Ludlam, J., Tant, K. M. M., Curtis, A., & Dolean, V. (2025). Ultrasonic travel-time tomography for approximating the local elastic tensor in complex media. In 2025 IEEE International Ultrasonics Symposium (IUS) (pp. 1-5). (2025 IEEE International Ultrasonics Symposium (IUS)). IEEE. https://doi.org/10.1109/ius62464.2025.11201431

Ultrasonic Travel-Time Tomography for Approximating the Local Elastic Tensor in Complex Media

James Ludlam
Department of Mathematics and Statistics
University of Strathclyde
Glasgow, UK
james.ludlam@strath.ac.uk

Katherine M. M. Tant
James Watt School of Engineering
University of Glasgow
Glasgow, UK
katy.tant@glasgow.ac.uk

Andrew Curtis
School of Geoscience
University of Edinburgh
Edinburgh, UK
andrew.curtis@ed.ac.uk

Victorita Dolean
Mathematics and Computer Science
TU Eindhoven
Eindhoven, The Netherlands
v.dolean.maini@tue.nl

***Abstract*—Ultrasonic non-destructive evaluation (UNDE) is vital for assessing the structural integrity of safety-critical infrastructure. However, accurate defect detection and characterisation using UNDE is particularly challenging in complex materials like austenitic steel welds, as the heterogeneous and locally anisotropic grain structures distort wave paths, causing traditional imaging methods based on homogeneous and isotropic assumptions to fail. We present a probabilistic framework to reconstruct spatially varying elastic tensor information from ultrasonic travel-time data using stochastic Stein Variational Gradient Descent. Unlike prior approaches, our method relaxes assumptions of uniformity of material properties across the domain and considers uncertainty induced by limited prior knowledge. We show that travel-time data alone cannot fully constrain high-dimensional domains, and accurate imaging requires informed priors on a series of anisotropy parameters we use as a proxy for the stiffness tensor - scale, strength, and orientation.**

***Index Terms*—ultrasound, tomography, materials characterisation, inverse problems, imaging, non-destructive evaluation**

I. Introduction

Ultrasonic non-destructive evaluation (UNDE), uses ultrasound waves to detect internal damage of industrial infrastructure without compromising the integrity of the components. Linear arrays of ultrasonic transducers are used to transmit mechanical waves through the object of interest. The resulting scattered wave field is simultaneously recorded and can be analysed to build up images of the interior of these parts [1].

However, inspection becomes significantly more complex when materials are inhomogeneous or anisotropic, as expected wave paths and arrival times become distorted [2]. A prime example is austenitic steel welds, widely used for their strength and corrosion resistance. During manufacturing, these welds develop locally anisotropic microstructures due to epitaxial crystal growth, which causes traditional imaging algorithms, that typically assume uniform wave velocity throughout the domain, to fail [3].

To correctly focus scattered wave energy in the ultrasonic imaging domain, we require good knowledge of the underlying spatial distribution of material properties in the object of interest, as these can impact the speed and direction of the propagating waves. Travel-time tomography methods can be used to invert the fastest time of arrivals between pairs of transmitters and sensors to construct a map of some spatially varying material property [4].

Within the field of UNDE, work to date has primarily focussed on constructing grain orientation maps, under the assumption that the elastic tensor is known and constant throughout the domain [5]–[8]. In this work, we relax this assumption and tackle the problem more generally to facilitate application across the NDE, medical and seismic imaging domains. Given the usually high-dimensional and non-linear nature of these problems, much of the related literature has focused on driving these tomography approaches with Markov Chain Monte Carlo (MCMC) methods, which are of course computationally expensive [5], [6], [8]. However, variational Bayesian Inversion approaches offer an efficient alternative [9]. In this work we apply the stochastic Stein Variational Gradient Descent [10] to invert for some parameterisation of the spatially varying elastic tensor in complex media from ultrasonic travel-time data. We show that the travel-time data itself is not enough to fully constrain this problem in spatial domains which require high dimensional parameterisations, and sufficient prior knowledge on one of the three anisotropy parameters we introduce (scale, anisotropy strength and orientation) is required.

II. Method

In a variational Bayesian framework, the travel-time tomography problem can be broken down into the following elements:

1) The observed data and prior beliefs.
2) The forward model and domain parameterisation.
3) Likelihood calculations and posterior inference.

