poor choices of the standard deviations  $\sigma_v, \sigma_{\text{prop}}^n, \sigma_c, \sigma_b$  will result in a slow exploration of the model space, negatively impacting the convergence of the Markov chain to the posterior probability distribution. There exists an extensive literature on how to best choose proposal probability distributions which efficiently search the model space [44–46]. In our work, they are tuned until the acceptance rates (that is the fraction of proposed samples that are accepted) lie somewhere between 23% and 44%, as declared optimal in [47]. Furthermore, to reduce the amount of tuning required, a delayed rejection scheme as used in [27] has been implemented for perturbations of type 1 and 3 (wave speed and seed location perturbations). The rejection of perturbations in a Metropolis–Hastings algorithm is of course essential in guaranteeing that the Markov chain converges to the intended posterior distribution. However, performance of the algorithm is also improved by decreasing the probability of remaining in the current state [48, 49], and so it is desirable to avoid persistent rejection of perturbations. Implementation of the delayed rejection scheme allows secondary perturbations with smaller standard deviations ( $\sigma_v^{\text{dr}} < \sigma_v$  and  $\sigma_c^{\text{dr}} < \sigma_c$ ) to be made on the rejection of the initial perturbation.

### 3.6. Sampling the posterior distribution

To generate a reliable estimate of the posterior probability distribution, the model must be evaluated iteratively as it explores the model space. After the initial sampling period (which is influenced by the random starting point), the Markov chain begins importance sampling, where the higher likelihood regions of the model space are more densely sampled. Ideally, the algorithm is terminated when convergence of the Markov chain is achieved: that is when our ensemble of models exhibits a density proportional to the posterior probability distribution. This convergence is problematic to detect as it is effectively the measure of how well our constructed sample of the distribution represents the underlying stationary distribution of the Markov chain (which is unknown in practice). Diagnostic tools to measure convergence of the chain are reviewed in [50] and it is concluded that there is no dependable method to ensure that the finite sample generated by the MCMC method truly represents the underlying stationary distribution. And so, in this work, once the acceptance ratios have been tuned (as discussed in section 3.5), the misfit and data noise are monitored throughout the running of the algorithm to ensure that they are converging to values which align with our expectations based on our prior knowledge of the residuals between an observed dataset and a modelled dataset. Then, once the chain has been terminated, the number of cells, system noise and wave speed at arbitrarily selected points in the domain are plotted versus iteration (trace plots). We assume that the algorithm has successfully converged to the posterior distribution if these trace plots exhibit stationarity (that is when their statistical properties are constant in time). Note that it does not make sense for us to study the parameters which define the Voronoi diagrams in this way due to the transdimensional nature of the algorithm [27]. Once we believe we have achieved convergence, the initial samples are discarded (this is known as the burn in period and by removing these samples we can discount any bias towards the initial model [51]) and we use the ensemble of remaining samples to characterise the posterior probability distribution. After the burn in period, the ensemble of accepted models is resampled at a decimation interval  $\kappa$ , where  $\kappa$  is the relaxation time of the random walk (the number of steps required before we can expect to obtain a model that is considered independent of the last). To produce the mean image of the material map we project the sampled partition models into a spatial domain and average the results across all of the samples at each point in space individually. Given the large number of samples, when the Voronoi tessellations are stacked, the cells overlap and the resulting mean regional material map is effectively a continuous function of the plane. This framework allows the study of the standard deviation of the wave speed throughout the image domain (the standard deviation across all models at each point in space can be calculated as easily as the mean) which can be exploited for uncertainty quantification studies.

## 4. Results

### 4.1. Reconstruction of simple geometries


When several material phases are randomly distributed in a component it can be difficult to visually assess the agreement between the reconstructed material map and the actual material map. Therefore we begin by studying a synthetic material featuring large monophasic regions. We examine the case of a 25 mm diameter disc with a longitudinal wave speed of  $4000 \text{ m s}^{-1}$  embedded in a host material with a longitudinal wave speed of  $6500 \text{ m s}^{-1}$  (see figure 5(a)). This example demonstrates that although we partition the domain using polygonal Voronoi cells, we can in fact reconstruct curved boundaries in the mean wave speed map. A phased array inspection (in the through-transmission format) was simulated in the finite

element package PZFlex [52] where the geometry was meshed with square elements with side lengths of approximately  $260 \mu\text{m}$ . Absorbing boundary conditions were employed on the vertical edges of the domain to ensure energy was not reflected back into the domain at these edges (we assume the domain we are imaging forms only a subsection of a larger component). Free boundary conditions were assigned to the top and bottom edges of the domain on which the arrays were placed. A 1.5 MHz sinusoidal pulse (giving rise to a wavelength of  $\lambda = 4.3 \text{ mm}$ ) was used to excite the system which consisted of two 32 element phased arrays placed directly opposite each other on either side of the rectangular component. The arrays had a pitch (element spacing) of 2 mm and the depth of the component was 50 mm. The simulated domain extended 5 mm either side of the array giving rise to an inspection domain with dimensions  $74 \text{ mm} \times 50 \text{ mm}$  (although only the  $64 \text{ mm} \times 50 \text{ mm}$  area directly between the arrays is reconstructed below). It is important to note that the PZFlex software models many of the physical phenomena present in the ultrasonic phased array inspection, for example mode conversion and diffraction, and it has previously been validated against experimental data [53]. We treat these aspects of the data as system noise. Coupled with the subjectivity involved in selecting the first time of arrival from the simulated time-domain data, the levels of uncertainty present in the simulated data imitate those present in experimentally collected data and so the addition of synthetic noise is not required here.


The resulting dataset was then interrogated twice: firstly, the inversion algorithm employed the straight ray assumption in the forward model, before raybending was included via the implementation of the MSFM algorithm. A deliberately high contrast in wave speed was chosen to assess the importance of using raybending in the forward model: in this case, a straight ray which would bisect the disc would be faster circumventing the disc entirely and it is interesting to examine how this affects the reconstructions.

The rj-MCMC algorithm was run for 20000 samples, where the first 2000 samples were discarded (the burn in period), and the remaining models were sampled at an interval of  $\kappa = 100$ . To ensure a fair comparison between the forward models, a random seed was used so that the initial Voronoi diagram was the same in both cases (shown in figure 5(b)). Considering the case where the straight ray assumption is made first, we obtain the material map plotted in figure 5(c), which shows the mean of the posterior probability distribution. Although a lower wave speed anomaly has been reconstructed in the centre of the domain, the area of this anomaly is much smaller than that present in the known map (see figure 7 and its accompanying quantitative analysis). By using the straight ray forward model in our inversion, we do not account for the possibility that the wave can find a faster path than the direct route. Therefore, the algorithm compensates for this by reducing the size of the low wave speed anomaly to expedite the waves progress along the straight ray paths and reduce the residuals between the observed data and the model. In figure 5(e), the mean of the posterior distribution for the case in which raybending is implemented in the forward model is plotted. By allowing the forward model to explore wave paths other than the straight line between the transmitter and receiver, we can much better reconstruct the full extent of the low wave speed anomaly.

