



HAL
open science

Weld map tomography for determining local grain orientations from ultrasound

Michal Kalkowski, Michael Lowe, Vykintas Samaitis, Fabian Schreyer, Sébastien Robert

► **To cite this version:**

Michal Kalkowski, Michael Lowe, Vykintas Samaitis, Fabian Schreyer, Sébastien Robert. Weld map tomography for determining local grain orientations from ultrasound. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 2023, 479 (2277), pp.20230236. 10.1098/rspa.2023.0236 . cea-04316187

HAL Id: cea-04316187

<https://cea.hal.science/cea-04316187>

Submitted on 30 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Weld map tomography for determining local grain orientations in welds from ultrasound

Michał K. Kalkowski^{a,b,*}, Michael J. S. Lowe^a, Vykintas Samaitis^c, Fabian Schreyer^d, Sebastien Robert^e

^a*Mechanical Engineering, Imperial College London, SW7 2AZ, UK*

^b*Institute of Sound and Vibration Research, University of Southampton, SO17 1BJ, UK*

^c*Prof. K.Baršauskas Ultrasound Research Institute, Kaunas University of Technology, K. Baršausko st. 59, Kaunas 51423, Lithuania*

^d*Materials Testing Institute (MPA) University of Stuttgart. Pfaffenwaldring 32, 70569 Stuttgart, Germany*

^e*CEA, LIST, Gif-sur-Yvette F-91191, France*

Abstract

The complex structure of inhomogeneous (e.g. austenitic stainless steel) welds poses a long-standing challenge in ultrasonic non-destructive testing. Elongated grains with spatially varying dominant orientations curve ultrasonic beams hindering the interpretation of inspection data. One way to tackle this problem is to include material information in imaging and signal analysis; however, such information is most often only gathered using destructive methods. This paper reports the development of a weld inversion strategy determining local grain orientations, based on a ray tomography principle. The considered approach does not rely on a macroscopic weld description but may incorporate it to facilitate inversion. Thanks to that, it is more general than other available approaches, but allows for reducing the complexity by accounting for known information at the same time. The methodology is demonstrated on both numerical and experimental examples. The experimental work focuses on mock-up samples from the nuclear industry and a mock-up specifically manufactured for the project sponsoring this research. The ‘ground truth’ for the latter comes from an EBSD evaluation, which is the most accurate (yet destructive) examination technique available. The paper concludes by presenting the benefit of the material information obtained from the reported inversion model for ultrasonic imaging.

1. Introduction

Ultrasonic inspection of inhomogeneous welds remains 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 complexity of ultrasound propagation. First, the ultrasonic beam no longer travels across straight paths, but follow the minimum propagation time principle resulting in curved,

*Corresponding author

Email address: m.kalkowski@soton.ac.uk (Michał K. Kalkowski)

6 complicated rays. Consequently, locating a defect based on the time of flight is no longer viable without
7 detailed material information and a suitable propagation model. The second effect is grain scattering. When
8 passing through the subsequent grain boundaries the ultrasonic energy is scattered leading to the loss of
9 amplitude of the main beam (attenuation) and elevated structural noise. Other issues include, e.g. beam
10 splitting, that is, emergence of a secondary beam at a grain boundary that gives rise to two strong scattered
11 waves.

12 The above mentioned issues may be at least partially accounted for if material information is available.
13 Such knowledge allows for correcting ultrasonic images, which could then consider true propagation times.
14 It would also help interpret conventional inspection acquisitions by revealing true beam paths and potential
15 scatterer locations [1]. Material information can be obtained from classical examination (e.g. metallography),
16 microscope imaging (EBSD) or assuming a certain structure of the weld (e.g. the gold-standard MINA model
17 [2] or the Ogilvy map [3]). Examination methods are destructive, costly and have no outlook to be applied
18 to components on the plant. On the other hand, simplified weld description models cover limited ranges of
19 cases and their parameters need to be tuned. An alternative is to use inversion based on ultrasound and
20 extract the orientation information from the ultrasonic measurements directly. This concept was one of the
21 premises underpinning the ADVISE project addressing the multi-faceted problem of ultrasonic inspection
22 of complex structured materials, of which this research is a part.

23 It is widely accepted that to support ultrasonic testing, one does not need to consider detailed grain
24 topology [4–6]. Instead, a macroscopic description is sufficient, where the grains are locally homogenised
25 with a uniform orientation over a chosen grid (often between 1 and 2 mm). Columnar grains with similar
26 preferential orientations can be approximated as a transversely isotropic material with the axis of mate-
27 rial symmetry co-linear with the dominant orientation. Historically, macroscopic weld descriptions were
28 developed first to improve the understanding of inspection results. While no inversion was performed, their
29 parameters were assumed based on the examination data and practical experience. Such geometric models
30 include the famous Ogilvy map [3], the layered arrangement [7] or the more basic approximation using only
31 two angles by Langenberg et al. [8]. A refined orientation layout, linked to welding parameters came with
32 the MINA model [2] which draws the information from the welder’s notebook and captures the solidification
33 processes in a phenomenological manner. All these approaches required assuming some if not all parameters,
34 especially for a weld of which little was known.