In our case, the observed data arises from the measured time of arrival between pairs of transmit-receive array elements at known locations. We apply uniform prior distributions to our parameters of interest, bounded by limits informed my the physics (see Section II-B). The Anisotropic Locally Interpolated Fast Marching Method (ALI-FMM) [11] has been selected as our forward model to approximate the relationship between the parameters of interest and the travel-time data (Section II-A). To parameterise the inspection medium, we introduce three anisotropy parameters as a proxy for the full elastic stiffness tensor to reduce the dimension of the inverse problem - see Section II-B). Finally in Section II-C, we introduce the stochastic Stein variational gradient descent [10] for our inference step.

A. The Forward Model

The travel-time tomography problem is driven by minimising the misfit between some observed travel-time data with modelled travel-times through an iteratively perturbed parameterisation of the wave propagation medium. In anisotropic media, the speed at which a wave travels is dependent on its direction of incidence. The group velocity curves which govern this relationship can be derived from the Christoffel equation given some stiffness tensor and density [12]. In this work, we model travel-times using the Anisotropic, Locally Interpolated Fast Marching Method (ALI-FMM) [11]. For a given parameterisation of some generally anisotropic and heterogeneous medium, where the group velocity curve is known at each point, this algorithm models the wave-front propagation through the domain and outputs the travel-time matrix for a known set of transmit-receive pairs.

B. Parameterising Anisotropy

In this paper, we restrict our attention to the 2D problem and assume that the materials of interest are cubic with 90° rotational symmetry. In this case, the stiffness tensor has only three independent elastic constants c_{11}, c_{12} and c_{44}. Solving the Christoffel equation [12] for a given set of elastic constants and density ρ yields the required group velocity curve for a given material. We can then assign an orientation θ which represents the primary axis of the material crystal at a given point, and can be visualised as a rotation of the velocity curve. Thus, to construct the group velocity curve at each point in the domain of interest we require knowledge of five parameters ($c_{11}, c_{12}, c_{44}, \rho$ and θ.). However, we can reduce the dimension of this inverse problem by reformulating the problem using three intuitive parameters:

i) Orientation (θ). We restrict attention to in plane orientation (we ignore out of plane tilt; this simplification is often justified in engineering applications such as welding and composites). As before, θ represents the primary axis of the material crystal and thus determines the rotation of the slowness curve at a given point in space.
ii) Strength (A). This parameter determines the strength of the anisotropy present, linearly interpolating between some reference isotropic and strongly anisotropic material.
iii) Scaling (S). The scaling parameter can be thought of the dilation of the velocity curve, where it is enlarged or reduced by some factor, whilst preserving its proportions and angles.

Here, θ is unchanged from our original parameterisation. The relationship between the elastic constants and the scaling and strength parameters can be written

$$c_{jk} = S\left((1-A)c^i_{jk} + Ac^a_{jk}\right). \tag{1}$$

where the superscripts i and a denote use of the stiffness matrix for some reference isotropic and anisotropic materials (discussed below). We now have $3M$ unknowns instead of $5M$, where M is the fixed number of cells in our spatial discretisation of the domain.

We can perturb our three parameters to adjust the fit of our model to the observed travel-times. The values these parameters can take are based on prior knowledge of physically relevant materials. In this paper, we introduce a reference anisotropic curve obtained by solving the Christoffel equation with some given stiffness tensor and density. A reference isotropic velocity curve is then constructed, where the elastic constants are chosen to produce a circular velocity curve with velocity equal to the average velocity of the reference anisotropic curve. The value pair $A = 0, S = 1$ represents the fixed reference isotropic velocity curve and $A = 1, S = 1$ the fixed reference anisotropic velocity curve. When $0 < A < 1$, the strength of the anisotropy lies between the two reference curves and when $A > 1$ the anisotropy is more exaggerated. However, as A continues to increase, the group velocity curve can eventually become concave at some angles (the group angle from phase angle no longer unique), so we enforce an upper limit (in this work set as $A = 1.5$). The scaling parameter can notionally take any positive value, however we restrict attention to the range $S \in [0.2, 5]$ to ensure the materials are physically relevant. Finally, we draw values of θ from a uniform distribution bounded by $0° \le \theta \le 90°$, informed by the rotational symmetry of the velocity curve (although it would be straightforward to include known statistics on the grain orientation distribution in the case these are known).

