

Research



Cite this article: Kalkowski MK, Lowe MJS, Samaitis V, Schreyer F, Robert S. 2023 Weld map tomography for determining local grain orientations from ultrasound. *Proc. R. Soc. A* **479**: 20230236.

<https://doi.org/10.1098/rspa.2023.0236>

Received: 4 April 2023

Accepted: 14 July 2023

Subject Areas:

mechanical engineering

Keywords:

non-destructive evaluation, ultrasonics, tomography, array imaging, austenitic welds

Author for correspondence:

Michał K. Kalkowski

e-mail: m.kalkowski@soton.ac.uk

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.6761807>.

Weld map tomography for determining local grain orientations from ultrasound

Michał K. Kalkowski^{1,2}, Michael J. S. Lowe¹,
Vykintas Samaitis³, Fabian Schreyer⁴ and
Sébastien Robert⁵

¹Mechanical Engineering, Imperial College London, London SW7 2AZ, UK

²Institute of Sound and Vibration Research, University of Southampton, Southampton SO17 1BJ, UK

³Prof. K. Baršauskas Ultrasound Research Institute, Kaunas University of Technology, K. Baršausko Street 59, Kaunas 51423, Lithuania

⁴Materials Testing Institute (MPA) University of Stuttgart, Pfaffenwaldring 32, 70569 Stuttgart, Germany

⁵Paris-Saclay University, CEA-List, Gif-sur-Yvette 91191, France

MKK, 0000-0001-7733-3178

The complex structure of inhomogeneous welds poses a long-standing challenge in ultrasonic non-destructive testing. Elongated grains with spatially varying dominant orientations can distort and split the ultrasonic beam, hindering inspection data interpretation. One way to tackle this problem is to include material information in imaging and signal analysis; however, such information is often only gathered using destructive methods. This paper reports the development of a physics-based weld inversion strategy determining grain orientations using a ray tomography principle. The considered approach does not rely on a macroscopic weld description but may incorporate it to facilitate inversion. Hence, it is more general than other available approaches. The methodology is demonstrated in both numerical and experimental examples. The experimental work focuses on mock-up samples from the nuclear industry and a sample manufactured during this research. The ‘ground truth’ for the latter comes from an EBSD evaluation—the most accurate (yet destructive) examination

© 2023 The Authors. Published by the Royal Society under the terms of the Creative Commons Attribution License <http://creativecommons.org/licenses/by/4.0/>, which permits unrestricted use, provided the original author and source are credited.

technique available. Across the considered specimens, our methodology yielded orientation maps with average errors well below 20° , leading to time-of-flight errors below $0.05\ \mu\text{s}$. Applying the result from inversion to ultrasonic imaging offered between 5 and 14 dB signal-to-noise ratio improvement for defect signatures.

1. Introduction

Ultrasonic inspection of inhomogeneous welds is a significant challenge in the nuclear industry. Their distinctive microstructure—long columnar grains with preferential orientations varying across the weld—is responsible for a surge in the complexity of ultrasound propagation. First, the ultrasonic beam no longer travels across straight paths but follows the minimum propagation time principle resulting in curved, complicated rays. Consequently, locating a defect based on time-of-flight (ToF) is only viable with detailed material information and a suitable propagation model. Another troublesome effect is grain scattering. When passing through the subsequent grain boundaries, the ultrasonic energy is scattered, leading to the loss of amplitude of the main beam (attenuation) and elevated structural noise. Other issues include, e.g. beam splitting, that is, the emergence of a secondary beam at a grain boundary that gives rise to two strong scattered waves.

The above mentioned issues may be partially accounted for if material information is available. Such knowledge allows for correcting ultrasonic images, which would then consider realistic propagation times. It would also help interpret conventional inspection acquisitions by revealing the actual beam paths and potential scatterer locations [1]. Material information can be obtained from classical examination (e.g. metallography) or EBSD imaging. However, these examination methods are destructive, costly and offer no outlook for the application on the plant. Alternatively, one may assume a specific structure of the weld (e.g. the gold-standard MINA model (Modelling anIsotropy from Notebook of Arc welding) [2], or the Ogilvy map [3]), but these simplified weld description models cover limited ranges of cases, and their parameters need to be tuned. These limitations can be circumvented with inversion based on ultrasound, allowing for extracting the orientation information from the ultrasonic measurements directly.

It is widely accepted that to support ultrasonic testing, one does not need to consider detailed grain topology [4–6]. Instead, a macroscopic description is sufficient, where the grains are locally homogenized to yield a single dominant orientation for each cell in a chosen grid (often between 1 and 2 mm). The behaviour of the columnar grains with similar preferential orientations can be approximated using a transversely isotropic material with the axis of material symmetry colinear with the dominant orientation [6]. Historically, macroscopic weld descriptions were first developed to improve the understanding of inspection results. Their parameters were assumed based on the examination data and practical experience. Such geometric models include the famous Ogilvy map [3], the layered arrangement [7] or the more basic approximation using only two angles by Langenberg *et al.* [8]. A refined orientation layout (referred to as ‘weld map’ hereafter) linked to welding parameters came with the MINA model [2], which draws the information from the welder’s notebook and captures the solidification processes in a phenomenological manner. All these approaches required assuming some if not all parameters, especially for a weld of which little was known.

The first attempts to use ultrasonic measurements to determine orientations were aimed at updating the weld description parameters mentioned above. For instance, Gueudre *et al.* [9] used evolutionary optimization to guess the ‘arbitrary’ MINA parameters (remelting rates and tilt angles) based on the echodynamic curve. Problems of this approach included the uncertainty related to transducer coupling when measuring amplitudes and the limited range of configurations it could cover. The same parameters were later determined from the ToF maps by Fan *et al.* [10], also based on a genetic algorithm. Inversion based on Bayesian updating and the Ogilvy map was proposed in [11], where Bayes inference helped estimate the parameters

of the Ogilvy model. Weld description constraints were removed in another Bayes-based work [12,13], where the orientations were identified on a per-cell basis. While delivering promising results, the Bayesian approaches incurred high computational costs, which was deemed one of their fundamental shortcomings.

Recent advances in the availability of high computing power led to the development of approaches based on soft computing. Singh *et al.* [14,15] proposed employing deep neural networks to determine the orientation maps based on a large simulated database of times of flight through different grain layouts. While this approach opens the possibility for real-time application, it still requires a computationally expensive training stage and does not exploit the physics of wave propagation through the heterogeneous material. An alternative method that takes advantage of modern optimization methods was developed by Ménard *et al.* [16,17]. Their focus is on optimizing the amplitude of the ultrasonic image around a postulated defect by adjusting the parameters of a macroscopic weld description. While the primary outcome of the method is the enhanced ultrasonic image (with improved defect signature location and the signal-to-noise ratio (SNR)), a grain orientation map comes as a by-product.

Although the developments reported in the previous paragraph provided solutions to some challenges, they were often based on 'black-box' optimizers, burdened by high calculation cost, requiring a pre-identified defect area or limited to a given macro weld description. To the authors' knowledge, physics-based inversion was not fully exploited for the applications of interest. Hence, we developed a ray inversion-based method for characterizing inhomogeneous welds based on ultrasonic array data acquired at the surface. Compared with the prior work, this is a more general approach, based on the physics of wave propagation.

ToF tomography is a known method applied in various contexts, including the biomedical field, e.g. breast tomography [18]. However, the materials considered in those problems are locally isotropic, which is not the case in complex welds. To account for the material complexity and other practical differences, e.g. the sensor arrangement, we extended the classical ray inversion, developing and validating a new approach suited to inhomogeneous welds called weld map tomography.

It is worth stating at the beginning that obtaining a precise orientation map is not a challenge we are pursuing in this paper. The purpose of the inversion is to achieve a workable engineering answer providing a significant improvement over the assumption of spatial homogeneity. Precision, especially in local detail, is simply not useful since it does not translate into enhanced ultrasonic images. This context is reflected in the construction of the weld map tomography algorithm, including the choice of starting map, regularization, convergence criteria and evaluation.

The method is presented in more detail in the remainder of this article. In the next section, we outline the Fermat principle-based forward model underpinning the inversion. Then, we recall the fundamentals of ray tomography for structures comprising isotropic materials and present the extension enabling local anisotropy to be accounted for. Then, we discuss the details of implementing inversion for welds, including our two-stage approach and regularization. Examples follow the description—first, a numerical example based on grain-scale finite-element simulation of a real weld case from EDF; second, a set of experimental examples of industry-relevant welds—two mock-ups from EDF and a weld specifically manufactured and characterized for this work.

2. Forward model

Weld map tomography matches the simulated ultrasound ToF with the experimental target by adjusting model parameters. The forward model, predicting times-of-flight for a given layout of grain orientations, is essential to the inversion. The most accurate yet computationally expensive route is to use time-domain finite-element simulations, which are extremely versatile for the material structure and geometry. An alternative, offering a potentially more efficient solution, is to use a ray tracing model. Previous research on ultrasound in inhomogeneous welds [1,10]

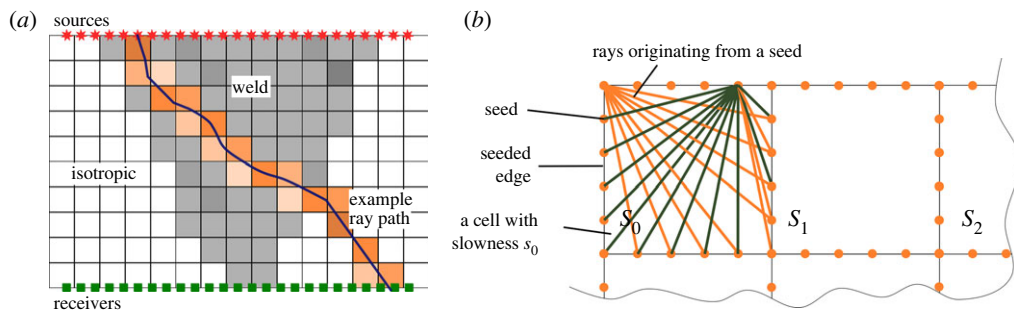


Figure 1. Schematic diagrams of the SRP model. (a) Computational representation of a weld in the SRP model. The orange-coloured cells represent those crossed by the depicted ray. (b) Seeding of cells in the SRP model. Two sets of cell rays are shown; for clarity, rays co-linear with the horizontal and vertical grid lines are not drawn.

used a ray-stepping algorithm based on scattering at fictitious grain boundaries according to a chosen step size. While accurate, it provided the dominant energy path for a given shoot angle (not a selected transmitter–receiver pair) and hence was difficult to be implemented efficiently. One needed to compute a large range of shooting angles and interpolate to obtain the arrival time at a desired location.

A crucial function of a forward model in this context is to determine the ToF between transducers directly, without the need for interpolation. We can achieve this by using the Fermat principle stating that the propagation path between two points is always such that the propagation time is minimal. This principle can be effectively used with efficient shortest path-finding algorithms popular from the origins of computer science. Zhang *et al.* [11] suggested using various incarnations of the Dijkstra method to model ultrasound in welds, which was further studied in more detail in [19,20]. The present model draws from those contributions but with some modifications to improve efficiency. Given its origins, the forward model used here is called the shortest ray path (SRP) model.