35 The first attempts to use ultrasonic measurements for determining orientations aimed at updating weld
36 description parameters mentioned above. For instance, Gueudre et al. [9] used evolutionary optimisation
37 to guess the ‘arbitrary’ MINA parameters (remelting rates and tilt angles) based on the echodynamic
38 curve. One problem with this approach is the uncertainty related to transducer coupling when measuring
39 amplitudes, as well as the limited range of configurations it could cover. The same parameters were later
40 determined from the time-of-flight maps by Fan et al. [10], also based on a genetic algorithm. Inversion

41 based on Bayesian updating and the Ogilvy map was conducted at the University of Bristol [11], where this
42 approach estimated parameters of the Ogilvy model. Weld description constraints were removed in another
43 Bayes-based work [12, 13], where the orientations were identified on a per-cell basis. While delivering
44 promising results, the Bayesian approaches incurred significant computational costs which was deemed as
45 one of their key shortcomings.

46 Recent advances in the availability of high computing power led to the development of approaches based
47 on soft computing. Singh et al. proposed employing deep neural networks to determine the orientation
48 maps based on a large simulated database of times of flight through different grain layouts [14, 15]. While
49 this approach gives the promise of real-time application, it still requires a computationally expensive train-
50 ing stage, and does not exploit the physics of wave propagation through the heterogeneous material. An
51 alternative method which takes the advantage of modern optimisation methods was developed by Menard
52 et al. [16, 17]. Their focus is on optimising the amplitude of the TFM image around a postulated defect
53 by adjusting the parameters of a macroscopic weld description. While the main outcome of the method is
54 the enhanced TFM image (with improved defect signature location and the signal-to-noise ratio), a grain
55 orientation map is delivered as a by-product.

56 Although the developments reported in the previous paragraph provided solutions to some challenges,
57 they were often based on ‘black-box’ optimisers, requiring significant calculation cost, pre-identified defect
58 area, or limited to a given macro weld description. To the authors’ knowledge, physics-based inversion was
59 not fully exploited for the applications of interest. Hence, we developed a ray inversion-based method for
60 characterisation of inhomogeneous welds based on ultrasonic array data acquired at the surface. Compared to
61 the prior work, this is a more general inversion approach with a stronger link to physics of wave propagation.
62 Time-of-flight tomography is a known method applied in a variety of contexts, including the biomedical field,
63 in e.g. breast tomography [18]. However, the materials considered in those problems are locally isotropic
64 which is not the case in complex welds. To account for the material complexity and other practical differences,
65 e.g. the sensor arrangement, we extended the classical ray inversion, and developed and validated a new
66 approach, suited to inhomogeneous welds, referred to as weld map tomography.

67 The different aspects of the developed method are presented in more detailed in the remainder of this
68 article. In the next section, we outline the Fermat principle-based forward model underpinning the inversion.
69 Then, we recall the fundamentals of ray tomography for structures comprising isotropic materials and present
70 the extension enabling local anisotropy to be accounted for. This is followed by discussing the details of the
71 implementation of inversion for welds, including our two-stage approach and regularisation. The description
72 is followed by examples - first a numerical example based on grain-scale finite element simulation of a real
73 weld case from EDF; second, a set of experimental examples of industry relevant welds - two mock-ups
74 from EDF and a weld specifically manufactured and characterised for the ADVISE project - based contact
75 transducer array measurements.

2. Forward model

Weld map tomography matches the simulated ultrasound time-of-flight 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, extremely versatile with regard to the material structure and geometry. An alternative, offering potentially more efficient solution is to use a ray tracing model. Previous research on ultrasound in inhomogeneous welds [1, 10] used a ray stepping algorithm following 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 in an efficient manner. One needed to compute a large range of shooting angles and interpolate to obtain the time-of-arrival at desired location.

A crucial functionality of a forward model in this context is to determine the time-of-flight between transducers directly, without the need for interpolation. This can be achieved by utilising 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. Nowers et al. [19, 20] suggested using various incarnations of the Dijkstra method to model ultrasound in welds. The present model draws from that contribution but with some modifications to improve efficiency. Given its origins, the forward model used here is called the shortest ray path model (SRP).

Nowers et al. [19, 20] suggest using a randomly distributed set of nodes and a selected search radius to maximise 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 arrangement of seeds in the forward model to maintain consistency across different stages of inversion. For the ray stepping model, 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. 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.

Fig. 1 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 for the ray model they are discretised in the same way - 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. Fig. 2 schematically depicts this setup.