C. Stochastic Stein Variational Gradient Descent

Stein Variational Gradient Descent (SVGD) [13] is a variational Bayesian inference method that approximates the posterior distribution using a set of particles $\{\mathbf{m}_i\}_{i=1}^n$, where each $\mathbf{m}_i \in \mathbb{R}^d$ represents a sample in the d-dimensional parameter space. These particles are updated deterministically to minimize the Kullback-Leibler (KL) divergence between their empirical distribution and the target posterior $p(\mathbf{m}|\mathbf{d}_{\text{obs}})$ [13].

Stochastic SVGD (sSVGD) [10] extends SVGD by introducing a stochastic term, combining SVGD's gradient-based efficiency with the exploratory behaviour of MCMC. Let $\mathbf{z}_t \in \mathbb{R}^{nd}$ be the stacked vector of all particles at iteration t. The sSVGD update rule is given by

$$\mathbf{z}_{t+1} = \mathbf{z}_t + \epsilon_t\left[\mathbf{K}\nabla_{\mathbf{z}_t}\log p(\mathbf{z}_t) + \nabla_{\mathbf{z}_t}\mathbf{K}\right] + \mathcal{N}(\mathbf{0}, 2\epsilon_t\mathbf{K}), \tag{2}$$

where ϵ_t is the step size at iteration t, $\nabla_{\mathbf{z}_t}\log p(\mathbf{z}_t)$ is the gradient of the log-posterior with respect to all particles, $\nabla_{\mathbf{z}_t}\mathbf{K}$ is the gradient of the kernel matrix, and $\mathcal{N}(\mathbf{0}, 2\epsilon_t\mathbf{K})$ is a multivariate Gaussian noise term with zero mean and covariance $2\epsilon_t\mathbf{K}$. We define the block kernel matrix $\mathbf{K} \in \mathbb{R}^{nd\times nd}$ as

$$\mathbf{K} = \frac{1}{n}\begin{bmatrix} k(\mathbf{m}_1,\mathbf{m}_1)\mathbf{I}_d & \dots & k(\mathbf{m}_1,\mathbf{m}_n)\mathbf{I}_d \\ \vdots & \ddots & \vdots \\ k(\mathbf{m}_n,\mathbf{m}_1)\mathbf{I}_d & \dots & k(\mathbf{m}_n,\mathbf{m}_n)\mathbf{I}_d \end{bmatrix}. \tag{3}$$

Here,

$$k(\mathbf{m}, \mathbf{m}') = \exp\left(-\frac{1}{h}||\mathbf{m} - \mathbf{m}'||^2\right), \tag{4}$$

is the kernel function, where $h = \tilde{d}^2/\log n$ and $\tilde{d}$ is the median of pairwise distances between particles.

The direction which the particle moves is determined by (2). The first term $\mathbf{K}\nabla_{\mathbf{z}_t} \log p(\mathbf{z}_t)$ is dependent on the likelihood function and so pulls the particles towards areas of high probability. The second term $\nabla_{\mathbf{z}_t}\mathbf{K}$ acts like a repulsive force between particles and ensures that the particles sample the probability space instead of converging to local maxima. The strength of these contributions are controlled by the kernel function in (4) where the value of h is chosen so that $\sum_j k(\mathbf{m}_i, \mathbf{m}_j) \approx n\exp\left(-\frac{1}{h}\tilde{d}^2\right) = 1$. This ensures that the contribution from a particle, its gradient and the influence from other particles are balanced. The final stochastic term helps particles escape local minima and improves exploration in high-dimensional spaces. Unlike traditional MCMC methods, sSVGD does not reject samples and maintains asymptotic convergence to the posterior as the number of iterations increases. A burn-in period is typically used to discard early samples before convergence.

III. Results

A. Target Data

To examine our ability to invert for orientation, anisotropy and scaling parameters with limited prior knowledge, we consider the simple layered structure plotted in Fig. 1. Transmit and receive elements are co-located on the boundary of the square domain, with 6 placed on each side at regular intervals. Note that travel-times arising from source-receiver pairs on the same edge are discounted as the first arrivals will correspond to the surface wave which we do not consider in this work. The top and bottom layers are assigned isotropic properties with velocities of $4500\,\mathrm{m\,s^{-1}}$ and $6100\,\mathrm{m\,s^{-1}}$ respectively. The middle layer is anisotropic with elastic constants c_{11} =203.6 GPa, c_{12} =133.5 GPa, c_{44} =129.8 GPa, density $\rho = 7850\,\mathrm{kg\,m^{-3}}$ [14] and orientation 30°. Since the top and bottom layers are isotropic, the orientation parameter θ has no effect on the constructed velocity curve. For the sake of plotting the target geometry we set this to 0° but expect to reconstruct a uniform distribution over orientation in these regions. In the target geometry $A = 0$ in the top and bottom layers and $A = 1$ in the middle layer. The scaling parameter for the top, middle and bottom layers is 0.74 (2dp), 1, and 1.37 (2dp) respectively. In this simulated example, the ALI-FMM was used to model the observed data, where the spatial domain was discretised over a 46×46 grid.