Figures 5(d) and (f) plot the standard deviation at each pixel in the reconstruction for their associated reconstructions shown in (c) and (e) respectively. The uncertainty loop that is present in the centre suggests that the problem is poorly constrained on the boundary of the anomaly; the pixels here move between the two regions. Such uncertainty loops have been conjectured to exist in almost all nonlinear tomography problems and have been demonstrated in both Love wave travel-time tomography [36] and in the diffusive regime in electrical resistivity tomography [54]. This phenomena is examined in more detail in figure 6, where the histograms depicting the posterior distributions of the wave speed at three points in the spatial domain have been plotted (these values are recorded at every 100th iteration). Plots (a) and (b) show the posterior



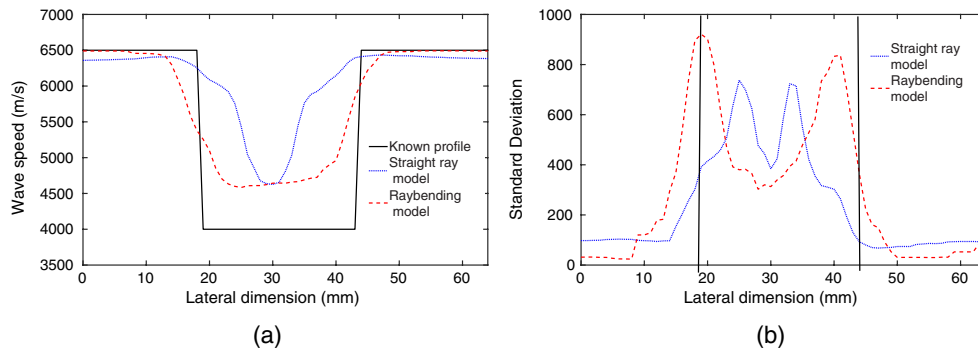
**Figure 5.** Material map reconstructions for a simplistic synthetic geometry where a low wave speed disc of diameter 25 mm is embedded in a high wave speed medium. Plot (a) shows the true geometry. Plot (b) is the initial, randomly assigned, material parametrisation. Plots (c) and (d) show the mean and standard deviation of the posterior distribution at each point in space when the waves are assumed to travel in straight lines between the transmitter and receiver. Plots (e) and (f) show the corresponding mean and standard deviation when raybending due to refraction is included in the forward model. The data used was through-transmission travel-times between a transmitting array spanning the top of the domain shown and a recording array spanning the base of the domain, each array having 32 evenly spaced individual elements.



**Figure 6.** Histograms depicting the posterior distributions of the wave speed (recorded at every 100th sample) at three points in the spatial domain: (a) at  $x = 32$  mm,  $y = 25$  mm (inside the disc), (b) at  $x = 60$  mm,  $y = 42$  mm, (in the host medium) and (c) at  $x = 22$  mm,  $y = 26$  mm (on the uncertainty loop).

distribution for a point lying in the disc and in the host medium respectively. The larger standard deviation present in plot (a) (measured at  $x = 32$  mm,  $y = 25$  mm) is reflected in the central region of the uncertainty map shown in figure 5(f) and it can be observed that the algorithm struggles to correctly reconstruct the value of the disc in figure 5(e). The lower uncertainty at the point located in the host medium (measured at  $x = 60$  mm,  $y = 42$  mm) is reflected in the much narrower probability density approximated in plot (b). More interestingly, when the posterior distribution for the wave speed at a point lying on the uncertainty loop is examined (for example at  $x = 22$  mm,  $y = 26$  mm), a bimodal distribution can be observed (as in figure 6(c)). It is clear that the algorithm sometimes considers this point to lie within the disc and sometimes considers it to lie outside of the disc.

By plotting a cross section of the material map through the centre of the disc, we compare the wave speed profiles obtained by both forward models against that of the known map. From figure 7(a), we can see that by implementing raybending (dashed line), we better capture the true diameter of the low wave speed anomaly (solid line) than when we use the straight ray assumption (dotted line). Quantitatively, if we subtract each of the reconstructed wave speed profiles from the known profile, we obtain root mean square errors (RMSE) of  $1054 \text{ m s}^{-1}$  when the straight ray assumption is employed and  $490 \text{ m s}^{-1}$  when the MSFM is implemented (errors of 42% and 20% respectively, relative to the  $2500 \text{ m s}^{-1}$  range of wave speeds present in the material map). On examination of the standard deviation across the reconstructed map plotted in figure 5(f), we observe another uncertainty loop on the boundary of the disc. A cross section of this is plotted (dashed line) alongside a cross section of figure 5(d) (dotted line) in figure 7(b). Note that we observe increased uncertainty on the boundary of the disc (the location of which are marked by the solid lines) when raybending is included in the forward model: since we have better reconstructed the difference in wave speed between the anomaly and the host medium, the boundary pixels experience a greater variation in values. The size of our low wave speed anomaly can be effectively measured by taking the distance between these two peaks. In the straight ray case we obtain a horizontal diameter measurement of 9 mm, giving rise to a relative error of 64%. By implementing the MSFM as the forward



**Figure 7.** Cross sections of (a) the reconstructed wave speed maps where the solid line represents the known wave speed profile, the dotted line arises from the straight ray assumption and the dashed line arises from the inclusion of raybending, and (b) the uncertainty maps where the dotted line arises from the straight ray assumption and the dashed line arises from the inclusion of raybending.