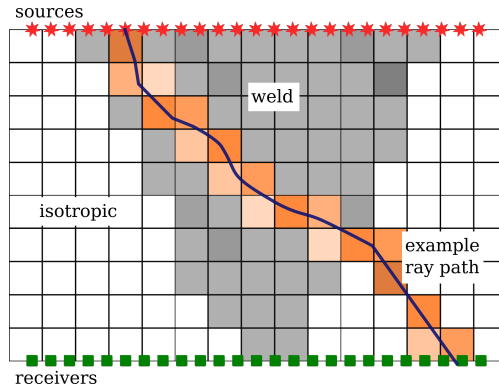


Figure 1: Computational representation of a weld in the SRP model. The orange-coloured cells represent those crossed by the depicted ray.

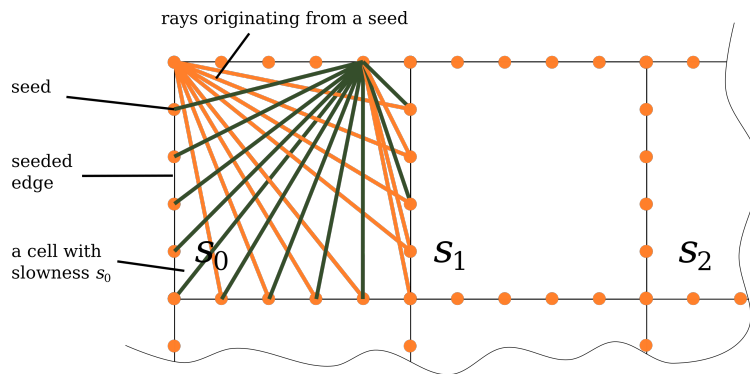


Figure 2: Seeding of cells in the SRP model. Two sets of cell rays are shown, for clarity, rays co-linear with the horizontal and vertical grid lines are not drawn.

110 The SRP model looks for the shortest propagation time between two chosen points. The propaga-
111 tion times across cells are calculated from the group slowness at given combination angles. The combina-
112 tion angles are a simple function of the local grain orientations and incident angles. Those propagation
113 times are first computed for the whole model and fed into a shortest ray path algorithm (we used the
114 `scipy.sparse.csgraph` [21] implementation).

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

123 The convergence study is not reported here for brevity, but for most considered cases, having 10 seeds for a
124 2 mm grid delivered good accuracy when compared with time-domain finite element simulations. This model,
125 despite its high efficiency and versatility also has several limitations, related to its underpinning principle -
126 the shortest propagation time between two points, not accounting for amplitude or energy. Consequently,
127 it is unable to simulate beam splitting and secondary arrival times which sometimes may carry most of the
128 energy. Further, it may output null paths, that is return ray paths associated with negligible energy transfer.
129 However, these limitations do not affect the inversion algorithm as the inversion relies on measured arrival
130 times. If arrival time for a certain transducer pair cannot be seen in measurements, it is not included in the
131 inversion.

132 3. Ray inversion

133 This section covers the main elements of the ray inversion algorithm. It starts with a description of the
134 better established methodology developed for materials with heterogeneous isotropy. The second subsection
135 introduces the modifications needed to allow for anisotropy and variable grain orientations. The final part
136 gives details about the implementation of this inversion model for an inhomogeneous weld.

137 3.1. *Heterogeneous isotropic medium*

The inversion model developed in the ADVISE project is motivated by the paper by Hormati et al. [18],
who used a bent-ray simulation to construct an ultrasonic breast tomography method. The fundamental
concept behind acoustic tomography is to recover local unknown properties from the characteristics of
sound propagation through a medium. In a heterogeneous medium (even if locally isotropic) the sound
takes generally curved paths between two points, depending on the local properties. It is convenient to use

slowness when describing propagation time, and in this case we focus on group slowness, as the method is based on tracking arrival times of finite wave packets generated by the transmitters. A generic time-of-flight is the integral of the group slowness along the propagation path

$$\tau = \int_{\Gamma} s \, dl \quad (1)$$

138 where τ is the time, Γ is the propagation path, s is the slowness and dl is the infinitesimal length along the
139 path.

The discretised domain used in the ADVISE SRP model allows for writing it as a discrete sum

$$\tau = \sum_{i=0}^N s_i l_i \quad (2)$$

where s_i and l_i are the local slowness and the local length of a ray passing a cell. As outlined earlier, the set of possible ray lengths is finite and the same for each cell. It is also worth noting, that a single ray typically passes only a small proportion of cells. Hence, if we cast this formulation into the matrix form

$$\tau = \mathbf{L}\mathbf{s} \quad (3)$$

140 with \mathbf{L} being the ray-length matrix and \mathbf{s} a vector describing slownesses across the whole model, one notices
141 that \mathbf{L} is a sparse matrix.