Nowers *et al.* [19,20] suggest using a randomly distributed set of nodes and a selected search radius to maximize the accuracy of the prediction. Our approach is based on the fact that the common macroscopic weld descriptions impose a regular grid structure, with local orientations varying from cell to cell. Propagation across each cell is determined by its assigned orientation and the incident angle. Moreover, the tomography algorithm itself is also typically programmed over a grid. Hence, we decided to use the regular grid in the forward model to maintain consistency across different stages of inversion, as shown in figure 1a. The edges of each grid cell are seeded according to a chosen resolution, and the paths connecting all points around the cell are candidate propagation path segments, which is illustrated in figure 1b. It is worth noting that a single segment only crosses a single cell; that is, segments belonging to different cells do not cross except at their start or end points.

Figure 1a shows an illustrative computational domain with a weld, isotropic parent materials to its left and right, and a set of sources and receivers. The cells generally represent different properties, but they are discretized in the same way for the ray model—using the same number of seeds along each edge. The number of seeds determines the angular accuracy of the model, i.e. the range of allowable angles for a ray passing through the cell. Figure 1b schematically depicts this set-up.

The SRP model looks for the shortest propagation time between two chosen points. The propagation times across cells are calculated from the group slowness at given combination angles (a simple function of the local grain orientations and incident angles). Those propagation times can be visualized as lengths of edges between nodes in a graph. Our model first computes the propagation times for the whole model. Then, with the entire network of edges defined, a

SRP algorithm looks for the path with the lowest cost (shortest propagation time). We used the `scipy.sparse.csgraph` [21] implementation of the SRP problem.

One immediate advantage of the way we represent the material is that rays within each cell have the same structure, e.g. they cover the same angles and have equal lengths. Therefore, they can be precomputed only once for a single cell and used for the whole model. Further, this arrangement benefits from the common assumption allowing for using a single ‘starting’ elastic tensor for the weld. The homogenized elastic tensor in material coordinates is the same for the whole weld, but it is rotated on a cell-by-cell basis. The chosen construction of the SRP model only requires calculating a single group slowness as a function of the combination angle and delivers a quick calculation of the entire ToF graph, offering significant computational advantages over the previously reported approaches.

The convergence study is not reported here for brevity (included in the electronic supplementary material [22]), but for cases we studied, 10 seeds for a 2 mm grid delivered good accuracy compared with time-domain finite-element simulations. A finer grid may be needed to represent the geometry of the weld (chamfer, backwall), and this needs to be verified for each individual case.

The SRP model, despite its high efficiency and versatility, has several limitations, too. They stem from its underpinning principle—focusing on the shortest propagation time between two points without considering amplitude or energy. Consequently, the SRP model cannot simulate beam splitting and secondary arrival times which sometimes may carry most of the energy. Further, it may output null paths, that is, return ray paths associated with negligible energy transfer. However, these limitations do not affect the inversion algorithm, as the inversion relies on measured arrival times. If the arrival time for a particular transducer pair cannot be identified from measurements, it is not included in the inversion. However, the SRP model and the inversion algorithm are blind to energy transfer. They may, in principle, suggest a null path for a transducer combination associated with energy transport. This has not been observed in this research but the associated uncertainty remains a subject of ongoing investigation.

3. Ray inversion

This section covers the main elements of the ray inversion algorithm, which is an extension of a method proposed for materials with heterogeneous isotropy by Hormati *et al.* [18]. The reader is encouraged to consult [18] about the details of the technique for ray angle-independent slowness.

Acoustic tomography aims to recover unknown local properties from the characteristics of sound propagation through a medium. Complex welds involve heterogeneous anisotropy, with the local group slowness being a function of the incident angle and the local grain orientation. This dependence needs to be included in the inversion. Below we assume that all cells are characterized by the same elastic tensor, appropriately rotated according to local grain orientations; hence, the weld map tomography aims at determining local grain orientations for a known elastic tensor. This is analogous to previously published research (e.g. [10–12]).

(a) Heterogeneous anisotropic medium

In a heterogeneous medium (even if locally isotropic), when travelling between two locations, the sound takes curved paths, depending on the local properties. The ToF of a finite-length wave packet between two points is the integral of the group slowness along the propagation path:

$$\tau = \int_{\Gamma} s(\phi) dl, \quad (3.1)$$

where τ is the time, Γ is the propagation path, s is the slowness, ϕ is the combined angle and dl is the infinitesimal length along the path. The notion of the combined angle ϕ is explained below. Since all cells in the SRP model share the same ‘basic’ tensor, the group slowness curve has the same shape but is rotated from cell to cell. The combined angle ϕ is the difference between




Figure 2. Group slowness for a transversely isotropic material—how the incident angle and grain orientation form a combined ray angle in material coordinates.

the incident angle and the local orientation in a cell. Figure 2 schematically depicts the angles involved and illustrates how the combined angle is formed.

The SRP model discretizes the domain, allowing for expressing propagation times in matrix form

$$\boldsymbol{\tau} = \mathbf{L}\mathbf{s}(\boldsymbol{\phi}), \quad (3.2)$$

with \mathbf{L} being the ray-length matrix and \mathbf{s} a vector describing group slownesses across the domain, depending on local orientations $\boldsymbol{\theta} = \boldsymbol{\gamma} - \boldsymbol{\phi}$ (referring to the notation in figure 2). A single ray typically passes through only a small proportion of all cells, so \mathbf{L} is a sparse matrix.

With a measured set of times-of-flight $\tau_{i,\text{exp}}$, the aim of the inversion is to determine grain orientations that minimize a cost function, which is defined as

$$C = \sum_{i=0}^M (\tau_{i,\text{model}} - \tau_{i,\text{exp}})^2, \quad (3.3)$$

where M is the total number of rays considered (i.e. the number of transmitter–receiver pairs) and $\tau_{i,\text{model}}$ are times-of-flight from the model.

After Hormati *et al.* [18], we use the nonlinear conjugate gradient approach [23] to obtain the optimal weld map. This approach requires finding the gradient of the cost function with respect to model parameters, that is, the local grain orientations:

$$\frac{\partial C}{\partial \theta_j} = \sum_{i=0}^M 2(\tau_{i,\text{model}} - \tau_{i,\text{exp}}) \left(\frac{\partial \tau_{i,\text{model}}}{\partial \theta_j} - \frac{\partial \tau_{i,\text{exp}}}{\partial \theta_j} \right), \quad (3.4)$$

where the sum is executed across all M rays (transmitter–receiver pairs), and j is the index of a grid cell in the model domain. The gradient depends on local derivatives of the propagation time over grain orientation. For a single path i , the propagation time can be written as a discrete sum

$$\tau_i = \sum_{j=0}^N s_j(\theta_j) l_{i,j}, \quad (3.5)$$

over all N cells, where s_j and $l_{i,j}$ are cell's group slowness and the length of ray i within cell j , respectively. The derivative of the propagation time of path i over grain orientation in cell j becomes:

$$\frac{\partial \tau_i}{\partial \theta_j} = \frac{\partial s_j}{\partial \theta_j} l_{i,j} + s_j \frac{\partial l_{i,j}}{\partial \theta_j}. \quad (3.6)$$

Here, we make a fundamental assumption (following [18]) that the ray length within a cell remains constant, irrespective of the local properties, i.e. $\partial l_{i,j} / \partial \theta_j = 0$. This is an approximation, but it simplifies the inversion considerably, removing the second component of the sum in equation

(3.6). It is worth noting that path variations will be implicitly accounted for in the dependence of group slowness on the combined angle and subsequent iterations of the gradient stepping algorithm. The former links to the fact that local group slowness is defined as a function of the combined angle ϕ in the inversion model; hence,

$$\frac{\partial s_j}{\partial \theta_j} = \frac{\partial s_j}{\partial \phi_j} \frac{\partial \phi_j}{\partial \theta_j} = -\frac{\partial s_j}{\partial \phi_j}. \quad (3.7)$$

Hence, the final expression for the derivative becomes

$$\frac{\partial \tau_i}{\partial \theta_j} = -\sum_{i=0}^M l_{i,j} \frac{\partial s_j}{\partial \phi_j} \Big|_{\phi_j=\gamma_j-\theta_j}, \quad (3.8)$$

with the sum executed over all rays (if a ray does not pass through cell j , its length $l_{i,j} = 0$).

Denoting the difference between the arrival times from the model and from the measurement as the residual r_i , and noting that the measured arrival times do not depend on the weld map in the model, we write

$$\frac{\partial C}{\partial \theta_j} = \sum_{i=0}^M -2r_i l_{i,j} \frac{\partial s_j}{\partial \phi_j} \Big|_{\phi_j=\gamma_j-\theta_j}. \quad (3.9)$$

It is worth noting that the angle-dependent term needs to be evaluated for each ray separately at every grid cell, which adds complexity and computational cost, but they are in part compensated for by the computational arrangement that takes advantage of the regular grid structure.

Using the tools outlined above, the inversion algorithm takes the following steps:

1. Choose the initial orientation distribution θ_0 .
2. Run the forward model for θ_0 to get the ToF estimate τ_0 and evaluate the cost function $C(\tau_0)$.
3. For each n of at most N iterations or until convergence:
 - 3.1. Calculate the gradient of the cost function with respect to the weld map $\nabla C(\theta_n)$ according to equation (3.9).
 - 3.2. Set the descent direction $\Delta \theta_n$

$$\Delta \theta_n = \begin{cases} -\nabla C(\theta_0), & \text{for the first iteration} \\ -\nabla C(\theta_0) + \kappa_n \Delta \theta_{n-1}, & \text{for subsequent iterations} \end{cases} \quad (3.10)$$

where $\kappa_n = \|\nabla C(\theta_n)\|_2^2 / \|\nabla C(\theta_{n-1})\|_2^2$

- 3.3. Update the weld map

$$\theta_{n+1} = \theta_n + \mu_n \Delta \theta_n, \quad (3.11)$$

where μ_n is the step size. We calculate the step size using the back-tracking line search algorithm detailed in [18] with control parameters $\alpha = 0.4$ and $\beta = 0.3$ (as defined in [18]).

- 3.4. Run the forward model for θ_{n+1} to get the ToF estimate τ_{n+1} and evaluate the cost function $C(\tau_{n+1})$.
 - 3.5. Check convergence or continue the loop. We defined convergence as three subsequent iterations that diminish the cost function by at most 2% with respect to the latest value.
4. Return the converged weld map.

The κ_n parameter determines the effect of the preceding step and can be calculated using the Fletcher–Reeves quoted above [23] or alternative formulae [24]. In general, the stepping may not decrease the cost function, in which case the back-tracking line search algorithm ensures that the cost function reduces enough. This algorithm iteratively decreases the step size by a chosen factor until the cost function drop is less than the quoted threshold.

From the computational viewpoint, each step of the gradient stepping approach requires solving the forward model once to obtain the gradient, and once for each iteration of the backtracking line search, should it be necessary. This arrangement offers a very efficient set-up, with the number of forward solutions usually not much greater than the number of steps, thanks to the convenient expression for the gradient developed above.

(b) Implementation for an inhomogeneous weld

The algorithm is suitable for generally anisotropic media; however, its application to inhomogeneous welds gives a specific flavour to several aspects, briefly outlined below. The key assumption behind this implementation is that the geometry of the weld must be known prior to inversion, which determines only grain orientations; the general methodology, however, is not bound by this assumption and the uncertainty in the geometry and tensor properties is a subject of ongoing research.