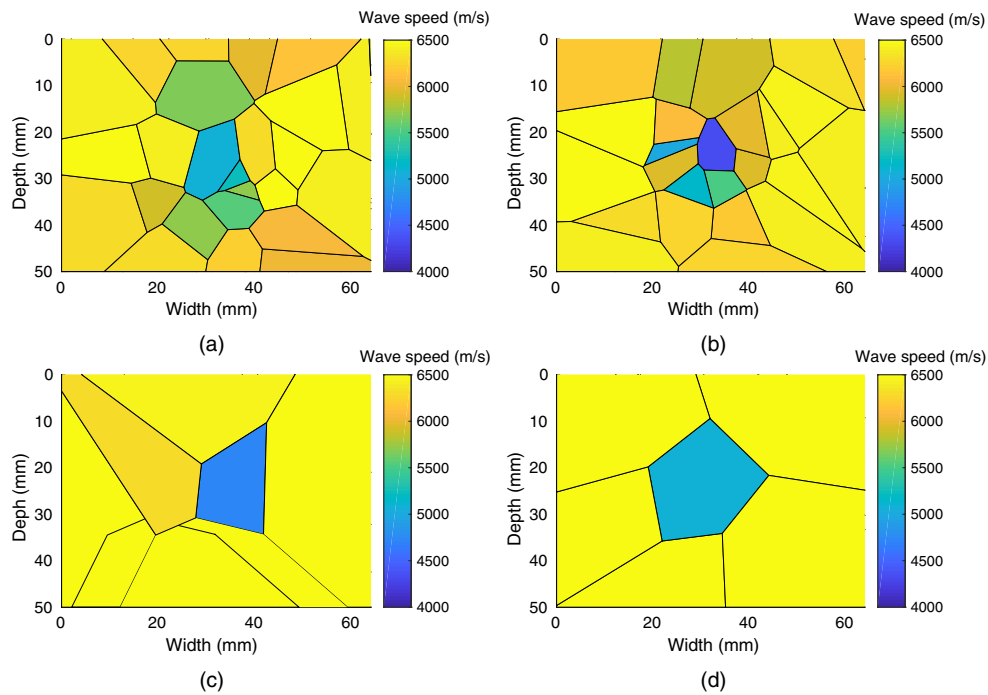
model we obtain an improved measurement of 22 mm (and a reduced relative error of 12%). The width of the peaks in figure 7(b) can also be viewed as a measure of resolution and could potentially be reduced (and thus the resolution enhanced) by increasing the coverage of the inspection domain (that is by employing more sources and receivers) and using a finer grid when computing the travel-time fields using the MSFM. However, both of these adjustments would significantly increase computational expense.

The 2000th and 20 000th samples of the posterior distribution for both cases (where the straight ray assumption is used and where raybending is modelled) are shown in figure 8. It can be observed that in the straight ray case (plots (a) and (b)), the domain is partitioned into more Voronoi cells. This can be attributed to the fact that the model fails to capture the behaviour of the wave front propagation. More explicitly, because the modelling assumptions are inconsistent with the recorded data, the data may appear to be self-contradictory with respect to the assumed physics. Hence the data should not be expected to be consistent with a simple model, even in the case where a simple structure exists in reality. In such cases, the algorithm partitions the domain into more cells in an attempt to better fit the data. However, when the MSFM is implemented within the forward model, fewer cells are required to describe the geometry (plots (c) and (d)). This demonstrates the natural parsimony of the Bayesian inversion; if two models fit the data equally well, the simpler model (in this case, the model with fewer cells) will be assigned higher probability [55]. Comparing the posterior distribution on the number of cells for both cases (see figure 9), we obtain a far narrower distribution when raybending is included in the model. We observe that the distribution is skewed to the right (due to the lower bound of  $M_{\min} = 4$  placed on the prior for the number of cells), has median 6 and an interquartile range (IQR) of 2. A broader distribution with median 25 and IQR of 6 is observed in the case where the straight ray assumption is made.


#### 4.2. Reconstruction of a random medium

To quantify the potential of the rj-MCMC algorithm for flaw image correction, a simulation of the through-transmission ultrasonic phased array inspection of a locally isotropic random media was run in the software package PZFlex [52]. An arbitrary Voronoi diagram with  $M = 100$  seeds was used to create the random geometry (it was then discretised for use as





**Figure 8.** Samples from the posterior distribution of Voronoi diagrams. Plots (a) and (b) are the 2000th and 20000th samples respectively arising from the case where straight rays are used in the forward model. Plots (c) and (d) are the 2000th and 20000th samples arising from the case where the effects of raybending are included in the forward model.



**Figure 9.** Histograms showing the posterior distributions on the number of Voronoi cells used to partition the spatial domain for (a) the case where straight rays between transmitters and receivers are used to calculate the first times of arrival and (b) where the MSFM has been used to consider the effects of raybending on the first times of arrival.

input for the PZFlex simulation). The geometry featured regions with dimensions commensurate with the size of the wavelength of the incident wave (approximately 3–5 mm) and these were randomly assigned wave speeds ranging between  $4500 \text{ m s}^{-1}$  and  $7000 \text{ m s}^{-1}$ , creating a highly scattering environment (see figure 10(a)). A 4 mm diameter circular void was embedded in the centre of the geometry (imaging of this defect is used as a proxy means of quantifying the successful reconstruction of the material map in section 4.3.1). A 1.5 MHz sinusoidal pulse was used to excite the system which consisted of two 32 element phased arrays placed directly opposite each other on either side of the rectangular component. The arrays had a pitch (element spacing) of 2 mm and the depth of the component was 40 mm, giving rise to an inspection domain with dimensions  $64 \text{ mm} \times 40 \text{ mm}$ . The geometry was input into PZFlex and meshed with square elements with side lengths of approximately  $\lambda/15 \approx 200 \mu\text{m}$ . The resulting FMC dataset was processed to yield a ToF matrix  $T_0$  (similar to those shown in figure 2) and this was taken as the input for the rj-MCMC algorithm.

The rj-MCMC algorithm was run for a single Markov chain, for 100 000 samples, where the first 20 000 samples were discarded and the remaining samples were retained at an interval of  $\kappa = 100$ . From the results shown in section 4.1, we observe that by including the effects of raybending in our forward model, the spatially varying wave speed can be more accurately reconstructed, so from this point forward we will only consider this case. Figure 10(b) shows our reconstructed map (the mean of the posterior distribution) and it can be observed that there is no indication of the defect. This can be attributed to the fact that the properties of the circular void do not fall within the finite range of wave speeds dictated by the prior, and so the algorithm cannot reconstruct it. The standard deviation of the wave speed at each pixel within our inspection domain is plotted in figure 10(c). Note that we observe high levels of uncertainty on borders between high and low wave speed regions, for the same reason we observed our uncertainty loops in section 4.1. There is also a clustering of high uncertainty near the centre of the map which could be a symptom of the algorithm's inability to assign the correct wave speed to the flaw region.

Figure 11 plots two Voronoi diagrams drawn from the posterior distribution. Plot (a) shows the model generated on the 35 000th iteration and plot (b) shows the model used for the final iteration of the model. Due to the lack of distinct features in the randomly generated geometry, and the transdimensional nature of the algorithm, it is difficult to visually compare the two. However, it can be observed that we do in fact have some consistency in regions shown to have low variance in figure 10, particularly around the regions which have reconstructed with high velocities (the yellow regions in figure 10(b)).