The inversion minimises a cost function with the aim of determining model parameters that match the simulated times-of-flight with those obtained from a measurement. The cost function used in this work takes the form

$$C = \sum_{i=0}^M (\tau_{i,\text{model}} - \tau_{i,\text{exp}})^2 \quad (4)$$

142 where M is the total number of rays considered (i.e. the number of transmitter-receiver pairs).

The optimisation relies on the gradient of the cost function with respect to the model parameters (here: slowness distribution). To calculate this gradient, one needs to find the derivative of the time taken by a ray with respect to the slowness of a cell covered by this ray. From Eq. (2), we can write

$$\frac{\partial \tau}{\partial s_j} = \sum_{i=0}^N \left[\frac{\partial s_i}{\partial s_j} l_i + s_i \frac{\partial l_i}{\partial s_j} \right] \quad (5)$$

A key assumption here is that the ray path remains constant, irrespective of the local properties, i.e. $\partial l_j / \partial s_j = 0$. This is an approximation, but it simplifies the inversion considerably. The path variations will also be implicitly accounted for in subsequent iterations of the gradient stepping algorithm. The path invariance assumption removes the second component of the above sum, while the first simplifies when one notes that $\partial s_i / \partial s_j$ is only non-zero if $s_i = s_j$. Hence, the final expression for the derivative becomes

$$\frac{\partial \tau}{\partial s_j} = l_j \quad (6)$$

143 meaning that the derivative of the propagation time over the local slowness is simply the local ray length
 144 (the length of a section passing the considered cell).

One can now formulate the gradient of the cost function with respect to the local properties, using the chain rule,

$$\frac{\partial C}{\partial s_j} = \sum_{i=0}^M 2(\tau_{i,\text{model}} - \tau_{i,\text{exp}}) \left(\frac{\partial \tau_{i,\text{model}}}{\partial s_j} - \frac{\partial \tau_{i,\text{exp}}}{\partial s_j} \right) \quad (7)$$

We denote the difference between the arrival times from the model and from the measurement as the residual r_i and we note that the measured arrival times do not depend on the slowness distribution in the model, which leads to

$$\frac{\partial C}{\partial s_j} = \sum_{i=0}^M 2r_i \frac{\partial \tau_{i,\text{model}}}{\partial s_j} \quad (8)$$

145 where the last term can be replaced by l_i for a heterogeneous isotropic medium.

Hormati et al. [18] suggest using the nonlinear conjugate gradient algorithm to step along the gradient toward the desired model properties. This approach calculates the step size at each iteration and is capable of following narrow paths (ill-conditioned), taking the criss-cross pattern. The algorithm starts from an initial guess and calculates the gradient based on the computed ray lengths as outlined above. Each subsequent step follows the local descent direction

$$\Delta \mathbf{s}_n = -\nabla C(\mathbf{s}_n) + \gamma_n \Delta \mathbf{s}_{n-1} \quad (9)$$

146 where the δ_n parameters determines the effect of the preceding step and can be calculated using the Fletcher-
 147 Reeves or alternative formulae. It may happen that the stepping will not decrease the cost function. In
 148 such cases, a back-tracking line search algorithm ensures that the cost function is decreased enough. The
 149 algorithm iteratively decreases the step size by a chosen factor, until the cost function drop is less than a
 150 chosen threshold.

151 From the computational viewpoint, each step of the gradient stepping approach requires solving the
 152 forward model once to obtain the gradient. If back-tracking line search is necessary, each iteration takes
 153 another forward solution. These offer a very efficient setup, with the number of forward solutions usually
 154 not much greater than the number of steps. This is thanks to the convenient expression for the gradient
 155 developed above.

156 3.2. Heterogeneous anisotropic medium

157 The ray inversion method described in the previous section assumes that each cell represents an isotropic
 158 medium. On the other hand, anisotropy is one of the key characteristics of inhomogeneous welds, and
 159 therefore, the algorithm needs to be extended to be suitable. For heterogeneous anisotropy the local group
 160 slowness is a function of the incident angle and the local grain orientation. Hence, the model parameters to




Figure 3: Group slowness for a transversely isotropic material - how the incident angle and grain orientation form a combined ray angle in material coordinates.

161 be determined using inversion are now local grain orientation, assuming that all cells can be described with
 162 one elastic tensor, appropriately rotated cell-by-cell.

163 To benefit from the characteristics of the grid - identical regular cells - one needs to express the local
 164 slowness for a given ray in material coordinates. Since all cells share the same ‘basic’ tensor, the group
 165 slowness curve is simply only rotated from cell to cell. Hence, the combined angle is a difference between
 166 the incident angle and the local orientation. Fig. 3 schematically depicts the angles involved and illustrates
 167 how the combined angle is formed.