(i) Material structure and elastic tensor for the weld

Most parent materials joined using the welds of interest to this paper can be modelled as isotropic (the effect of the heat-affected zone is beyond the scope of this paper). Hence, a typical inversion set-up considered in this work is solving only for grain orientations within the weld, with the parent material isotropic properties assumed *a priori* [10,11]. Furthermore, the 'basic' transversely isotropic elastic tensor governing the columnar grain structure of the weld is considered constant, i.e. only the orientations are allowed to vary [9–11]. This assumption allows for simplifying the computational domain and discretizing only the weld region, considerably reducing the size of the graph over which the shortest paths are sought.

(ii) Transducer configuration

Transducer configuration is a challenge for weld map tomography. The classical tomography set-up (optical and acoustic) assumes a circular array of transmitter–receivers, ensuring complete coverage of view angles [18]. Any set-up different from the full view incurs a loss in the ability to recover the properties of the medium [25]. Since weld map tomography relies on measured arrival times, one can realistically consider three configurations: (i) a through-transmission set-up, with transducers on both the top and bottom surfaces; (ii) a pulse-echo set-up, with only one array on top of the weld; or (iii) a tandem configuration with two arrays located on the top surface on both sides of the weld. The illustration of the implication of the reduced view for the through-transmission and the pulse-echo set-up can be found, e.g. in [25].

The through-transmission set-up is straightforward for the considered model, as the sources and receivers are placed along the top and bottom edges of the computational domain. However, the pulse-echo and the tandem set-up cannot be modelled directly. To account for the reflection from the backwall, one needs to consider an extended domain in the SRP model, including an auxiliary part mirrored around the backwall [10]. Such an arrangement, based on the Fermat principle, allows for modelling propagation from the backwall to the transducer. Nodal positions in the auxiliary domain will have the y -coordinates and orientation angles appropriately flipped. Figure 3 schematically depicts the arrangement when through-backwall propagation is considered.

The considered configurations offer the possibility of characterizing welds without sectioning but may not be suitable for *in situ* applications (where weld caps are present). One way to circumvent that is to use an enhanced experimental procedure, similar to that reported in [26,27]. Also, other array configurations, such as the tandem one used in [10], or the single-side access as in [17] could be used with the current modelling/inversion scheme. Both development directions are the subject of our current research.




Figure 3. Auxiliary mirrored domain for inversion through the backwall: (a) the set-up and (b) illustration of the reduction of the view angle between through-transmission and pulse-echo.

(iii) Incorporating the known information and regularization

Weld map tomography is an ill-posed problem with multiple minima. Compared with the isotropic set-up, the added complexity of the anisotropic weld context comes from the fact that the group slowness is not unique with respect to the ray angle. Hence, incorporating additional information is necessary for the inversion to succeed. Knowledge about typical characteristics of weld maps can offer a plausible initial guess and direct selecting regularization schemes.

For the initial guess, we suggest taking advantage of simple weld map models as initial guesses for inversion. In particular, we used the Ogilvy map [3] as the starting orientation layout. The inversion then targets variations of the weld map from this plausible guess, offering a higher chance of convergence compared with the case where the starting angles are random.

Regularization involves adding terms to the cost function equation (3.3), which are then included in the gradient calculation at each step. In this work, we used generally available regularization schemes, following intuitions from observations of metallography/EBSD examinations of welds [2,6,9,17,28–30]. They show that the variation of angles is typically smooth, with just a few abrupt changes. Moreover, the orientations following welding passes tend to form regions with similar dominant orientations that vary smoothly across the cross-section. This naturally leads to smoothness and total variation regularization terms, which were adopted in this paper.

The analysis of destructive examinations, including the papers quoted above, also suggests that a dominant orientation of grains often exists. For the most uncomplicated symmetric cases, this will be the vertical direction, i.e. perpendicular to the top surface. For other welds, it may be the direction almost parallel to the chamfer. This dominant orientation affects the propagation path considerably, and one approach to take advantage of this knowledge is to divide the inversion into two stages. Starting from the conventional Ogilvy map, in the first stage, the algorithm only adjusts the dominant orientation, that is, a single angle. Then, having reached convergence, it proceeds to updating local orientations across the weld. This approach required a modification of the Ogilvy model (for which the weld was symmetric, with vertically oriented grains close to the centreline), referred to as the ‘generalized Ogilvy map’, similar to that used in [17]. The modification introduces the dominant angle ζ , which represents the effect observed in examinations. The details of the implementation are given in appendix A.

In the first stage, the inversion algorithm starts from a ‘default’ Ogilvy map ($T_L = T_R = n_L = n_R = 1$; $\zeta = 0$; the symbols are explained in appendix A). The gradient of the cost function is now evaluated with respect to the dominant orientation ζ :

$$\frac{\partial C}{\partial \zeta} = \sum_{i=0}^M -2r_i l_{i,j} \frac{\partial s_j}{\partial \phi_j} \Big|_{\text{at } \phi_j = \gamma_j - \theta_j} \frac{\partial \theta_j}{\partial \zeta} \Big|_{\text{at } \zeta_k'} \quad (3.12)$$




Figure 4. The macroscopic weld description with dominant orientations for each zone marked (a), and the generated synthetic microstructure (b) for the numerical example of weld map tomography. The orientations are quoted in degrees, measured from the vertical direction, with positive angles anticlockwise.

with ζ_k being the current dominant orientation angle. The derivative of the weld map with respect to the dominant orientation ζ can be pre-calculated numerically before inversion, based on the Ogilvy map model from appendix A, and presents negligible computational cost. The remaining steps of the inversion procedure remain the same, including the convergence criterion, except for the fact that now only the scalar value of dominant orientation is updated, instead of the entire weld map. Grain orientations are calculated from the Ogilvy map at each step, according to the updated dominant orientation ζ .

This approach was found particularly effective in our work and represented a pragmatic balance between the flexibility of ultrasonic tomography and the prior knowledge about the structure and the behaviour of welds of interest.

4. Numerical example

(a) Numerical microstructure

Having outlined the conceptual principles of weld map tomography, the attention now turns to examples. First, we present a numerical example in which the target times-of-flight came from grain-scale time-domain finite element simulations. The calculations were executed on a graphical processing unit (GPU) using Pogo, a fast GPU finite-element solver developed at Imperial College London [31]. The example is motivated by a real industrial inspection case at EDF, France [29,32]. The grain structure of the weld bent and split the incident ultrasonic beam, so the reflection from the backwall at some locations could be misinterpreted as a defect. In [32], the author developed a macroscopic weld description to support the interpretation of the inspection, based on metallography (figure 4). The procedure is described in [32]. After mechanical polishing, the samples were submitted to chemical etching with a solution of aqua regia, that is, 35% nitric acid and 35% chlorhydric acid. This preparation is particularly well suited for revealing the epitaxial growth of austenitic weld. After this treatment, the sample was observed using an optical microscope. The analysis led to dividing the weld into several regions with dominant orientations, which are depicted in figure 4. In this paper, we used that description to create a numeric microstructure, allowing us to capture the rich physics of ultrasound propagation through complex welds in simulation.

The numerical microstructure was generated using a free polycrystal generation package—Neper [33].

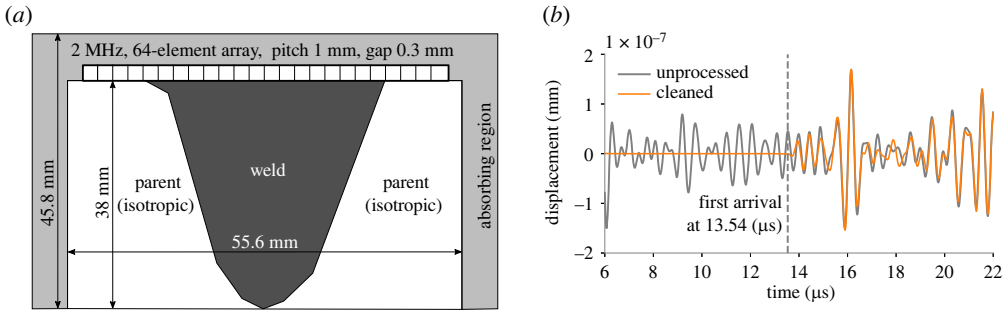


Figure 5. FE modelling of the numerical example. (a) Schematic diagram. (b) Illustrative simulated time trace—array element no. 41 transmitting and element no. 36 receiving.

First, we created a columnar grain microstructure with vertically oriented grains by stretching a standard Poisson–Voronoi tessellation along one dimension [33,34]. Then, for each orientation region depicted in figure 4a, the starting model was rotated and cropped, to match the orientation and the shape of the region. After numerical microstructures for all regions were created, they were stitched into a complete weld representation shown in figure 4b. For each region, we generated a set of crystallographic orientations, uniformly randomly distributed along the dominant orientation (referred to as ‘fibre’ orientations in Neper).

(b) Grain-scale finite-element model—virtual experiment

The generated microstructure was the basis for setting up the grain-scale finite-element model, acting as a virtual experiment, depicted schematically in figure 5a. The weld region was surrounded by isotropic regions, representing the parent plates ($c_{11} = 260$ GPa, $c_{12} = 100$ GPa, $c_{44} = 80$ GPa, $\rho = 8000$ kg m⁻³). For each grain belonging to the weld, we rotated the crystal elasticity matrix (cubic crystal, $c_{11} = 206$ GPa, $c_{12} = 133$ GPa, $c_{44} = 119$ GPa, $\rho = 8000$ kg m⁻³) according to Euler angles generated in neper. We simulated the action of a 64-element 2 MHz array (pitch: 1 mm, gap: 0.3 mm) on top of the weld by a series of vertical nodal forces (one force per array element) separated by the pitch distance and acting on the top boundary of the weld with a three-cycle toneburst signal. The simulation was a numerical equivalent to a full matrix capture (FMC). To minimize the impact of the finite computational domain and eliminate reflections, we used absorbing regions around the whole domain (except the backwall). The discretization used a regular quad mesh (CPE4-type linear elements) with an element size of 0.025 mm, which was a very conservative choice providing more than 100 elements per longitudinal wavelength. The time step was chosen based on the Courant number of 0.9, giving 3.7 ns. The model had nearly 12 million degrees of freedom and was solved efficiently using Pogo [31] on a GPU card (illustrative solution times for one array shot with over 6800 time steps: standard gaming card Nvidia GTX 1080Ti—15 s, Nvidia RTX 2080Ti—10 s).

Weld map tomography, as used in this paper, is based on backwall-reflected first arrivals, that is, longitudinal waves. To simplify extracting them from the simulation, we executed an auxiliary calculation with an absorbing boundary instead of the backwall. Subtracting the auxiliary from the original FMC dataset helped eliminate backscatter and led to clean first arrivals at the top surface. Such an approach is impossible in practice but helped focus solely on the inversion algorithm rather than peripheral issues in the numerical example. Figure 5b shows an example time trace, unprocessed and cleaned signals, and the estimated arrival time. The estimation used an implementation of the Akaike information criterion [35], which is particularly effective if the first arrival is not the strongest in amplitude. For experimental (or uncleaned) datasets, obtaining an automated universal procedure is a research challenge on its own, and to keep the focus on inversion, we used a procedure involving manual input in other cases.