B. Inversion Parameters

To run the sSVGD, we use a fixed 6×6 grid of square cells to parameterise the domain (a much lower resolution grid than that used to generate the target data). The prior probability density functions (pdfs) are as set out in Section II-B. The gradients for the material models are calculated by tracing rays backward from each receiver to each source using the gradient of the travel time field [11], [15]. We initiate 40 particles drawn from the prior pdf and run 1500 iterations of the update step given in (2). The step size ϵ_l is selected using the AdaDelta method [16] implemented with a general learning rate of 2, 1, 0.5 for iterations 0-50, 51-100 and 101-1500, and the first 500 iterations are discarded as the burn-in period.

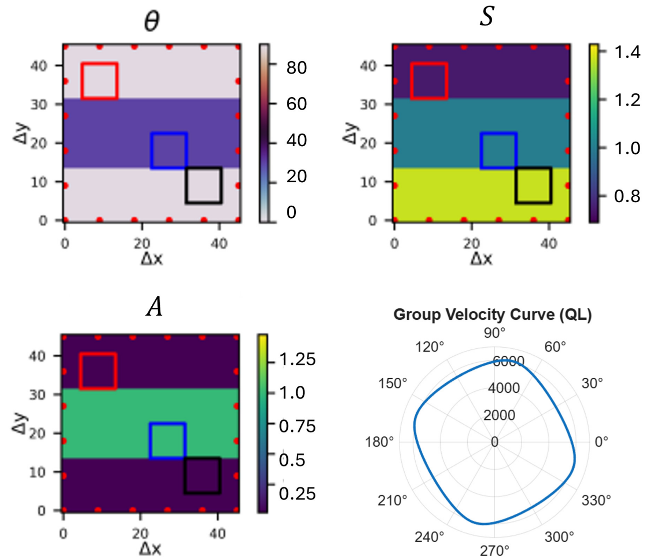

Fig. 1. The true material model plotted for the 3 inversion parameters: orientation θ (top left); scaling (S) (top right); anisotropy strength (A) (bottom left). Co-located source receivers indicated by red crosses. The bottom right image show the true group velocity curve of the central layer (with it 30° rotation applied).

C. Results

Fig. 2 plots the distribution of constructed velocity curves at three distinct points in the domain, marked by the red, blue and black boxes in Fig. 1. The two columns were generated by two separate runs of the sSVGD, with the only difference being in the definition of the prior distributions. The left hand column of Fig. 2 arises from the uninformative priors set out in Section II-B. In the top plot, we have the distribution of curves at a point in the top, isotropic layer (red box) which has a velocity of $4500\,\mathrm{m\,s^{-1}}$. Significant leakage of the anisotropic behaviour of the central layer can be observed and this has the effect on increasing the estimated velocity in this layer in some directions. The middle left hand plot shows the distribution on velocity curves in the central anisotropic layer with 30° orientation (blue box). This reconstruction fails to fully capture the full degree of anisotropy present in the model (the curve representing the mean of the distribution is more rounded thank the true curve). However, the dominant orientation is well estimated by the mean and the scaling factor effectively captures the dilation of the curve. This can be seen more clearly in the left hand column of Fig. 3, where histograms representing the marginal posterior distributions on θ, A and S at this point are plotted. The right hand column of Fig. 2 plots the distributions of velocity curves at