The algorithm requires an expression for the gradient of the cost function with respect to local orientations. Let us first consider the derivative of the time-of-flight along a given ray. The slowness is now a function of both the incident angle and local grain orientation; however, under the path invariance assumption stated in the preceding section, only the dependence on grain orientation is retained. Hence, we write

$$\frac{\partial \tau}{\partial \theta_j} = \frac{\partial \tau}{\partial s_j} \frac{\partial s_j}{\partial \phi_j} \frac{\partial \phi_j}{\partial \theta_j} \quad (10)$$

where $\partial \tau / \partial s_j = l_j$, as previously shown, $\partial \phi_j / \partial \theta_j = -1$, since $\phi_j = \gamma_j - \theta_j$, and $\partial s_j / \partial \phi_j$ can be precomputed at the beginning, and interpolated as a callable function, as it stays the same for all cells. The derivative of the ray time-of-flight over local orientation finally is

$$\frac{\partial \tau}{\partial \theta_j} = -l_j \frac{\partial s_j}{\partial \phi_j} \Big|_{\text{at } \phi_j = \gamma_j - \theta_j} \quad (11)$$

This expression determines the form of the cost function gradient expression which is the core of the inversion model:

$$\frac{\partial C}{\partial \theta_j} = \sum_{i=0}^M -2r_i l_j \frac{\partial s_j}{\partial \phi_j} \Big|_{\text{at } \phi_j = \gamma_j - \theta_j} \quad (12)$$

168 It is worth noting that the new term needs to be evaluated for each ray separately at every pixel, which
 169 adds complexity and computational cost to the isotropic solution, but these are in part minimised by the
 170 computational arrangement that takes the advantage of the regular grid structure.




Figure 4: Auxiliary mirrored domain for inversion through the backwall: a) the setup; b) illustration of the reduction of the view angle between through-transmission and pulse-echo.

171 3.3. Implementation for an inhomogeneous weld

172 The weld inversion algorithm described above and the the forward SRP model were developed and coded
 173 in Python. The algorithm is generally applicable to generally anisotropic media, however its application to
 174 inhomogeneous welds gives a specific flavour to several aspects. They are briefly outlined below.

175 3.3.1. Material structure and elastic tensor for the weld

176 Most parent materials joined using the welds of interest to this paper can be modelled as isotropic.
 177 Hence, a typical inversion setup considered in the ADVISE project is solving only for grain orientations
 178 within the weld, with the parent material isotropic properties assumed a priori. The ‘basic’ transversely
 179 isotropic elastic tensor governing the columnar grain structure of the weld is considered constant, i.e. only
 180 the orientations are allowed to vary. This assumption allows for simplifying the computational domain and
 181 discretising only the weld region, considerably reducing the size of the graph along which the shortest path
 182 is of interest.

183 3.3.2. Transducer configuration

184 Transducer configuration is a challenge for weld map tomography. The classical tomography setup (both
 185 optical and acoustic) assumes a circular array of transmitter-receivers, ensuring a full coverage of view angles.
 186 Any setup different from the full view incurs considerable loss in the ability to recover the properties of the
 187 medium. Since weld map tomography relies on measured arrival times, one can realistically consider three
 188 configurations – (i) a through-transmission setup, with transducers on both the top and bottom surfaces; (ii)
 189 a pulse echo setup, with only one array on top of the weld; or (iii) a tandem configuration with two arrays
 190 located on the top surface on both sides of the weld. The illustration of the implication of the reduced view
 191 for the through-transmission and the pulse-echo setup can be found e.g. in [22].

192 The through-transmission setup is straightforward for the considered model, as the sources and receivers
 193 are places along the two edges of the computational domain. However, the pulse-echo and the tandem setup
 194 cannot be modelled directly. To account for the reflection from the backwall with the SRP model, one needs

195 to considered an extended domain, mirrored around the backwall. Note that not only the nodes but also
196 material properties, including orientation angles need to be flipped. This technique is a direct consequence
197 of the Fermat principle invoked in the preceding section. Fig. 4 schematically depicts the arrangement when
198 through-backwall propagation is considered.

199 *3.3.3. Incorporating the known information and regularisation*

200 Weld map tomography is an ill-posed problem with multiple minima. Compared to the isotropic setup,
201 the added complexity of the anisotropic weld context comes from the fact that the group slowness is not
202 unique with respect to the ray angle. Hence, without incorporating additional information the inversion is
203 likely to fail. Such input is possible both in the form of an educated initial guess/starting point for inversion
204 and regularisation.

205 The grain layout of inhomogeneous welds is linked to the manufacturing procedure and the physics of
206 solidification. These factors are at the origin of the previously mentioned weld models, such as MINA or the
207 Ogilvy map. Whilst fast and effective, these models are constrained by the assumptions and simplifications.
208 We suggest to take advantage of these models by using them as initial guesses for inversion. With the
209 expectation that the inversion will target variations of the angles from this starting point, there is a higher
210 chance of convergence compared to the case where the starting angles are random.