To examine a more industrially relevant scenario, the simulation described above was rerun, this time featuring a single array which was responsible for both the transmission and reception of the ultrasonic signals (a pulse-echo inspection). The times of flight were extracted from the dataset by cross correlating the input signal with the received signals and these were then used as input for the rj-MCMC inversion. All priors and proposal distributions for the inversion were identical to those in the through-transmission scenario and the resulting map and its standard deviation are shown in figures 12(a) and (b) respectively. The reason the maps generated by the two data acquisition geometries differ can be attributed to the fact that the through-transmission scenario affords increased coverage of the domain and that larger errors are incurred in the estimation of the time of arrival data in the pulse-echo case. This is also reflected in the level of system noise inferred by the rj-MCMC algorithm. The posterior distributions on the noise parameters are plotted in figure 13, where plot (a) arises from the through-transmission case and plot (b) arises from the pulse-echo case. It can be observed that the mean noise level is significantly higher in the pulse-echo case. Furthermore, from figure 14, it can be observed that the median number of cells used to partition the domain (we

examine the median as the distribution in the pulse-echo case appears to be left-skewed) is lower in the through-transmission case than in the pulse-echo case (110 compared to 149), suggesting that the algorithm struggles more to fit the noisier data of the pulse-echo inspection than that of the through-transmission inspection.

#### 4.3. Flaw imaging using the reconstructed material map

The TFM is the benchmark imaging algorithm within the ultrasonic NDT industry [25]. The technique also arises in other fields where data is collected by an array of sensors: in seismology it is referred to as migration [56] and in the medical field it is known as delay and sum beamforming [57]. The algorithm calculates the distance from each transmit element to each pixel and then from the pixel to each receive element. When assuming a constant wave speed throughout the domain, the time this journey takes can be estimated and related to a point on the relevant A scan. Each pixel is then coloured by summing these amplitudes over the set of transmit/receive pairs and can be written [25]

$$w_{ij} = \left| \sum_{t_x=1}^N \sum_{r_x=1}^N A_{t_x, r_x} \left( \frac{\sqrt{(x_{t_x} - x_i)^2 + y_j^2} + \sqrt{(x_{r_x} - x_i)^2 + y_j^2}}{c_l} \right) \right|, \quad (1)$$

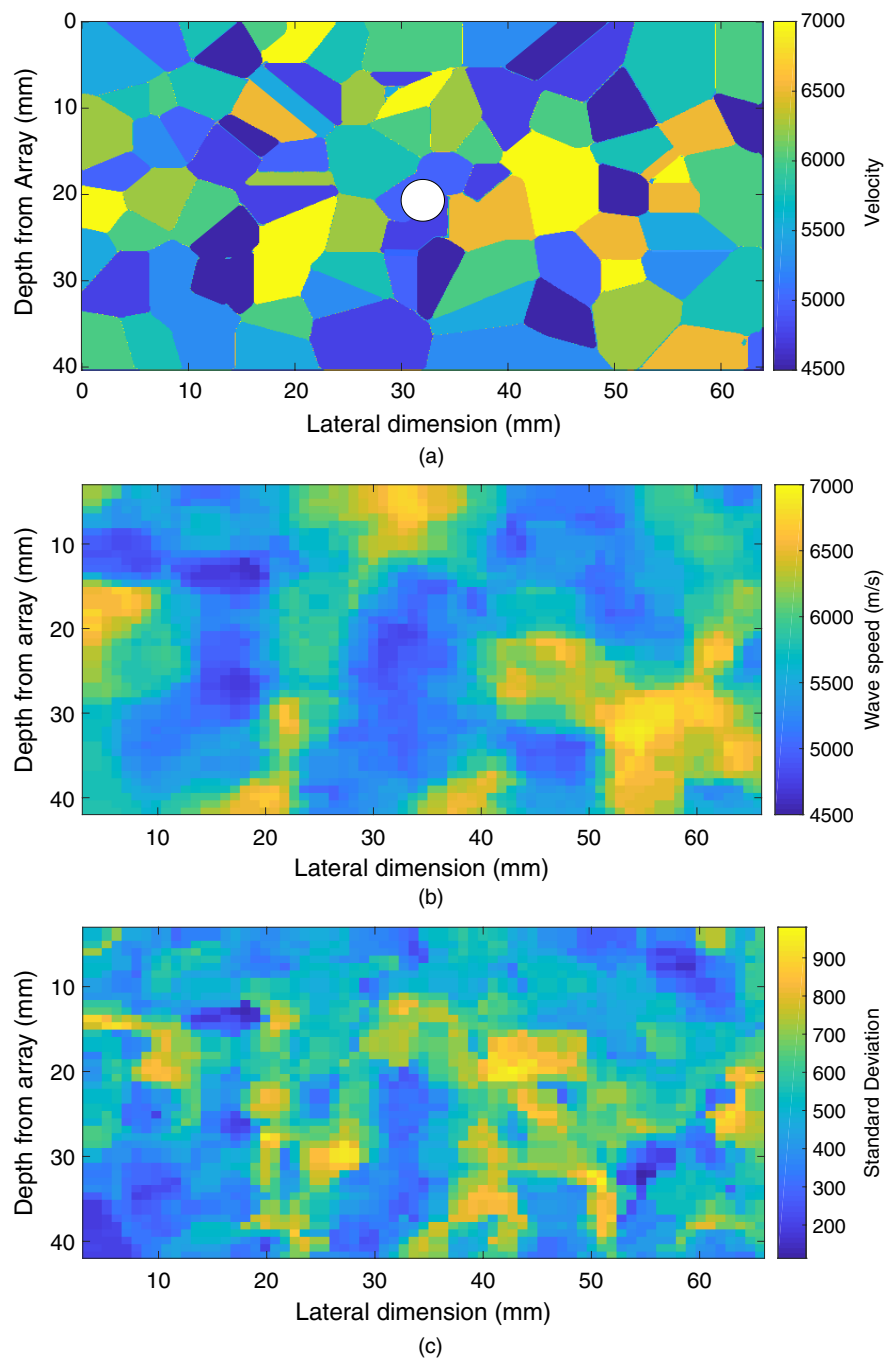
where  $c_l$  is the estimated constant longitudinal wave speed in the host material,  $x_{t_x}$  and  $x_{r_x}$  are the  $x$ -coordinates of the transmit and receive array elements and  $w_{ij}$  is the intensity of the image in the pixel  $(x_i, y_j)$ .