(c) Forward ray tracing model

The inversion requires first setting up the forward SRP ray tracing model. The domain was cast on a regular 1 mm grid to accommodate weld geometry. The edges of each cell were seeded with 10 nodes, and, to take the computational advantage described in §2, the algorithm pre-calculated the potential paths across a single cell. The pulse-echo configuration was modelled by mirroring the domain around the backwall, as described in §2.

The forward model uses homogenization; that is, it does not consider individual grains but assigns local material properties based on the average grain orientation within an SRP model grid cell. Regions with columnar grains aligned to the same dominant angle are known to behave like a transversely isotropic material [2,4,6]. The material constants for the SRP model were determined from the grain-scale model using Voigt homogenization [36], knowing that the dominant orientation aligns with the direction of the smallest stiffness [6]. The homogenization yielded the following transversely isotropic material constants: $c_{11} = 248$ GPa, $c_{12} = 91$ GPa, $c_{13} = 133$ GPa, $c_{33} = 206$ GPa, $c_{44} = 119$ GPa, $c_{66} = 77$ GPa, and the same density as for the parent material $\rho = 8000$ kg m⁻³.

The vertically oriented grains have the 33 (soft) material axis aligned with the global z -axis in the SRP model domain. Orientation angles are measured from the vertical, positive anticlockwise. Note that the SRP model is defined in the y - z domain, with z being vertical, while the FE model is in the x - y domain. This required rotating the reference frame for the elasticity matrix to achieve equivalent behaviour (in FE, the ‘soft’ 33 material direction aligns with the y global direction).

(d) Inversion

We used the two-stage inversion described above. First, the algorithm started with a ‘standard’ Ogilvy map ($T_L = T_R = n_L = n_R = 1$, $\alpha_L = \alpha_R = 20$) with a dominant vertical orientation ($\zeta = 0^\circ$), converging to -4.7° after six iterations. Then, in the per-cell inversion stage, we used smoothing and total variation regularization, achieving convergence after a further 23 iterations.

Figure 6 shows the obtained weld map and the evolution of the cost function. The vertical line marks the point where the first stage, in which only the dominant angle was varied for a general Ogilvy map, was completed. Individual orientations were then updated to the point where for the three consecutive iterations, the cost function dropped by less than 2% (this criterion was adopted for both stages). The weld map is compared with the original microstructure in figure 4, together with the orientation maps at the beginning, and after the first stage finishes.

Detailed error maps for the absolute angular and ToF errors are presented in figure 7, whereas the histograms are in figure 8. The mean, median and maximum absolute angular errors ($|\theta_{\text{original}} - \theta_{\text{inversion}}|$) are 7.28, 5.71 and 33.06°, respectively. From a practical viewpoint, these are promising results compared with previously reported average errors exceeding 20° [10]. We also compared the converged ToF map with that chosen as a target for inversion in figure 7b. The mean, median and maximum absolute time errors ($|\tau_{\text{original}} - \tau_{\text{inversion}}|$) are 0.03, 0.02 and 0.27 μs , respectively.

The numerical example demonstrates the potential of weld map tomography and quantifies its performance. While errors exist, they are below the previously reported level [10] (as expected from a simulation-based study) and related to the limitations of ToF tomography (the inability to resolve sharp transitions and smaller orientation zones) [37,38], simplifications in the forward model and the ill-posedness of the problem (multiple local minima). As stated at the beginning, we are not pursuing a perfect orientation map since small local details do not play a major role. Rather, the interest is in obtaining a map that provides a meaningful improvement to ultrasonic images, which will be demonstrated in the next section.

5. Experimental examples

In this section, we report exercising weld map tomography on three experimental samples: two industrial mock-ups from EDF and one specifically manufactured for this research. All samples




Figure 6. The result of weld map tomography based on time traces from grain-scale finite-element simulations: (a) obtained grain orientations; (b) the cost function; (c) starting grain orientations; and (d) grain orientations after the first stage of inversion.




Figure 7. Error maps of weld map tomography based on time traces from grain-scale finite-element simulations: (a) orientation errors and (b) time of flight errors.




Figure 8. Error histograms of weld map tomography based on time traces from grain-scale finite-element simulations: (a) orientation errors and (b) time of flight errors.




Figure 9. Schematic diagram for experimental weld map tomography on an EDF Inconel 182 V-weld with uniform orientations: (a) photograph and (b) experimental set-up.

had artificial defects created for testing imaging algorithms. Verifying the effectiveness of the inversion is not straightforward, as destructive material examinations usually describe a different part of the weld than that used in the ultrasonic measurement. On the other hand, the community agrees that grain orientations do not vary significantly along the weld except for arc ends [2]. However, from the practical inspection perspective, the most convincing evaluation of the obtained weld structure information is its application to imaging. Hence, for all examples, we demonstrate how the result of weld map tomography improves ultrasonic array images and the signatures of target defects.

(a) Weld with uniform orientations

First, we consider a homogeneous Inconel 182 wide chamfer weld from EDF, with approximately uniform orientations, as used in [16]. Figure 9 shows a photograph and a schematic diagram of the experimental pulse-echo set-up used for inversion. As opposed to Ménard *et al.* [16], we used contact measurements, with the array (5 MHz 64-element; pitch 0.6 mm, element: 0.5 mm) placed on the surface of the weld, centrally with respect to the middle defect. This is the simplest test case, where the target is a single orientation angle. Hence, this is a single-stage inversion with the objective function depending only on one variable. The gradient of the objective function was calculated as a sum of contributions from all cells.

The forward model underpinning inversion discretized the domain into 2×2 mm cells, with each edge seeded with 10 nodes. For the parent material, we used $\rho_{\text{parent}} = 8260 \text{ kg m}^{-3}$,

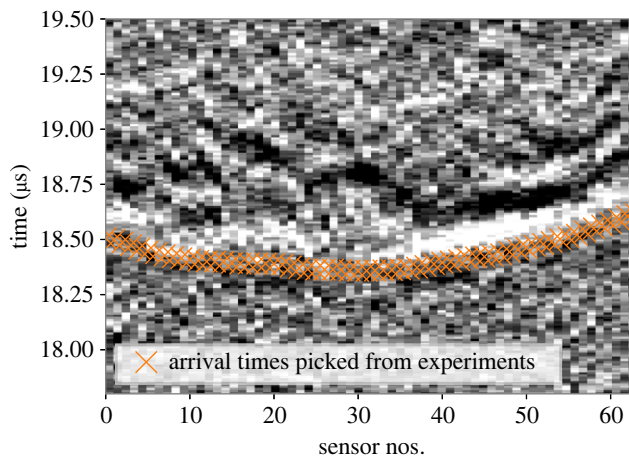


Figure 10. Extracting time-of-arrival as the target for inversion from experimental data (the map combines time traces across all sensors from single-transducer excitation; the markers show the identified times-of-arrival).

$c_{11} = 260$ GPa, $c_{12} = 100$ GPa and $c_{44} = 80$ GPa. The weld material had the same density, and the elasticity matrix with $c_{11} = 265$ GPa, $c_{13} = 140$ GPa, $c_{33} = 247$ GPa, $c_{44} = 112$ GPa and $c_{66} = 86$ GPa. The latter was determined from the data provided by EDF.

As signalled in the preceding section, extracting arrival times from experimental data collected on coarse-grained structures is challenging. Scattering noise and the low amplitude of the backwall reflection preclude the use of simple automated methods, such as following the maximum of the Hilbert envelope. Figure 10 shows an example of received signals across all array elements for a single excitation. The orange markers represent extracted arrival time, whose slight variations may affect the inversion significantly. The discussion about algorithms for determining first arrival times is outside this paper's scope. Owing to their complicated nature, we used manual tracking based on subsequent visual analyses of time traces across all receivers excited from one source at a time. Developing a more robust and automated procedure needs to be pursued in future.

The experimental times-of-flight across the weld showed a small difference between the array's two extremities, e.g. elements 1 and 64 (with the same element transmitting and receiving). Assuming that the sample was composed of grains with uniform orientation (confirmed by available data and previously published work [16]), one would expect to observe equal arrival times. While the difference could be related to several factors, we modelled it by distributing the sources along a sloped line at a slight angle ($\approx 1.2^\circ$) with respect to the backwall. This allowed us to execute the inversion, assuming that a single value for grain orientation can represent the considered material.

Figure 11 shows the inversion results based on experimental time traces recorded at CEA. Weld map tomography started from a uniform vertical orientation of grains and converged to -8.11° after nine iterations. Determining the exact angular error is impossible for this example. However, the metallography performed on an off-cut of the same weld suggests the dominant orientation of -8° , which confirms that the inversion reached a realistic value. However, quantitative error analysis is possible for arrival times, as depicted in figure 11*b*. The mean absolute error was $0.021 \mu\text{s}$, the median absolute error $0.015 \mu\text{s}$ and the maximum absolute error $0.093 \mu\text{s}$. The highest errors are between transducers close to array extremities, suggesting that the geometry of the sample was not captured correctly. Despite that, the result of the tomography is close to the reference provided for the mock-up. We attribute this robustness to the fact that only a single orientation angle is sought in this example, minimizing the impact of locally mismatched definitions.

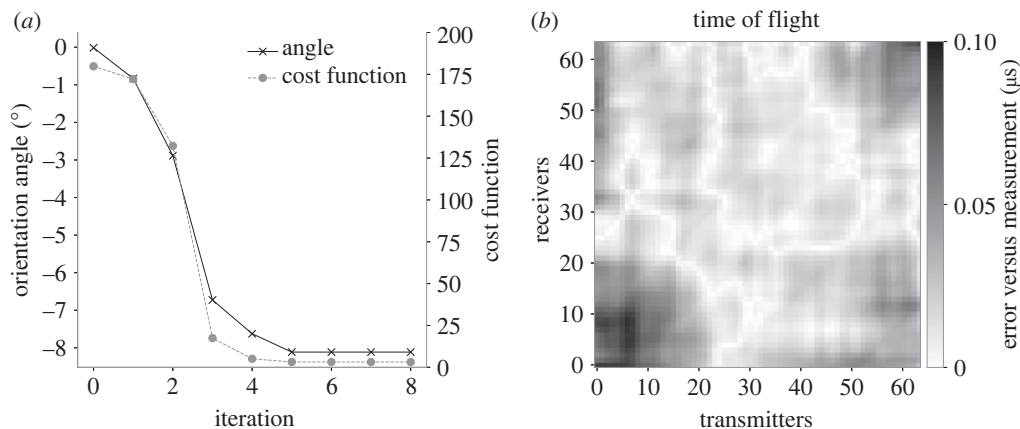


Figure 11. Results of weld map tomography for the uniform Inconel weld: (a) cost function and orientation angle at different iterations and (b) a map of absolute ToF errors ($|\tau_{\text{original}} - \tau_{\text{inversion}}|$).