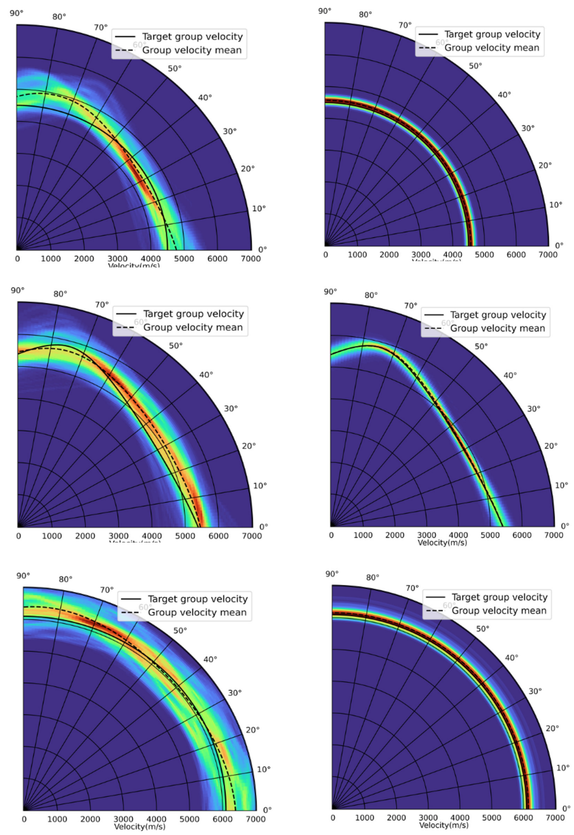

Fig. 2. Reconstructions of the velocity curves in the top, middle and bottom layers as denoted by the red, blue and black boxes in Fig. 1. The left hand column corresponds to reconstructions achieved with uninformative uniform priors on θ, S, A. The right hand column arises from the case when A is assumed known.

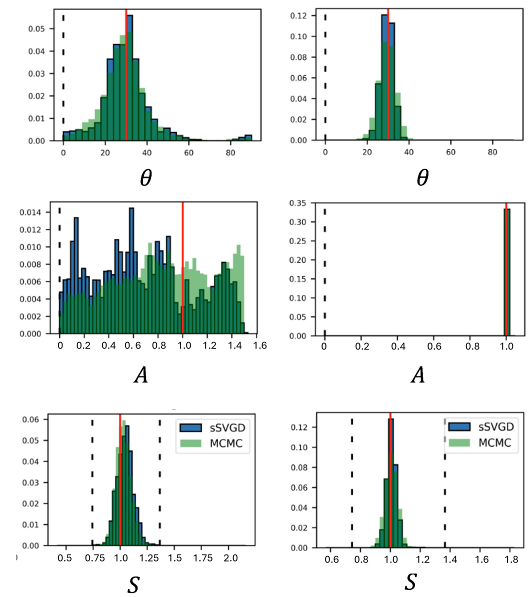

Fig. 3. Histograms approximating the marginal posterior distributions on θ, A and S of in the central layer (blue box in Fig. 1) constructed using the sSVGD (blue bars). The left hand column corresponds to reconstructions achieved with uninformative uniform priors on θ, S, A. The right hand column arises from the case when A is assumed known. The green shaded areas represent the benchmark distributions reconstructed by running a Metropolis Hastings MCMC.

the same three points, however we now assume the level of anisotropy is known and fixed in each layer. Orientation θ and scaling S are allowed to vary as before. It is very clear that this additional prior knowledge allows us to much more successfully constrain the problem and that the mean velocity curve well approximates the true values, with little uncertainty. This is clearly reflected in the right hand column of Fig. 3, where the marginal distributions are plotted for this case.

The green shaded histograms in Fig. 3 arise from running a Metropolis-Hastings MCMC inversion using the same domain parameterisation and priors as the sSVGD to benchmark our results. Thirty Markov chains, each with 55,000 samples, were run. Once convergence was believed to have been achieved, the initial 5000 samples were discarded (the burn in period) and the ensemble of accepted models was resampled at a decimation interval of 100. This MCMC approach required a total of 1,650,000 forward evaluations, and the sSVGD manages to well replicate the posterior estimations using only 60,000 forward evaluations. Given that the MCMC should converge to the true posterior, it's inability to correctly estimate the anisotropy in the case it is not fixed (left hand column) suggests that the travel-time data does not sufficiently constrain this problem.

It is worth noting that improvements in the reconstruction of the true velocity curves are observed when any one of the three parameters is assumed known and fixed. However, in these cases (where either θ or S are assumed known) estimating the parameter A remains particularly challenging, with its marginal posterior consistently showing significant uncertainty. We believe this is due to the lower bound constraint at $A = 0$. which causes the posterior to peak at the true value (zero) but exhibit a positive tail. This asymmetry likely causes the anisotropic layer to compensate for the elevated anisotropy inferred in the isotropic layers - an effect driven by their asymmetric distribution. As a result, the model accommodates higher anisotropy in the isotropic layers by allowing the anisotropic layer to absorb the discrepancy, leading to greater uncertainty in the estimation of A.