To account for our wave speed map, we employ the MSFM once more. Treating each transmitting element in turn as a source, a set of  $N$  travel-time fields  $\tau^j(x, y)$ ,  $j = 1, \dots, N$ , through our wave speed reconstruction are generated, providing us with the travel-times to each pixel in the imaging domain from each source. Since source-receiver reciprocity holds, these travel-time fields also represent the time taken for the wave to travel from each pixel back to each receiver. To create an image of our inspection domain, the intensity at each pixel can therefore be written

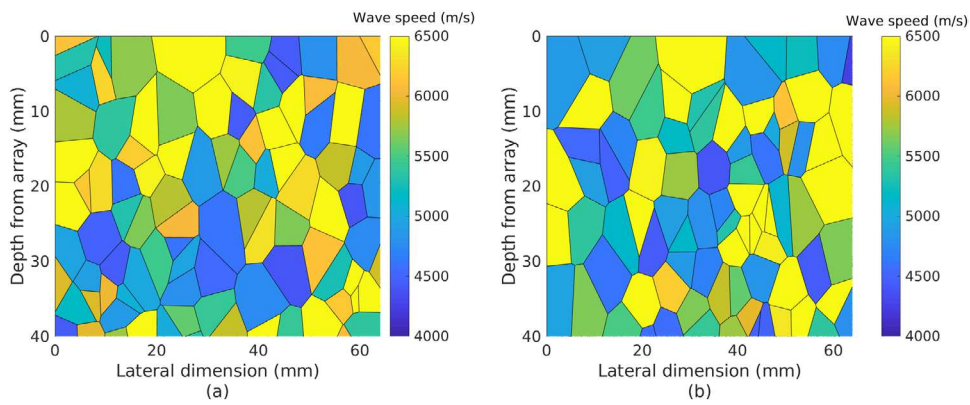
$$w_{ij} = \left| \sum_{t_x=1}^N \sum_{r_x=1}^N A_{t_x, r_x} (\tau^{t_x}(x_i, y_j) + \tau^{r_x}(x_i, y_j)) \right|, \quad (2)$$

where  $\tau^{t_x}(x_i, y_j)$  is the time taken for the wave to travel from source  $t_x$  to the pixel  $(x_i, y_j)$  and, similarly,  $\tau^{r_x}(x_i, y_j)$  is the time taken for the wave to travel from the pixel  $(x_i, y_j)$  to the receiver  $r_x$ . This modified TFM approach, which incorporates the wave speed map and raybending, will be referred to as TFM+ throughout the remainder of this paper.

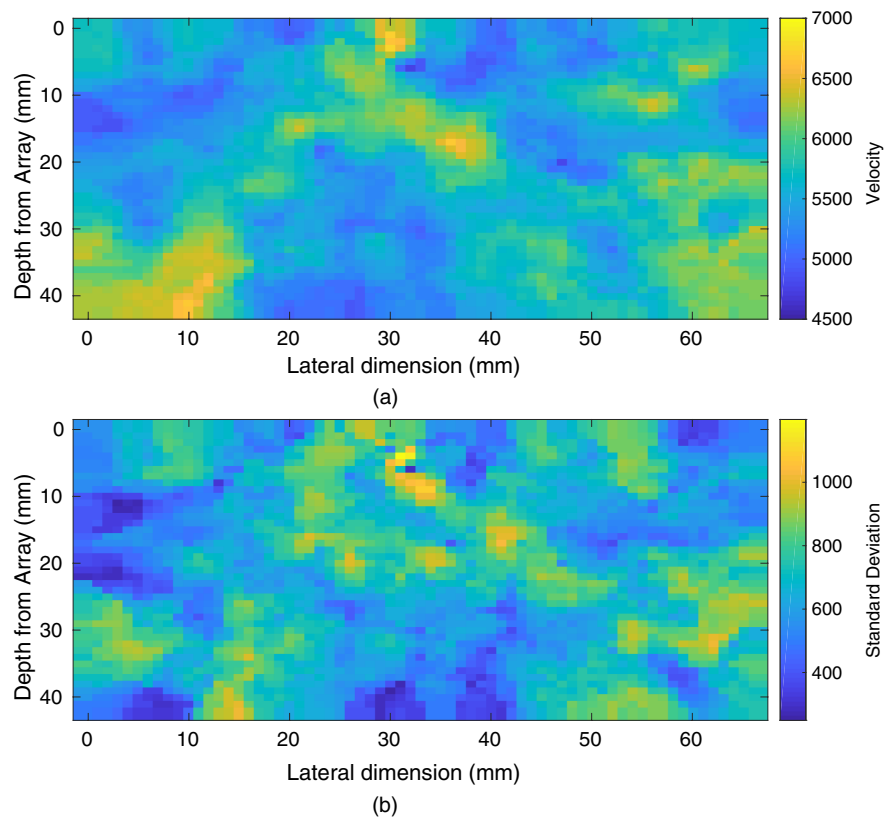
**4.3.1. Flaw reconstruction results.** The success of the wave speed maps shown in section 4.2 is now measured via their use in conjunction with the TFM+ algorithm as developed in equation (2). The RMS longitudinal wave speed through this heterogeneous medium was estimated as  $5081 \text{ m s}^{-1}$  with a standard deviation of  $474 \text{ m s}^{-1}$ . This was calculated using the times corresponding to the back wall echo in the A-scans where transmission and reception took place on the same element. This high standard deviation provides some measure of the extreme heterogeneity of the medium. The flaw reconstruction displayed in figure 15(a) arises from application of the standard TFM (see equation (1)) where the RMS longitudinal wave speed  $c_l = 5081 \text{ m s}^{-1}$  has been assumed over the entire imaging domain. Image (b) shows a close up of the region centred on the known location of the flaw (the exact size and location is represented by the black circle in both images which are plotted on a decibel scale and




**Figure 10.** Material map reconstructions for a random medium generated by data arising from a simulated through-transmission inspection. Plot (a) shows the known material geometry based on a Voronoi tessellation as input to the finite element simulation. The white disc represents the embedded flaw which is assigned void properties. Plot (b) shows the mean of the posterior distribution of the reconstructed wave speed map and plot (c) depicts the corresponding standard deviation.




**Figure 11.** Voronoi tessellations drawn from the posterior distribution for the reconstruction of a random media from through-transmission data. Plot (a) shows iteration 35 000 and plot (b) shows iteration 100 000.



**Figure 12.** Material map reconstructions for a random medium generated by data arising from a simulated pulse-echo inspection. Plot (a) shows the mean of the posterior distribution of the reconstructed wave speed map and plot (b) depicts the corresponding standard deviation.




**Figure 13.** Histograms showing the posterior distributions on the system noise (recorded at every 100th iteration) for (a) the through-transmission inspection case and (b) the pulse-echo inspection case.



**Figure 14.** Histograms showing the posterior distributions on the number of cells for (a) the through-transmission inspection case and (b) the pulse-echo inspection case.


normalised with respect to the highest amplitude). Although some scattering takes place in this region, the energy is dispersed and it is not possible to differentiate this from noise. The maximum intensity in the image occurs at a distance of 10.4 mm from the known centre of the flaw (this can be seen in the bottom right corner of image (b)). As a proof of concept that the material map can enhance our ability to successfully image a flaw, the TFM+ algorithm was applied to the same dataset, incorporating a coarse discretisation (1 mm  $\times$  1 mm grid size) of the true map as shown in figure 10(a). The results are shown in figures 15(c) and (d) and it can be observed that the images display an increased intensity within the region occupied by the flaw. The distance between the location of maximum intensity and the known centre of the flaw is now only 0.8 mm, a ten-fold improvement on that achieved by the standard TFM.