The practically relevant verification of the inversion is to perform imaging with the obtained material information. We used the classical total focusing method (TFM) [39], which synthetically focuses the signal in transmission and reception at every point in the region of interest. The imaging function $I(x, y)$ is defined as

$$I(x, y) = \left| \sum_{i=1}^M \sum_{j=1}^M [u_{ij}(\tau_{i \rightarrow (x,y)} + \tau_{j \rightarrow (x,y)}) + J\hat{u}_{ij}(\tau_{i \rightarrow (x,y)} + \tau_{j \rightarrow (x,y)})] \right|, \quad (5.1)$$

where (x, y) is the position in the physical domain, M is the number of elements in the array, u_{ij} is the response measured by element i with element j exciting, $\tau_{i \rightarrow (x,y)}$ is the traveltime between element i and point (x, y) (called a delay law) and \hat{u}_{ij} is the Hilbert transform of u_{ij} [40]. Note that the analytic signal $u_{ij} + J\hat{u}_{ij}$ is the typical output of ‘hilbert’ functions in the major scientific computing packages (e.g. *scipy*, *MATLAB*).

Figure 12 shows TFM images calculated using the dataset used for inversion. For comparison, we show the ‘conventional’ image evaluated under the assumption that the material is isotropic and the ‘updated’ image, which uses the results from inversion presented earlier. Both images are normalized to the maximum amplitude of the middle defect (individually). The circular markers indicate the true positions of side-drilled holes. To quantify the difference between the two images we calculated the SNR using the following procedure. First, the algorithm searched for a maximum within ± 5 mm of the known defect position. Then, we defined a signal region I_{sig} (solid line square), and the noise probe region I_{noise} (dashed line square), which were squares with sides 4 and 10 mm, respectively, centred around the maximum. The SNR was calculated using the maximum amplitude in the signal region and the root-mean-square amplitude in the noise region (not including the signal part)

$$\text{SNR} = 20 \log_{10} \frac{\max(I_{\text{sig}})}{\sqrt{I_{\text{noise}}^2}}. \quad (5.2)$$

The results in figure 12 confirm that using the correct material information reduces the overall level of noise in the image, improves the location of defect signatures (by 3.2, 1.1 and 4.6 mm from the left to the right defect), and enhances the SNR (by 5.9, 14.4 and 7.5 dB, respectively), making defect signatures more focused and better pronounced. Moreover, it confirms the effectiveness of weld map tomography and its benefits to imaging.

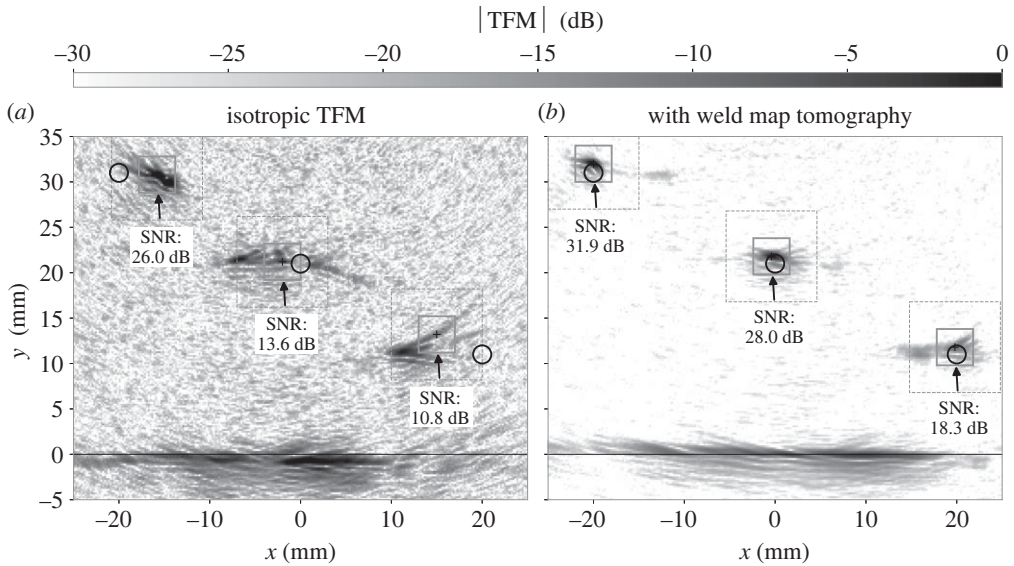


Figure 12. Total focusing method imaging with material information for the Inconel weld: (a) standard isotropic delay laws and (b) with delay laws coming from weld map tomography. The SNRs of the defect signatures are included for comparison.

(b) Weld with a non-planar backwall

In this section, we test weld map tomography on a 316L stainless steel industrial mock-up from EDF. The sample has a typical V-shaped configuration and a non-planar backwall profile. Two target defects (side-drilled holes of 1 mm diameter) located at the centreline (approx. 10 and 20 mm from the top surface—see figure 15b for reference) provide a reference for imaging. Compared with the previous sample, the growth of complexity is twofold: first, we update individual local grain orientations; second, the backwall profile has a non-planar profile.

The measurements used a contact 5 MHz 64-element array with 0.6 mm pitch and 0.5 mm element width, positioned centrally above the weld. Extracting backwall-reflected arrival times from the recorded dataset required a semi-manual procedure described in §5a. An illustrative map of arrival times extracted from time traces recorded by the array after one of the elements was excited is shown in figure 13.

As presented in this paper, weld map tomography only focuses on determining local grain orientations, that is, material information. Thus, the geometry was fixed and assumed known, which is a straightforward task for a mock-up. To maintain the brevity of this paper, we omit the simplified representation of the backwall profile.

The forward model represents the pulse-echo set-up using the mirrored domain as in the previous examples, with a modification to account for the non-planar backwall. Mirroring the domain with respect to the $y = 0$ mm makes the backwall nodes not coincident anymore. We address it by joining corresponding original and mirrored backwall nodes with single zero-propagation time edges. This arrangement allows the SRP algorithm to recognize the physical connection between corresponding backwall nodes of original and mirrored domains, despite them being separate in the computational domain.

The forward model used in the inversion discretized the domain into 2×2 mm cells, seeded with 10 nodes per each side. For the parent material, we used $\rho_{\text{parent}} = 7800 \text{ kg m}^{-3}$, $c_{11} = 276 \text{ GPa}$, $c_{12} = 102 \text{ GPa}$ and $c_{44} = 87 \text{ GPa}$. The weld material had the same density, and the elasticity matrix with $c_{11} = 245 \text{ GPa}$, $c_{12} = 110 \text{ GPa}$, $c_{13} = 145 \text{ GPa}$, $c_{22} = 253 \text{ GPa}$, $c_{33} = 221 \text{ GPa}$, $c_{44} = 110 \text{ GPa}$ and $c_{66} = 75 \text{ GPa}$ [17]. We exercised the two-stage inversion, identical to that in §4, where the first




Figure 13. Extracting time-of-arrival as the target for inversion from experimental data (the map combines time traces across all sensors from single-transducer excitation (no. 21); the markers show the identified times-of-arrival).

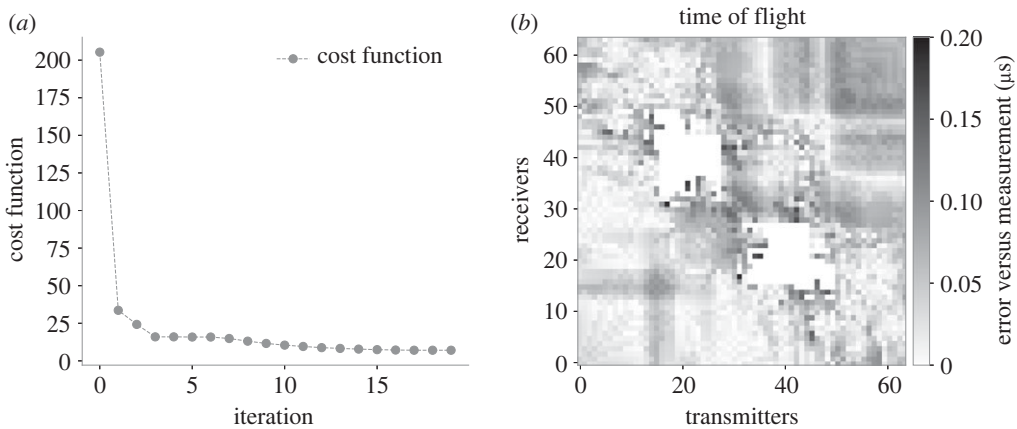


Figure 14. Results of weld map tomography for the stainless steel weld with a non-planar backwall: (a) cost function and orientation angle at different iterations; and (b) a map of absolute ToF errors ($|\tau_{\text{original}} - \tau_{\text{inversion}}|$).

stage looks for the optimal dominant orientation angle for the generalized Ogilvy model, and the second stage updates local orientations to decrease the cost function further until convergence.

Figure 14 shows the results of the inversion, which started from an Ogilvy map with 0° dominant orientation. After six iterations, the minimization converged to -20.5° and moved to the second stage, where individual angles were further updated, up to a total of 19 iterations. The evolution of the cost function clearly suggests that the dominant grain orientation had the most significant effect on the inversion.

The arrival time errors are higher than in the previous example, with the mean absolute error of $0.034 \mu\text{s}$, the median absolute error of $0.029 \mu\text{s}$ and the maximum absolute error of $0.185 \mu\text{s}$. It is worth noting that the ‘white spots’ in figure 14b indicate that for these transducer pairs, it was impossible to extract arrival times from the measured data. We expect this to be related to the more complex shape of the backwall and the very low amplitude of some rays, owing to the grain orientation layout, as mentioned, e.g. in [41]. While such a reduction in the available data is undesired for the inversion, this case confirms that weld map tomography does not require the entire dataset to work. We also observe that the maximum errors occurred for transducer pairs

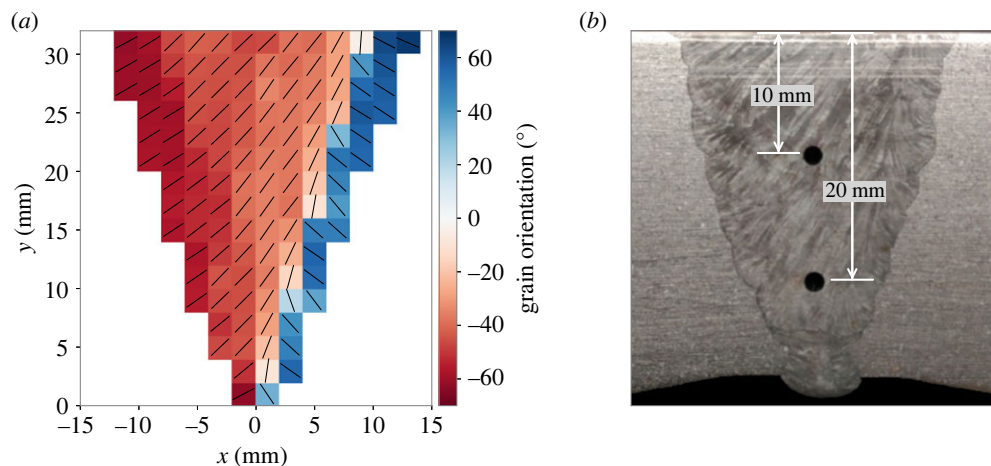


Figure 15. Results of weld map tomography for the stainless steel weld with a non-planar backwall: (a) updated local grain orientations and (b) the metallography image of the mock-up.

close to the undefined arrival time regions, bound by high uncertainty. Nevertheless, the mean error is still low and represents a practically promising performance. In figure 15, we offer a visual comparison of the updated orientations with the metallography of the mock-up, confirming the effectiveness of the inversion.