211 Likewise, metallography/EBSD examinations of welds show that the variation of angles is typically
212 smooth with just a few abrupt changes. The orientations, following welding passes, tend to form regions
213 with similar dominant orientations which vary smoothly across the cross section. Hence in many cases
214 adding total variation and smoothness regularisation terms helps achieve convergence and reduce the risk
215 of being trapped in non-physical local minima.

216 The analysis of destructive examinations suggest that often there exists a dominant orientation of grains.
217 For the simplest symmetric cases, this will be the vertical direction, i.e. perpendicular to the surface. For
218 other welds it may be a line almost parallel to the chamfer. This dominant orientation affects the propagation
219 path considerably, and one approach to take advantage of this knowledge is to divide the inversion into two
220 stages. Starting from the conventional Ogilvy map, in the first stage the algorithm is only adjusting the
221 dominant orientation, that is, a single angle. When converged, it proceeds to updating local orientations
222 across the weld. This approach was found particularly effective in our work and represents a pragmatic
223 balance between the flexibility of ultrasonic tomography and the knowledge about the structure and the
224 behaviour of welds of interest. This approach required a modification of the Ogilvy model, referred to as the
225 ‘generalised Ogilvy map’, similar to that used in [17], which is described in more detail in Section Appendix
226 A.

227 4. Numerical example

228 4.1. Microstructure and forward models

229 Having outlined the conceptual principles of weld map tomography, the attention now turns to examples.
230 First, we present a numerical example in which the target times-of-flight came from grain-scale time-domain
231 finite element simulations. The calculations were executed on a graphical processing unit (GPU) using Pogo,
232 a fast GPU finite element solver developed at Imperial College London [23]. The example is motivated by a
233 real industrial inspection case at EDF, France. The grain structure of the weld bent and split the incident
234 ultrasonic beam, so that the reflection from the backwall at some locations could be misinterpreted as a
235 defect. At the time, to support the interpretation of the inspection, a macroscopic weld description was
236 created, based on metallography. We used that macroscopic descriptions to create a numeric microstructure,
237 allowing us to capture the rich physics of ultrasound propagation through complex welds in simulation.

238 The model was generated using free polycrystal generation package – **neper** [24]. First, we created a
239 vertically oriented columnar grain microstructure by stretching a standard Poisson-Voronoi tessellation along
240 one dimension. Then for each orientation zone, the starting model was rotated and cropped according to
241 the weld description from EDF. Each zone was accompanied with an individually generated set of crystallo-
242 graphic orientations which were randomly distributed along the dominant orientation (referred to as ‘fiber’
243 orientations in **neper**). The final step was comprised of stitching the microstructures for each zone into a
244 complete weld representation. The starting macroscopic weld descriptions and the grain-scale model are
245 shown in Fig. 5.

246 The generated microstructure was the basis for setting up the grain-scale finite element model, which
247 is the virtual experiment in this section, depicted schematically in Fig. 6. The weld region was surrounded
248 by isotropic regions, representing the parent plates ($c_{11} = 260$ GPa, $c_{12} = 100$ GPa, $c_{44} = 80$ GPa, $\rho =$
249 8000 kg m⁻³). For each grain belonging to the weld, we rotated the crystal elasticity matrix (cubic crystal,
250 $c_{11} = 206$ GPa, $c_{12} = 133$ GPa, $c_{44} = 119$ GPa, $\rho = 8000$ kg m⁻³) according to Euler angles generated in
251 **neper**. We simulated the action of a 64-element 2 MHz array on top of the weld by a series of vertical
252 nodal forces (one force per element) separated by the pitch distance and acting on the top boundary of
253 the weld with a 3-cycle toneburst time history. To minimise the impact of the finite computational domain
254 and eliminate reflections, we used absorbing regions around the whole domain (except the backwall). The
255 discretisation used a regular quad mesh (CPE4-type linear elements) with element size of 0.025 mm, which
256 was a very conservative choice providing more than 100 elements per longitudinal wavelength. The time
257 step was chosen based on the Courant number of 0.9, giving 3.7 ns. The model had nearly 12 million degrees
258 of freedom, but was solved efficiently using Pogo [23] on a GPU card (illustrative solution times for one
259 array shot with over 6800 time steps: standard gaming card Nvidia GTX 1080Ti – 15 s, Nvidia RTX 2080Ti
260 – 10 s).