IV. Conclusions

This paper presented a probabilistic methodology for constructing the spatially varying group velocity curves in a heterogeneous and locally anisotropic domain from ultrasonic travel time data. Three parameters, which together are effectively a proxy for the elastic stiffness tensor, were introduced as the key inversion parameters. We demonstrated that even in a simple synthetic case where data is collected across all boundaries of the domain, travel-time data fails to fully constrain this problem. However, by adding additional prior information on at least one of these parameters, we can achieve a significant improvement in the accuracy of the reconstructed group velocity curves.

REFERENCES

[1] C. Holmes, B. W. Drinkwater, P. D. Wilcox, "Post-processing of the full matrix of ultrasonic transmit–receive array data for non-destructive evaluation," NDT & E Int., vol. 38, no. 8, pp. 701-11, Dec 2005.

[2] O. Nowers, D. J. Duxbury, j. Zhang, B. W. Drinkwater, "Novel ray-tracing algorithms in NDE: Application of Dijkstra and A$*$ algorithms to the inspection of an anisotropic weld," NDT & E Int., vol. 61, pp. 58-66, Jan 2014.

[3] C. Nageswaran, C. Carpentier, Y. Y. Tse, "Microstructural quantification, modelling and array ultrasonics to improve the inspection of austenitic welds," Insight-NDT and Condition Monitoring, vol. 51, no. 12, pp.660-6, Dec 2009.

[4] P. Stefanov, G. Uhlmann, A. Vasy, H. Zhou, "Travel time tomography," Acta Mathematica Sinica, Eng. Ser., vol. 35, no. 6, pp. 1085-114, Jun 2019.

[5] J. Zhang, A. Hunter, B. W. Drinkwater, P. D. Wilcox, "Monte Carlo inversion of ultrasonic array data to map anisotropic weld properties," IEEE TUFFC, vol. 59, no. 11, pp.2487-97, Nov 2012.

[6] K. M. M. Tant, E. Galetti, A. J. Mulholland, A. Curtis, A. Gachagan, "Effective grain orientation mapping of complex and locally anisotropic media for improved imaging in ultrasonic non-destructive testing," Inv. Prob. in Sci. and Eng., vol. 28, no. 12, pp. 1694-1718, Dec 2020.

[7] J. Singh, K. M. M. Tant, A. J. Mulholland, C. N. MacLeod, "Deep learning based inversion of locally anisotropic weld properties from ultrasonic array data. Applied Sciences," vol. 6, no. 12, pp. 532, Jan 2022.

[8] M. K. Kalkowski, M. J. Lowe, v. Samaitis, F. Schreyer, S. Robert, "Weld map tomography for determining local grain orientations from ultrasound," Proc. of the Roy. Soc. A, vol. 6, no. 479, pp. 20230236, Sep 2023.

[9] X. Zhang, A. Curtis, "Seismic tomography using variational inference methods," J. Geophys. Res. Solid Earth, vol. 125, no. 4, pp. e2019JB018589, Apr 2020.

[10] V. Gallego, D. R, Insua, "Stochastic Gradient MCMC with Repulsive Forces," NIPS Workshop on Bayesian Deep Learning, Online, Nov 2018.

[11] J. Ludlam, K. Tant, V. Dolean, A. Curtis, " Travel times and ray paths for acoustic and elastic waves in generally anisotropic media," Journal of Comp. Phys. vol. 494, pp. 112500, Dec 2023.

[12] G. D. Connolly, M. J. Lowe, J. A. Temple, S. I. Rokhlin, "The application of Fermat's principle for imaging anisotropic and inhomogeneous media with application to austenitic steel weld inspection," Proc. Roy. Soc A, vol. 8, no. 65, pp. 3401-23, Nov 2009.

[13] Q. Liu, D. Wang, "Stein variational gradient descent: A general purpose bayesian inference algorithm," Advances in neural information processing systems, vol. 29, 2016.

[14] G. Harvey, A. Tweedie, C. Carpentier, P. Reynolds, "Finite element analysis of ultrasonic phased array inspections on anisotropic welds," AIP Conf Proc, vol. 1335, no. 1, pp. 827-834, Jun 2011.

[15] N. Rawlinson, M. Sambridge, "Wave front evolution in strongly heterogeneous layered media using the fast marching method," Geophys. J. Int. vol. 156, no.3, pp.631-47, Mar 2004.

[16] M. D. Zeiler, "Adadelta: an adaptive learning rate method," arXiv preprint arXiv:1212.5701, Dec 2012.