The final verification involves calculating the TFM image using the same recorded dataset with and without material information provided by weld map tomography, as shown in figure 16, which also includes the image computed with orientations from the first stage inversion only. Updated material information reduces the noise in the image and improves the prominence and focus of defect signatures. While the noise is not eliminated entirely, the SNRs of the target defects improve by 12.3 and 4.4 dB, respectively. There is no improvement between the first stage and the final orientation map, which could be expected from the shape of the cost function. In this case, the first stage of inversion would suffice, but such a conclusion can only be made after inversion and imaging. We acknowledge artefacts close to the root of the weld; however, their origin, and ways of reducing them, remain the subject of future work. Nevertheless, our proposed inversion technique outputs local orientations agreeing well with the metallography image and enhancing the ultrasonic image of target defects, despite the complex backwall geometry.

(c) ADVISE mock-up

The final example considered is an austenitic stainless steel mock-up manufactured for the project sponsoring this work. Project partners designed and manufactured a weld of industrial relevance, but with a careful recording of welding parameters, materials and other details necessary for weld formation and phenomenological MINA [2] modelling.

The mock-up was manufactured by joining two stainless steel pipes (1.4550) with a wall thickness of 39 mm (outside diameter: 517 mm)—figure 17. In the flat position, the circumferential weld was manufactured using the shielded metal arc method with ESAB OK 61.30 filler (1.4316). The cylindrical specimen was then sectioned into several parts along the circumference (figure 17), which served other research tasks. In this paper, one section was used for characterization (mock-up A), and another in subsequent imaging (mock-up B). The relevance of characterization on A for imaging on B follows the common assumption that the variation of the grain orientations in the welding direction is small.

A separate section of the mock-up was used for material examinations, both traditional metallography and EBSD. For the EBSD measurement, the weld was further sectioned into four

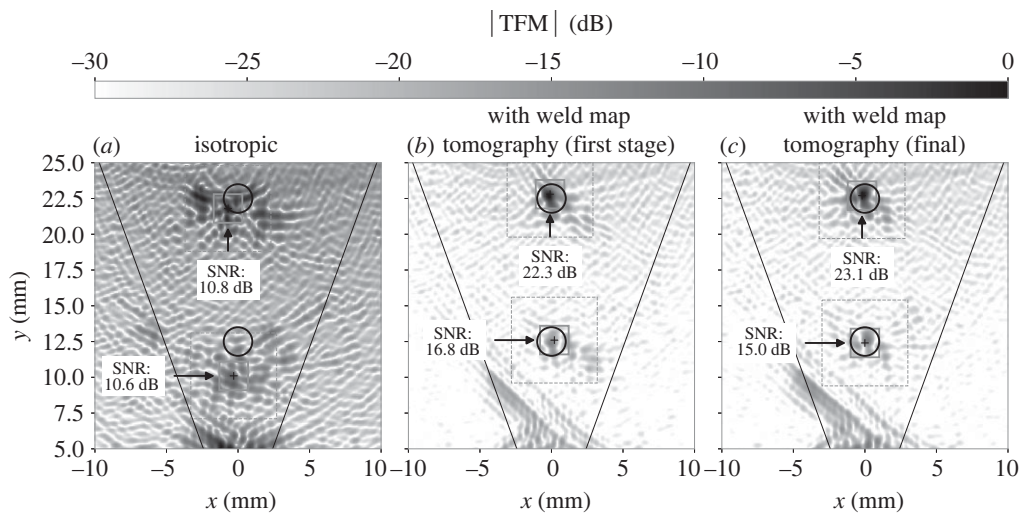


Figure 16. Total focusing method imaging with material information for the stainless steel weld with a non-planar backwall: (a) standard isotropic delay laws; (b) with delay laws coming from the first stage of weld map tomography; and (c) with delay laws coming from the final updated weld map. The SNRs of the defect signatures are indicated for comparison.




Figure 17. The ADVISE mock-up: a photograph and schematic drawings of a cut-out section.

parts, and each was covered by a grid of smaller measurement areas in the microscope. Finally, the obtained dataset, comprising over 100 ‘tiles’ was stitched and analysed using Dream.3D [42], an open-source package for microstructure synthesis, to deliver a numerical microstructure that can be visualized, analysed and subsequently used in grain-scale modelling or the verification of experimental inversion.

To facilitate the measurements on both A and B mock-ups, their top and bottom surfaces were machined to achieve flat surfaces (figure 18). This modification helped ensure good contact between the transducer and the sample and allowed us to focus on the inversion rather than the transducer issues. In addition, material evaluations at the University of Stuttgart indicated several natural defects within the weld. Additional targets (holes) were drilled along the chamfer of mock-up B at CEA.

Array measurements using a 2.25 MHz 64-element array with a pitch of 0.75 mm (in contact with the sample, directly on top of the weld, coupled through silicone gel) were taken at KTU using the Gekko system from Eddyfi Technologies. The acquisitions were processed using the

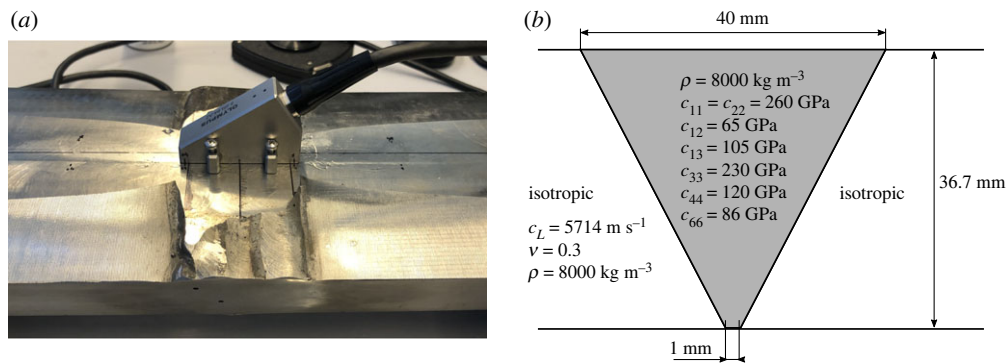


Figure 18. ADVISE mock-up and the transducer over mock-up B: (a) a photograph and (b) dimensions with material properties.

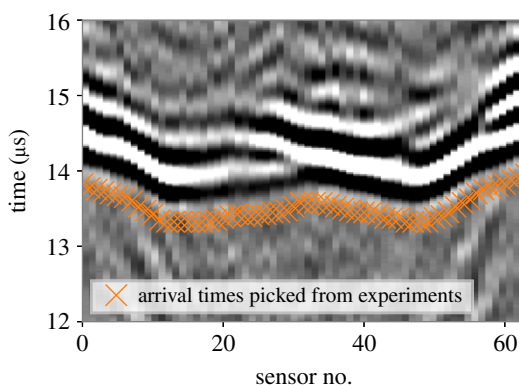


Figure 19. Extracting time-of-arrival as the target for inversion from experimental data (the map combines time traces across all sensors from single-transducer excitation (no. 33); the markers show the identified times-of-arrival).

semi-manual time-of-arrival extraction procedure described in §5a, yielding the target for the weld map tomography algorithm. An illustrative map of time traces for one transmitter, together with the extracted arrival times, is shown in figure 19.

The forward model used in inversion assumed known material properties of the sample and known geometry. We used the EBSD dataset to execute homogenization and determine the macroscopic transversely isotropic elastic tensor based on the grain structure and the crystal elastic tensor. These parameters are listed in figure 18.

Two-stage weld map tomography, analogous to that in §4, converged after 26 iterations, with the first stage (updating the dominant orientation of the generalized Ogilvy map) taking four iterations. Figure 20 shows the cost function and the final map of absolute time of flight errors (mean absolute error: $0.044 \mu\text{s}$, median absolute error: $0.036 \mu\text{s}$ and maximum absolute error: $0.193 \mu\text{s}$). To further illustrate the evolution of the ToF map during inversion, we show the initial, the final and the target results for two sources in figure 21. In the illustrated cases, initial ToF errors were up to $0.5 \mu\text{s}$ for some receivers while reducing to less than $0.1 \mu\text{s}$ in the final result.

Figure 22 compares grain orientations determined using weld map tomography and the reconstruction of the EBSD measurement taken at the University of Stuttgart. This visual comparison gives a convincing indication of the effectiveness of the method. A detailed error map, which uses grain orientations extracted from EBSD data, and an error histogram are presented in figure 23. We note that the errors in orientation angles are mainly concentrated around the chamfer, where the EBSD data had a low confidence index. The mean absolute error in

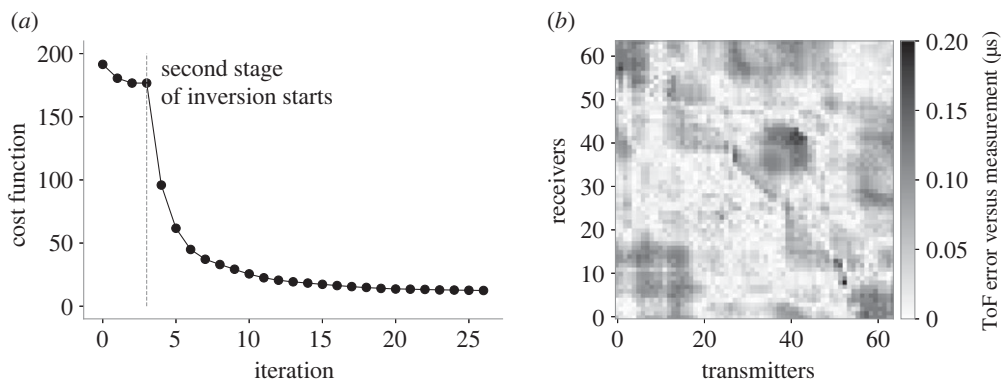


Figure 20. Results of weld map tomography for the stainless steel ADVISE mock-up: (a) cost function and orientation angle at different iterations; and (b) a map of absolute ToF errors ($|\tau_{\text{original}} - \tau_{\text{inversion}}|$).




Figure 21. Evolution of the ToF between the initial guess and the converged weld map tomography output versus the target for the ADVISE mock-up: (a) source at transducer no. 13 and (b) source at transducer no. 54.

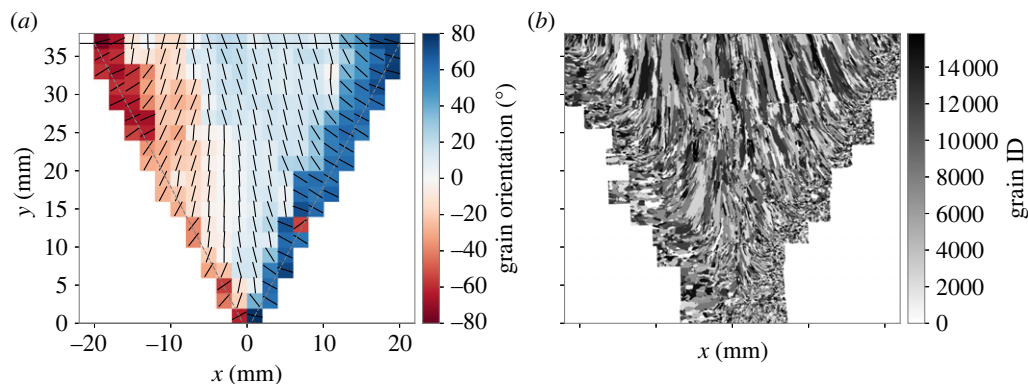


Figure 22. Results of weld map tomography for the stainless steel ADVISE mock-up: (a) updated local grain orientations and (b) microstructure reconstruction based on EBSD measurements.