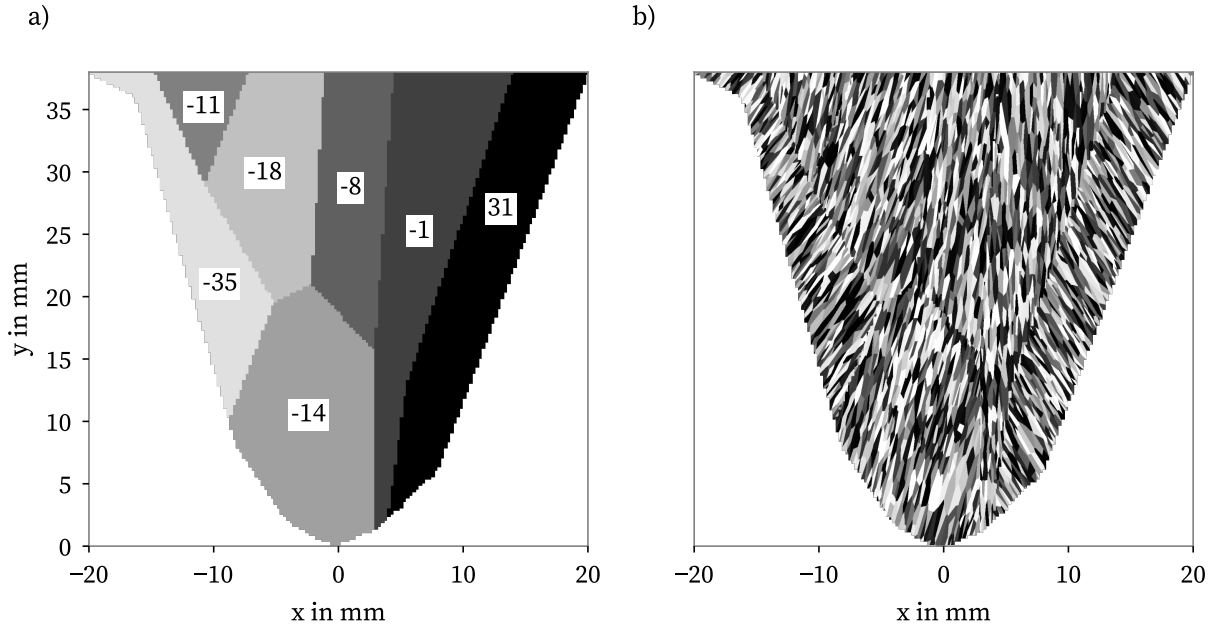


Figure 5: The macroscopic weld description with dominant orientations for each zone marked (a), and the generated synthetic microstructure (b) for the numerical example of weld map tomography. The orientations are quoted in degrees, measured from the vertical direction, with positive angles going anti-clockwise




Figure 6: Schematic diagram for the FE model for the numerical example of weld map tomography.

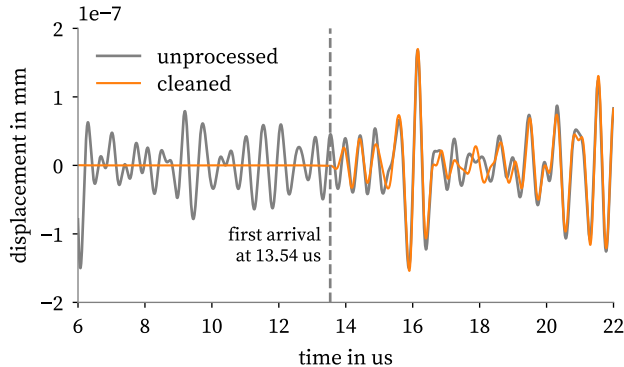


Figure 7: Illustrative simulated time trace for the numerical example - array element #41 transmitting and element #36 receiving.

261 Weld map tomography, as used in this paper, is based on backwall-reflected first arrivals, that is, longi-
 262 tudinal waves. To simplify extracting them from simulation, we executed an auxiliary calculation in which
 263 the backwall was replaced by an absorbing boundary. Substituting the two datasets led to clean first ar-
 264 rivals at the top surface. Such an approach is of course impossible in practice, but helps focus solely on the
 265 inversion algorithm, rather than peripheral issues, in a numerical example. Fig. 7 shows an example time
 266 trace and both unprocessed and cleaned signals, and the estimated arrival time. The estimation used an
 267 implementation of the Akaike Information Criterion [25] which helps if the first arrival is not the strongest in
 268 amplitude. For experimental datasets, obtaining an automated universal procedure is a research challenge
 269 on its own, and to keep the focus on inversion, we took the advantage of manual input.

270 To run the inversion on the simulated ultrasonic data, we first set up the SRP ray tracing model. The
 271 domain was cast on a regular 1 mm grid to accommodate weld geometry. The edges of each cell were
 272 seeded with ten nodes, and, to take the computational advantage described in Section 2, the algorithm
 273 pre-calculated the potential paths across a cell which were then used in all subsequent tracing. The domain
 274 was mirrored around the backwall, so that the pulse-echo configuration can be simulated. The SRP model is
 275 based on homogenisation, that is, it does not consider individual grains, but assigns local material properties
 276 based on the average grain orientation. Regions with columnar grains aligned to the same dominant angle are
 277 known to behave like a transversely isotropic material [6]. We exercised a homogenisation calculation based
 278 on the numerically generated microstructure, knowing that the dominant orientation should align with the
 279 direction of the smallest stiffness. The homogenisation yielded the following transversely isotropic material
 280 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,
 281 and the same density as for the parent material $\rho = 8000$ kg m⁻³. The vertically oriented grains have the
 282 33 material axis aligned with the global z axis. Orientation angles are measured from the vertical, positive
 283 anticlockwise. A short note on the orientation convention is included in Section Appendix B. Note also