Now that it has been shown that an accurate map of the underlying material properties of a component can be used in conjunction with the TFM+ algorithm as described in equation (2) to obtain improved images, we apply the TFM+ using our reconstructed material maps. Firstly, we use the map obtained from through-transmission ultrasonic measurements



**Figure 15.** TFM reconstructions of the 4mm void embedded in the random, heterogeneous medium as depicted in figure 10(a). Image (a) arises from application of the standard TFM algorithm with a constant wave speed  $c_l = 5081 \text{ m s}^{-1}$  assumed throughout the domain. Image (b) is a close up of a  $20 \text{ mm}^2$  region centred on the known location of the flaw. Image (c) depicts the results when the true material map is used in conjunction with the TFM+ imaging algorithm and image (d) is a close up of the same  $20 \text{ mm}^2$  region centred on the known location of the flaw. In each image, the black circle depicts the actual size, shape and location of the defect. Each plot was normalised with respect to its highest amplitude and plotted on a decibel scale.


and the result is shown in figure 16(a) (a cropped region centred on the flaw is depicted in image (b)). Notice in image (a) that we have better concentration of energy near the known location of the defect than that observed in figure 15(a). If we take the highest amplitude point in image (b) as an indication of the flaw location, this is 1.2mm away from the known centre of the flaw, which is a substantial improvement on the 10.4mm error obtained when the standard TFM is employed. The map reconstructed from the data arising from the more industrially relevant, pulse-echo scenario was also exploited by the TFM+ and the results are shown in figures 16(c) and (d). The image exhibits more noise than that obtained using the through-transmission data, but still results in a six-fold improvement over the basic TFM in terms of flaw location, exhibiting an error of only 1.8mm. However, it must be noted that although computation times are not quoted in this paper (our implementation of the rj-MCMC could benefit from extensive optimisation and parallelisation and so a comparison to real-time TFM [58] would be difficult), it is clear that the improvement in flaw characterisation and sizing comes at significant computational cost. The method will be of most value therefore when imaging flaws in difficult materials where TFM imaging based on a homogeneous wave speed assumption fails. The method will also be of interest to those interested in recovering the interior texture of a heterogeneous material.



**Figure 16.** TFM+ reconstructions of the 4 mm void embedded in the random, heterogeneous medium as depicted in figure 10(a). Image (a) depicts the results when the map of the spatially varying wave speed is reconstructed from data arising from a through-transmission inspection is used in conjunction with the TFM+ imaging algorithm. Image (b) is a close up of the 20 mm<sup>2</sup> region centred on the known location of the flaw. Image (c) depicts the results when the wave speed map is reconstructed from data collected by a pulse-echo inspection is used in conjunction with the TFM+ imaging algorithm. Image (d) is a close up of the same 20 mm<sup>2</sup> region centred on the known location of the flaw. In each image, the black circle depicts the actual size, shape and location of the defect. Each plot was normalised with respect to its highest amplitude and plotted on a decibel scale.

**4.3.2. ROC curve analysis.** A ROC curve plots the performance of a binary classifier system as its discrimination threshold is varied [32, 59]. Previously, within the NDT community, ROC curves have been constructed by running numerous experiments and/or simulations of different inspection scenarios, featuring assorted defects at different locations throughout the component and evaluating the probability of detection for each case [60]. Here, we instead develop a methodology to address the subjectivity involved in the analysis of the image of a single flaw, and present a quantitative technique to compare two imaging algorithms' performance on a single dataset. To formulate ROC curves which represent the success of our TFM and TFM+ images, we develop a binary classifier system which determines whether an image segment lies within the flaw domain. We first use the known material map (as input into the finite element simulation) in conjunction with the TFM+ algorithm to generate a reliable image of the flaw (see figure 15(c)). This image is then partitioned by a regular grid with side-length  $\delta = 1$  mm. The grid cells which contain pixels above the  $-5$  dB threshold are used to define the flaw domain  $\Omega_F$  (at  $-5$  dB, there are no other artefacts in the image). The number of grid cells assigned to the flaw domain is recorded as  $n^f$  and the number of remaining grid cells is referred to as  $n^f = n^t - n^p$ , where  $n^t$  is the total number of grid cells. These quantities are used to normalise the equivalent calculations when we compare the two images arising





**Figure 17.** ROC curves for the images arising from the standard TFM and from TFM+ where the map is recovered from pulse-echo (PE) ultrasonic measurements and through-transmission (TT) ultrasonic measurements.

**Table 1.** Area under the curve values of ROC curves shown in figure 17, using the same key.

Method	AUC
TFM	0.538496
TFM+ (TT)	0.703419
TFM+ (PE)	0.650705

from the TFM and the TFM+ algorithms. The ROC curve is produced by implementing the following steps:

- (1) The images are partitioned using the same grid used to define  $\Omega_F$ .
- (2) The number of grid cells which lie above a threshold  $D$  and lie within  $\Omega_F$  is denoted  $n^+$ . The number of grid cells which meet this threshold but do not lie within  $\Omega_F$  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.

Here, we choose an initial threshold,  $D = -1$  dB, and calculate the probability of detection and false positive rates at intervals down to  $-40$  dB. The ROC curves are plotted in figure 17 for images arising from the standard TFM and the TFM+ where the reconstructed maps are generated by both through-transmission and pulse-echo ultrasonic inspections. Perfect classification performance is achieved when the ROC curve passes through the point (0,1), indicating 100% probability of detection with a false positive rate of 0. This can be quantified using the area under the curve (AUC) [61], the values of which are recorded in table 1. It can be seen

from figure 17 and the respective AUC values that by using the TFM+ , we have an improved ability to correctly identify the flaw. Using the map arising from the through-transmission inspection (denoted by TT) in conjunction with the TFM+ exhibits the highest probability of detection to false positive ratio. The map arising from the pulse-echo inspection (denoted by PE) also provides a significant improvement on the image constructed by the standard TFM. However, as the decibel level is lowered, this image is subjected to more low level noise than both the standard TFM and TFM+ arising from the through-transmission data and this is reflected in its sudden increase in false positive rate (shown by the plateau around PoD = 0.6). As discussed in section 4.2, it is likely that the map arising from the pulse-echo inspection is not as accurate as that arising from the through-transmission data capture as the pulse-echo data acquisition geometry provides sparser coverage of the domain and the extraction of the ToF data is subject to larger errors.