Figure 23. Absolute orientation error of the weld map tomography output, with respect to orientations extracted from EBSD data; note low confidence in benchmark data around the chamfer: (a) error map and (b) histogram.

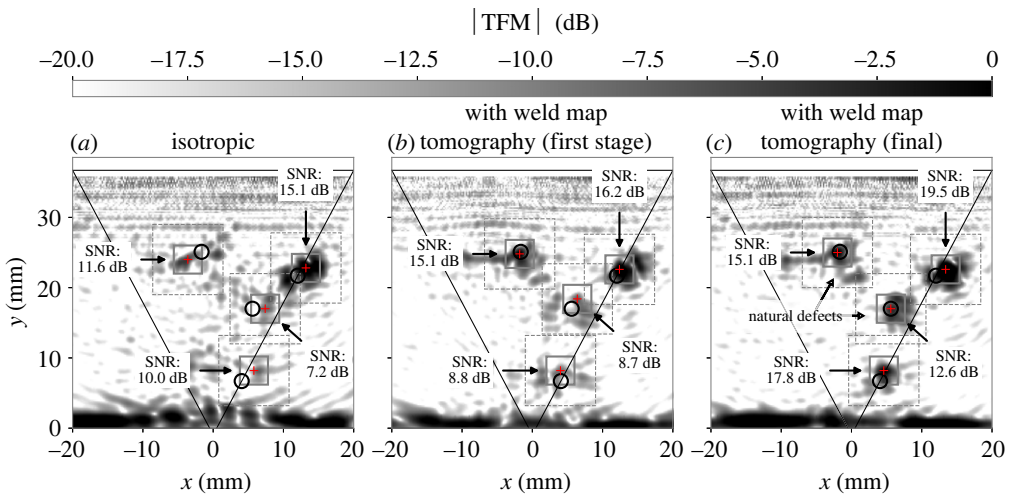


Figure 24. Total focusing method imaging with material information for the ADVISE mock-up: (a) standard isotropic delay laws; (b) with delay laws coming from the first stage of weld map tomography; and (c) with delay laws coming from the final updated weld map. The SNRs of the defect signatures are indicated for comparison.

orientations is 19.2° . If we disregard the cells adjacent to the chamfer (low confidence in the EBSD benchmark), the mean absolute error reduces to 13.8° . Nevertheless, as in previously presented examples, the final confirmation of the efficacy of weld map tomography is illustrated by applying its output to imaging.

Figure 24 compares images calculated for mock-up B. The effect of the proposed inversion technique is clear and convincing. First, the sample had two known target defects (side-drilled holes) along the right chamfer—figure 24 shows that only one is visible in the conventional TFM image. Including material information from inversion raised the bottom chamfer defect from the background noise. The SNRs of the chamfer defects improved by around 4 and 8 dB. Conversely to the example in §5b, the weld map obtained during the first stage of our algorithm provides only partial improvement. Including the second stage is vital in this example, especially for defects closer to the backwall. This observation can be linked to the evolution of the cost function, which experiences a significant drop after the second stage of inversion begins.

In addition to the intentionally introduced side-drilled holes, two natural defect signatures emerged in figure 24*b*: one close to the centreline and another near the chamfer. Post-manufacture X-ray scans of the weld indicated the presence of voids in some parts of the weld, and it was confirmed with ultrasonic imaging. More importantly, the absence of these natural defects' signatures in the conventional image proves the effectiveness of weld map tomography in supporting imaging and is compelling proof of the importance of accounting for material information.

6. Discussion and conclusion

This paper proposed a method for determining local grain orientations in complex, thick welds from ultrasonic array data. The reconstruction followed the ray tomography principle, accounting for the heterogeneous anisotropy of the welds of interest. The inversion used a SRP model as the forward solver, and the paper detailed the principles and construction of the model. We discussed the importance of the *a priori* information about the weld and the advantages of using simple geometrical descriptions for the initial stages of inversion. The algorithm was first showcased using a numerical example based on an industrial inspection scenario, with grain-scale time-domain finite-element simulations acting as a virtual experiment. The results of the inversion showed good agreement with the numerically generated microstructure.

Encouraged by this result, we demonstrated weld map tomography on three mock-ups—two supplied by the industry and one manufactured specifically for this research. The outcome of the method agreed well with destructive material examinations (both metallography and EBSD) and provided the input for updating the delay laws for imaging. The converged mean absolute ToF errors for experimental cases was between 0.021 and 0.044 μs , and mean absolute orientation error of 13.8° (for the last case, which offered orientation ground-truth data). In all examples, TFM images using material information from inversion showed significant improvement in both the location of defect signatures and SNR, confirming the efficacy of the proposed algorithm—between 6 and 14 dB for the uniform orientation case, between 5 and 12 dB for the sample with non-planar backwall, and between 4 and 8 dB for the final ADVISE mock-up. While computational efficiency was not the goal of this research, the presented algorithm achieves good results, with a single iteration taking approximately a few seconds on a personal laptop.

While the results are encouraging, several aspects are still pending investigation. Also, towards the end of the manuscript, it may be helpful to summarize the limitations of our algorithm. From a practical perspective, the main obstacle is the requirement to know the geometry of the sample and the weld precisely—this is only sometimes possible in *in situ* inspections. Second, extracting arrival times from time-domain data is challenging in the presence of high structural noise and may require more time than the actual inversion. While arrival time extraction has received considerable attention in the literature, in geophysics in particular, a robust, automated procedure for NDE imaging is still to be identified. Other remaining challenges, such as the uncertainty analysis, different transducer configurations related to limited accessibility (through-weld cap [26,27], tandem, single-side access), and weld complexities (e.g. buttering), are related to the implementation and deployment of the method and are the subject of our ongoing investigation [43].

Data accessibility. The numeric microstructure was generated using the open-source program Neper: <https://github.com/rquey/neper>. Grain-scale finite element simulation of ultrasound was performed using Pogo (demo version available from: www.pogo.software/). All data associated with considered examples together with Jupyter notebooks allowing to reproduce figures in the paper are publicly available from the Zenodo repository: <https://doi.org/10.5281/zenodo.8030163> [43].

Data provided in the electronic supplementary material [22].

Authors' contributions. M.K.K.: conceptualization, data curation, formal analysis, investigation, methodology, software, validation, visualization, writing—original draft, writing—review and editing; M.J.S.L.: conceptualization, funding acquisition, investigation, methodology, project administration, resources, supervision, writing—original draft, writing—review and editing; V.S.: conceptualization, data curation,

investigation, methodology, validation, writing—original draft, writing—review and editing; F.S.: data curation, investigation, methodology, resources, validation, writing—original draft, writing—review and editing; S.R.: data curation, investigation, methodology, resources, validation, writing—original draft, writing—review and editing.

All authors gave final approval for publication and agreed to be held accountable for the work performed therein.

Conflict of interest declaration. We declare we have no competing interests.

Funding. This research was a part of the ADVISE project (www.advise-h2020.eu) funded from the Euratom Research and Training Programme 2014–2018 under grant agreement no. 755500.

Acknowledgements. Special thanks to EDF R&D and Dr Andreas Schumm for providing mock-ups for examinations and the data underpinning the numerical example and to Dr Peter Huthwaite for inspiring discussions and suggestions at the beginning of the project.

Appendix A. Generalized Ogilvy map

The geometric description of grain orientations proposed by Ogilvy [3] is a simple but powerful method. Throughout recent decades it became the gold standard for modelling local orientations in austenitic stainless steel welds. This paper uses a slightly extended version of the Ogilvy map, accounting for the dominant orientation angle, which is relatively common in published examinations (e.g. figure 9a). What is presented below is conceptually analogous to an approach used in [17].

For completeness, let us start from the classical Ogilvy map [3]. We consider a simple V-weld of thickness a , left and right weld angles α_L and α_R , and weld bead lengths (positive) d_L and d_R . For a domain defined in y - z coordinates with z being vertical, the Ogilvy model requires four parameters: T_L and T_R , representing the tangent of the orientation angles at the left and right chamfers, and η_L and η_R , representing the rate of change of the angle while moving horizontally from the chamfer towards the centreline (where the grains are always vertical). The set-up, the coordinate system and the parameters are illustrated in figure 25. Our convention assumes the lengths, weld angles and η parameters to be always positive. Orientation tangent T_L is expected to be positive, while T_R is negative. It is common to assume symmetry, with $T_L = -T_R$, $\eta_L = \eta_R$, $\alpha_L = \alpha_R$ and $d_L = d_R$.

Grain orientations θ , defined as the rotation from the vertical positive anticlockwise, are calculated using the following expressions

$$\theta(y, z) = \begin{cases} \arctan\left(\frac{T_L|d_L + z \tan \alpha_L|}{y^{\eta_L}}\right) - \frac{\pi}{2}, & \text{if } y < 0 \\ \arctan\left(\frac{T_R|d_R + z \tan \alpha_R|}{y^{\eta_R}}\right) + \frac{\pi}{2}, & \text{if } y > 0 \\ 0, & \text{if } y = 0 \end{cases} \quad (\text{A } 1)$$

The generalized Ogilvy map, as we call it in this paper, adds one more parameter ζ to capture the dominant orientation representing the direction of privileged grain growth during cooling [17]. This parameter is the angle between the original z -axis and the new axis along which the grains would align. In the model, we account for that by temporarily rotating the y - z system by ζ (figure 25b). If $|\zeta|$ is greater than α_L or α_R it should be capped at the weld angle for the coordinate rotation (for some configurations, e.g. very narrow welds, the limit may need to be less than the weld angle). After the rotation, each point in the original y - z reference frame has its auxiliary counterpart in the y^ζ - z^ζ frame. Then, the grain orientations are calculated as

$$\theta(y, z) = \begin{cases} \arctan\left(\frac{T_L|d_L + z^\zeta \tan(\alpha_L - \zeta)|}{(y^\zeta)^{\eta_L}}\right) - \frac{\pi}{2} + \zeta, & \text{if } y^\zeta < 0 \\ \arctan\left(\frac{T_R|d_R + z^\zeta \tan(\alpha_R + \zeta)|}{(y^\zeta)^{\eta_R}}\right) + \frac{\pi}{2} + \zeta, & \text{if } y^\zeta > 0 \\ \zeta, & \text{if } y^\zeta = 0 \end{cases} \quad (\text{A } 2)$$




Figure 25. Schematic diagram for: (a) the classical Ogilvy model and (b) auxiliary rotation of the coordinate system to account for dominant orientation ζ .