## 5. Conclusions

A stochastic, transdimensional approach for flaw imaging within the ultrasonic NDT industry has been proposed. By parametrising the heterogeneous medium by Voronoi diagrams and using the reversible jump MCMC method, locally isotropic inhomogeneous wave speed maps were reconstructed from simulated ultrasonic phased array data. It was shown that by using the MSFM to model the monotonically advancing wave front, the dependency of the data uncertainty on ray length could be reduced. This is advantageous as it negates the requirement to retrospectively raytrace through the calculated travel-time field which can become prohibitively computationally expensive. To demonstrate the reliability of the reconstructed maps, simple geometries were considered. Visually, these maps appeared to reconstruct key facets of the medium's structure and, due to the dense and regular coverage provided by the data acquisition geometry, the MCMC algorithm converged to the posterior distribution relatively quickly. The methodology was then applied to a synthetic dataset arising from a randomly partitioned medium, exhibiting wavelength sized regions of highly contrasting wave speeds (between  $4500 \text{ m s}^{-1}$  and  $7000 \text{ m s}^{-1}$ ) and a centrally embedded 4 mm diameter void (also commensurate with the wavelength). It is more difficult to visually assess the similarity between the known and reconstructed maps in this case (where the medium exhibits no distinct patterns or features) and so the success of the material map's reconstruction was measured by using it in conjunction with a flaw imaging algorithm and taking the improvement in flaw reconstruction as a proxy measurement. The NDT community's benchmark imaging algorithm, the TFM, was modified to consider the reconstructed maps, accounting for the varying wave speed throughout the component and the subsequent refraction induced by the structure (this method is referred to as TFM+). By correcting for the inhomogeneous wave speed map, a six-fold improvement in flaw location was achieved. To ensure that the comparison of imaging algorithms was not biased by the subjective choice of image thresholding, a methodology to calculate ROC curves for each image was developed and it was thus shown that by using the TFM+ , the probability of detection rates were enhanced in both the through-transmission and pulse-echo modes of operation.

## Funding statement

This work was funded through the UK Research Centre in NDE by the Engineering and Physical Sciences Research Council (grant number EP/L022125/1) and the iNEED project (grant number EP/P005268/1). The authors would also like to thank the Edinburgh

Interferometry Project (EIP) sponsors (ConocoPhillips, Schlumberger Cambridge Research, Statoil and Total) for supporting this research.

## ORCID iDs

K M M Tant  <https://orcid.org/0000-0003-4345-7054>

## References

- [1] Drinkwater B W and Wilcox P D 2006 Ultrasonic arrays for non-destructive evaluation: a review *NDT & E Int.* **39** 525–41
- [2] Blitz J and Simpson G 1995 *Ultrasonic Methods of Non-Destructive Testing* vol 2 (Netherlands: Springer)
- [3] Smith N B and Webb A 2010 *Introduction to Medical Imaging: Physics, Engineering and Clinical Applications* (Cambridge: Cambridge University Press)
- [4] Kemnitz P, Richter U and Klüber H 1997 Measurements of the acoustic field on austenitic welds: a way to higher reliability in ultrasonic tests *Nucl. Eng. Des.* **174** 259–72
- [5] Chassignole B, Duwig V, Ploix M-A, Guy P and El Guerjouma R 2009 Modelling the attenuation in the ATHENA finite elements code for the ultrasonic testing of austenitic stainless steel welds *Ultrasonics* **49** 653–8
- [6] Nageswaran C, Carpentier C and Tse Y 2009 Microstructural quantification, modelling and array ultrasonics to improve the inspection of austenitic welds *Insight* **51** 660–6
- [7] Zhang J, Hunter A, Drinkwater B W and Wilcox P D 2012 Monte Carlo inversion of ultrasonic array data to map anisotropic weld properties *IEEE TUFFC* **59** 2487–97
- [8] Abrahams I D and Wickham G R 1992 The propagation of elastic waves in a certain class of inhomogeneous anisotropic materials. I. The refraction of a horizontally polarized shear wave source *PRSA* **436** 449–78
- [9] Connolly G D, Lowe M J S, Temple J A G and Rokhlin S I 2010 Correction of ultrasonic array images to improve reflector sizing and location in inhomogeneous materials using a ray-tracing model *J. Acoust. Soc. Am.* **127** 2802–12
- [10] Spies M 2004 Analytical methods for modeling of ultrasonic nondestructive testing of anisotropic media *Ultrasonics* **42** 213–9
- [11] Harvey G, Tweedie A, Carpentier C and Reynolds P 2011 Finite element analysis of ultrasonic phased array inspections on anisotropic welds *AIP Conf. Proc.* **1335** 827–34
- [12] Lane C 2013 *The Development of a 2D Ultrasonic Array Inspection for Single Crystal Turbine Blades* (Berlin: Springer)
- [13] Sharples S D, Clark M and Somekh M G 2006 Spatially resolved acoustic spectroscopy for fast noncontact imaging of material microstructure *Opt. Express* **14** 10435–40
- [14] Moysan J, Apfel A, Corneloup G and Chassignole B 2003 Modelling the grain orientation of austenitic stainless steel multipass welds to improve ultrasonic assessment of structural integrity *Int. J. Press. Vessels Pip.* **80** 77–85
- [15] Mark A F, Fan Z, Azough F, Lowe M J S and Withers P J 2014 Investigation of the elastic/crystallographic anisotropy of welds for improved ultrasonic inspections *Mater. Charact.* **98** 47–53
- [16] Fan Z, Mark A F, Lowe M J S and Withers P J 2015 Nonintrusive estimation of anisotropic stiffness maps of heterogeneous steel welds for the improvement of ultrasonic array inspection *IEEE TUFFC* **62** 1530–43
- [17] Gueudré C, Le Marrec L, Moysan J and Chassignole B 2009 Direct model optimisation for data inversion. Application to ultrasonic characterisation of heterogeneous welds *NDT & E Int.* **42** 47–55
- [18] Grangeat P 2013 *Tomography* (New York: Wiley)
- [19] Rose J L, Jiao D and Spanner J Jr 1996 Ultrasonic guided wave NDE for piping *Mater. Eval.* **54** 1310–3