When applying the dominant orientation parameter ζ to an originally symmetric weld (with $T_L = -T_R$), one should adjust these parameters to retain relevance to realistic cases. A rule-of-thumb adjustment based on comparing such maps with metallography images suggests multiplying initial T_L and T_R parameters by $(1 + 0.01\zeta)$ and $(1 - 0.01\zeta)$, respectively.

When used in weld map tomography as the first stage of inversion, the Ogilvy map takes the initial T and η parameters equal to one.

References

1. Connolly GD, Lowe MJS, Temple JAG, Rokhlin SI. 2009 The application of Fermat's principle for imaging anisotropic and inhomogeneous media with application to austenitic steel weld inspection. *Proc. R. Soc. A* **465**, 3401–3423. (doi:10.1098/rspa.2009.0272)
2. Moysan J, Apfel A, Corneloup G, 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. (doi:10.1016/S0308-0161(03)00024-3)
3. Ogilvy JA. 1985 Computerized ultrasonic ray tracing in austenitic steel. *NDT Int.* **18**, 67–77. (doi:10.1016/0308-9126(85)90100-2)
4. Chassignole B, El Guerjouma R, Ploix MA, Fouquet T. 2010 Ultrasonic and structural characterization of anisotropic austenitic stainless steel welds: towards a higher reliability in ultrasonic non-destructive testing. *NDT & E Int.* **43**, 273–282. (doi:10.1016/j.ndteint.2009.12.005)
5. Hunter AJ, Drinkwater BW, Zhang J, Wilcox PD, Thompson DO, Chimenti DE. 2011 A study into the effects of an austenitic weld on ultrasonic array imaging performance. In *Review of Progress in Quantitative Nondestructive Evaluation* (eds DO Thompson, DE Chimenti), vol. 30, pp. 1063–1070. San Diego, CA: AIP Conference Proceedings.
6. Mark AF, Fan Z, Azough F, Lowe MJS, Withers PJ. 2014 Investigation of the elastic/crystallographic anisotropy of welds for improved ultrasonic inspections. *Mater. Charact.* **98**, 47–53. (doi:10.1016/j.matchar.2014.09.012)
7. Spies M. 2000 Modeling of transducer fields in inhomogeneous anisotropic materials using Gaussian beam superposition. *NDT & E Int.* **33**, 155–162. (doi:10.1016/S0963-8695(99)00036-5)
8. Langenberg KJ, Hannemann R, Kaczorowski T, Marklein R, Koehler B, Schurig C, Walte F. 2000 Application of modeling techniques for ultrasonic austenitic weld inspection. *NDT & E Int.* **33**, 465–480. (doi:10.1016/S0963-8695(00)00018-9)
9. Gueudre C, Le Marrec L, Moysan J, Chassignole B. 2009 Direct model optimisation for data inversion. Application to ultrasonic characterisation of heterogeneous welds. *NDT & E Int.* **42**, 47–55. (doi:10.1016/j.ndteint.2008.07.003)
10. Fan Z, Mark AF, Lowe MJS, Withers PJ. 2015 Nonintrusive estimation of anisotropic stiffness maps of heterogeneous steel welds for the improvement of ultrasonic array inspection. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **62**, 1530–1543. (doi:10.1109/TUFFC.2015.007018)

11. Zhang J, Hunter A, Drinkwater BW, Wilcox PD. 2012 Monte Carlo inversion of ultrasonic array data to map anisotropic weld properties. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **59**, 2487–2497. (doi:10.1109/TUFFC.2012.2481)
12. Tant KMM, Galetti E, Mulholland AJ, Curtis A, Gachagan A. 2018 A transdimensional Bayesian approach to ultrasonic travel-time tomography for non-destructive testing. *Inverse Probl.* **34**, 095002 (doi:10.1088/1361-6420/aaca8f)
13. Tant KMM, Galetti E, Mulholland AJ, Curtis A, Gachagan A. 2020 Effective grain orientation mapping of complex and locally anisotropic media for improved imaging in ultrasonic non-destructive testing. *Inverse Probl. Sci. Eng.* **28**, 1694–1718. (doi:10.1080/17415977.2020.1762596)
14. Singh J, Tant K, Mulholland A, MacLeod C. 2022 Deep learning based inversion of locally anisotropic weld properties from ultrasonic array data. *Appl. Sci.* **12**, 532 (doi:10.3390/app12020532)
15. Singh J, Tant K, Curtis A, Mulholland A. 2022 Real-time super-resolution mapping of locally anisotropic grain orientations for ultrasonic non-destructive evaluation of crystalline material. *Neural Comput. Appl.* **34**, 4993–5010. (doi:10.1007/s00521-021-06670-8)
16. Ménard C, Robert S, Miorelli R, Lesselier D. 2020 Optimization algorithms for ultrasonic array imaging in homogeneous anisotropic steel components with unknown properties. *NDT & E Int.* **116**, 102327 (doi:10.1016/j.ndteint.2020.102327)
17. Ménard C, Robert S, Lesselier D. 2021 Ultrasonic array imaging of nuclear austenitic v-shape welds with inhomogeneous and unknown anisotropic properties. *Appl. Sci.* **11**, 6505 (doi:10.3390/app11146505)
18. Hormati A, Jovanović I, Roy O, Vetterli M. 2010 Robust ultrasound travel-time tomography using the bent ray model. In *Medical imaging 2010: ultrasonic imaging, tomography, and therapy* (eds J D'hooge, SA McAleavey), vol. 7629, pp. 165–176. San Diego, CA: SPIE.
19. Nowers O, Duxbury DJ, Zhang J, Drinkwater BW. 2014 Novel ray-tracing algorithms in NDE: application of Dijkstra and A* algorithms to the inspection of an anisotropic weld. *NDT & E Int.* **61**, 58–66. (doi:10.1016/j.ndteint.2013.08.002)
20. Nowers O, Duxbury DJ, Drinkwater BW. 2016 Ultrasonic array imaging through an anisotropic austenitic steel weld using an efficient ray-tracing algorithm. *NDT & E Int.* **79**, 98–108. (doi:10.1016/j.ndteint.2015.12.009)
21. Virtanen P *et al.* 2020 SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* **17**, 261–272. (doi:10.1038/s41592-019-0686-2)
22. Kalkowski MK, Lowe MJS, Samaitis V, Schreyer F, Robert S. 2023 Weld map tomography for determining local grain orientations from ultrasound. Figshare. (doi:10.6084/m9.figshare.c.6761807)
23. Fletcher R, Reeves CM. 1964 Function minimization by conjugate gradients. *Comput. J.* **7**, 149–154. (doi:10.1093/comjnl/7.2.149)
24. Hager WW, Zhang H. 2006 A survey of nonlinear conjugate gradient methods. *Pac. J. Optim.* **2**, 35–58.
25. Huthwaite P. 2012 *Quantitative imaging with mechanical waves*. PhD thesis, Imperial College London.
26. Long R, Russell J, Cawley P. 2010 Through-weld ultrasonic phased array inspection using full matrix capture. In *AIP Conf. Proc., Kingston, RI, July 2009*, vol. 1211, pp. 918–925. San Diego, CA: AIP Conference Proceedings.
27. Russell J, Long R, Duxbury D, Cawley P. 2012 Development and implementation of a membrane-coupled conformable array transducer for use in the nuclear industry. *Insight Non-Destruct. Test. Cond. Monit.* **54**, 386–393. (doi:10.1784/insi.2012.54.7.386)
28. Nageswaran C, Carpentier C, Tse YY. 2009 Microstructural quantification, modelling and array ultrasonics to improve the inspection of austenitic welds. *Insight Non-Destruct. Test. Cond. Monit.* **51**, 660–666. (doi:10.1784/insi.2009.51.12.660)
29. Ploix MA, Guy P, Chassignole B, Moysan J, Corneloup G, Guerjouma RE. 2014 Measurement of ultrasonic scattering attenuation in austenitic stainless steel welds: realistic input data for NDT numerical modeling. *Ultrasonics* **54**, 1729–1736. (doi:10.1016/j.ultras.2014.04.005)
30. Gueudré C, Mailhé J, Ploix MA, Corneloup G, Chassignole B. 2019 Influence of the uncertainty of elastic constants on the modelling of ultrasound propagation through multi-pass austenitic welds. impact on non-destructive testing. *Int. J. Press. Vessels Pip.* **171**, 125–136. (doi:10.1016/j.ijpvp.2019.02.011)
31. Huthwaite P. 2014 Accelerated finite element elastodynamic simulations using the GPU. *J. Comput. Phys.* **257**, 687–707. (doi:10.1016/j.jcp.2013.10.017)

32. Chassignole B. 2000 *Influence of the metallurgical structure of austenitic stainless steel welds on ultrasonic non destructive testing*. These de doctorat Lyon, INSA.
33. Quey R, Dawson PR, Barbe F. 2011 Large-scale 3D random polycrystals for the finite element method: generation, meshing and remeshing. *Comput. Methods Appl. Mech. Eng.* **200**, 1729–1745. (doi:10.1016/j.cma.2011.01.002)
34. Huang M, Sha G, Huthwaite P, Rokhlin SI, Lowe MJS. 2021 Longitudinal wave attenuation in polycrystals with elongated grains: 3D numerical and analytical modeling. *J. Acoust. Soc. Am.* **149**, 2377–2394. (doi:10.1121/10.0003955)
35. St-Onge A. 2011 Akaike information criterion applied to detecting first arrival times on microseismic data. In *SEG Technical Program Expanded Abstracts 2011, SEG Technical Program Expanded Abstracts*, pp. 1658–1662. Society of Exploration Geophysicists.
36. Hill R. 1952 The elastic behaviour of a crystalline aggregate. *Proc. Phys. Soc. A* **65**, 349 (doi:10.1088/0370-1298/65/5/307)
37. Williamson PR, Worthington MH. 1993 Resolution limits in ray tomography due to wave behavior: numerical experiments. *Geophysics* **58**, 727–735. (doi:10.1190/1.1443457)
38. Li C, Duric N. 2013 Resolution limitation of travel time tomography: beyond the first Fresnel zone. In *SPIE medical imaging* (eds JG Bosch, MM Dooley), p. 86751D. Lake Buena Vista (Orlando Area), FL: SPIE.
39. Holmes C, Drinkwater BW, Wilcox PD. 2005 Post-processing of the full matrix of ultrasonic transmit–receive array data for non-destructive evaluation. *NDT & E Int.* **38**, 701–711. (doi:10.1016/j.ndteint.2005.04.002)
40. Marmonier M, Robert S, Laurent J, Prada C. 2022 Real-time 3D imaging with Fourier-domain algorithms and matrix arrays applied to non-destructive testing. *Ultrasonics* **124**, 106708 (doi:10.1016/j.ultras.2022.106708)
41. Connolly GD. 2009 *Modelling of the propagation of ultrasound through austenitic steel welds*. PhD thesis, Department of Mechanical Engineering, Imperial College London.
42. Groeber MA, Jackson MA. 2014 DREAM.3D: a digital representation environment for the analysis of microstructure in 3D. *Integr. Mater. Manuf. Innov.* **3**, 56–72. (doi:10.1186/2193-9772-3-5)
43. Kalkowski MK, Lowe MJS, Samaitis V, Schreyer F, Robert S. 2023 Dataset and code for: Weld map tomography for determining local grain orientations from ultrasound. *Zenodo*. (doi:10.5281/zenodo.8030163)