- [20] Leonard K R and Hinders M K 2005 Lamb wave tomography of pipe-like structures *Ultrasonics* **43** 574–83
- [21] Nagy P B, Simonetti F and Instanes G 2014 Corrosion and erosion monitoring in plates and pipes using constant group velocity Lamb wave inspection *Ultrasonics* **54** 1832–41
- [22] Leonard K R, Malyarenko E V and Hinders M K 2002 Ultrasonic Lamb wave tomography *Inverse Problems* **18** 1795–808
- [23] Huthwaite P and Simonetti F 2011 High-resolution imaging without iteration: a fast and robust method for breast ultrasound tomography *J. Acoust. Soc. Am.* **130** 1721–34
- [24] Ogilvy J A 1986 Ultrasonic beam profiles and beam propagation in an austenitic weld using a theoretical ray tracing model *Ultrasonics* **24** 337–47
- [25] Holmes C, Drinkwater B W and Wilcox P D 2005 Post processing of the full matrix of ultrasonic transmit receive array data for non destructive evaluation *NDT & E Int.* **38** 701–11
- [26] Green P J 1995 Reversible jump Markov chain Monte Carlo computation and Bayesian model determination *Biometrika* **82** 711–32
- [27] Bodin T and Sambridge M 2009 Seismic tomography with the reversible jump algorithm *Geophys. J. Int.* **178** 1411–36
- [28] Bodin T, Sambridge M and Gallagher K 2009 A self-parametrizing partition model approach to tomographic inverse problems *Inverse Problems* **25** 055009
- [29] Galetti E, Curtis A, Baptie B, Jenkins D and Nicolson H 2016 Transdimensional Love-wave tomography of the British Isles and shear-velocity structure of the East Irish Sea Basin from ambient-noise interferometry *Geophys. J. Int.* **208** 36–58
- [30] Okabe A, Boots B, Sugihara K and Chiu S N 2009 *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams* (New York: Wiley)
- [31] Hassouna M S and Farag A A 2007 Multistencils fast marching methods: a highly accurate solution to the eikonal equation on cartesian domains *IEEE TPAMI* **29** 1563–74
- [32] Fawcett T 2006 An introduction to ROC analysis *Pattern Recognit. Lett.* **27** 861–74
- [33] Schmerr L 1998 *Fundamentals of Ultrasonic Nondestructive Evaluation a Modeling Approach* (New York: Plenum)
- [34] Marioli D, Narduzzi C, Offelli C, Petri D, Sardini E and Taroni A 1992 Digital time-of-flight measurement for ultrasonic sensors *IEEE TIM* **41** 93–7
- [35] Van Pamel A, Brett C R, Huthwaite P and Lowe M J S 2015 Finite element modelling of elastic wave scattering within a polycrystalline material in two and three dimensions *J. Acoust. Soc. Am.* **138** 2326–36
- [36] Galetti E, Curtis A, Meles G A and Baptie B 2015 Uncertainty loops in travel-time tomography from nonlinear wave physics *Phys. Rev. Lett.* **114** 148501
- [37] Sethian J A 1999 *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision and Materials Science* vol 3 (Cambridge: Cambridge University Press)
- [38] Popovici A M and Sethian J A 2002 3D imaging using higher order fast marching traveltimes *Geophysics* **67** 604–9
- [39] Sethian J A 1999 Fast marching methods *SIAM Rev.* **41** 199–235
- [40] Sethian J A and Popovici A M 1999 3D traveltimes computation using the fast marching method *Geophysics* **64** 516–23
- [41] Rawlinson N and Sambridge M 2004 Wave front evolution in strongly heterogeneous layered media using the fast marching method *Geophys. J. Int.* **156** 631–47
- [42] Bodin T, Sambridge M, Rawlinson N and Arroucau P 2012 Transdimensional tomography with unknown data noise *Geophys. J. Int.* **189** 1536–56
- [43] Gelman A, Carlin J B, Stern H S, Dunson D B, Vehtari A and Rubin D B 2014 *Bayesian Data Analysis* vol 2 (Boca Raton, FL: CRC press)
- [44] Al-Awadhi F, Hurn M and Jennison C 2004 Improving the acceptance rate of reversible jump MCMC proposals *Stat. Probab. Lett.* **69** 189–98
- [45] Brooks S P, Giudici P and Roberts G O 2003 Efficient construction of reversible jump Markov chain Monte Carlo proposal distributions *J. R. Stat. Soc. Series* **65** 3–39
- [46] Rosenthal J S 2011 Optimal proposal distributions and adaptive MCMC *Handbook Markov Chain Monte Carlo* **4** 92–113
- [47] Gelman A, Roberts G O and Gilks W R 1996 Efficient metropolis jumping rules *Bayesian Stat.* **5** 599–608

- [48] Tierney L and Mira A 1999 Some adaptive Monte Carlo methods for Bayesian inference *Stat. Med.* **18** 2507–15
- [49] Green P J and Mira A 2001 Delayed rejection in reversible jump Metropolis–Hastings *Biometrika* **88** 1035–53
- [50] Cowles M K and Carlin B P 1996 Markov chain Monte Carlo convergence diagnostics: a comparative review *J. Am. Stat. Assoc.* **91** 883–904
- [51] Aster R C, Borchers B and Thurber C H 2013 *Parameter Estimation and Inverse Problems* 2nd edn (Oxford: Elsevier)
- [52] PZFlex. Thornton Tomasetti Defence Ltd. 6th Floor South, 39 St Vincent Place, Glasgow, Scotland, G1 2ER, United Kingdom
- [53] Dobson J, Tweedie A, Harvey G, OLeary R, Mulholland A, Tant K and Gachagan A 2016 Finite element analysis simulations for ultrasonic array NDE inspections *AIP Conf. Proc.* **1706** 040005
- [54] Galetti E and Curtis A 2018 Transdimensional electrical resistivity tomography *J. Geophys. Res.: Solid Earth* (<https://doi.org/10.1029/2017JB015418>)
- [55] Malinverno A 2002 Parsimonious Bayesian Markov chain Monte Carlo inversion in a nonlinear geophysical problem *Geophys. J. Int.* **151** 675–88
- [56] Miller D, Oristaglio M and Beylkin G 1987 A new slant on seismic imaging: migration and integral geometry *Geophysics* **52** 943–64
- [57] Wang Z, Li J and Wu R 2005 Time-delay-and time-reversal-based robust capon beamformers for ultrasound imaging *IEEE TMI* **24** 1308–22
- [58] Dziewierz J, Gachagan A, Lord N and Mulholland A 2012 An application-specific design approach for 2D ultrasonic arrays *51st Annual Conf. of the British Institute of Non-Destructive Testing 2012 (NDT 2012)*
- [59] Metz C E 1978 Basic principles of ROC analysis *Seminars in Nuclear Medicine* vol **8** (Amsterdam: Elsevier) pp 283–98
- [60] Jarvis R, Cawley P and Nagy P B 2017 Performance evaluation of a magnetic field measurement NDE technique using a model assisted probability of detection framework *NDT & E Int.* **91** 61–70
- [61] Bradley A P 1997 The use of the area under the ROC curve in the evaluation of machine learning algorithms *Pattern Recognit.* **30** 1145–59