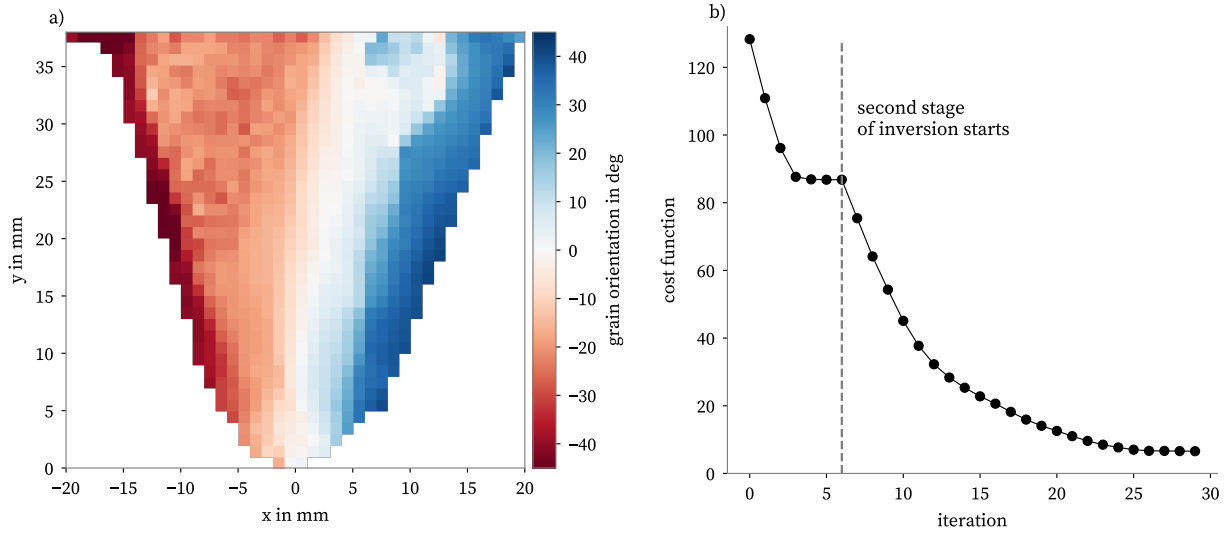


Figure 8: The result of weld map tomography based on time traces from grain-scale finite element simulations: a) obtained grain orientations; b) the cost function.

284 that the SRP model is set up in the y - z domain with z being the vertical, while the FE model is in the x - y
 285 domain.

286 4.2. Inversion

287 We used the two-stage inversion described above. The algorithm started with a ‘standard’ Ogilvy map
 288 ($T_L = T_R = n_L = n_R = 1$, $\alpha_L = \alpha_R = 20$) with a vertical dominant orientation ($\zeta = 0^\circ$), converging to
 289 -4.7° after six iterations. In the per-cell inversion stage we used smoothing and total variation regularisation,
 290 achieving convergence after a further 23 iterations.

291 Fig. 8 shows the obtained weld map and the evolution of the cost function. The vertical line marks
 292 the point where the first stage, in which only the dominant angle was varied for a general Ogilvy map, was
 293 completed. Individual orientations were then updated to the point where for the three consecutive iterations
 294 the cost function dropped by less than 2% (this criterion was adopted for both stages). The weld map can
 295 be compared to the original microstructure in Fig. 5.

296 Detailed error maps for the absolute angular error and the time-of-flight error are presented in Fig. 9.
 297 The mean, median, and maximum absolute angular errors ($|\theta_{\text{original}} - \theta_{\text{inversion}}|$) are 7.28, 5.71, and 33.06°,
 298 respectively. From a practical viewpoint, these are good results, compared to commonly reported average
 299 errors exceeding 20° [10]. We also compared the converged time-of-flight map to that chosen as a target
 300 for inversion in Fig. 9(b). The mean, median, and maximum absolute time errors ($|\tau_{\text{original}} - \tau_{\text{inversion}}|$) are
 301 0.03, 0.02, and 0.27 μs , respectively.

302 The numerical example demonstrates the potential of weld map tomography and quantifies its perfor-
 303 mance. While errors exist, they are below the commonly reported level (as expected from a simulation-based




Figure 9: Error maps of weld map tomography based on time traces from grain-scale finite element simulations: a) orientation errors; b) time of flight errors.

304 study) and related to the limitations of time-of-flight tomography itself.

305 5. Experimental examples

306 In this section we report exercising weld map tomography on three experimental samples - two industrial
 307 mock-ups from EDF and one mock-up specifically manufactured for the ADVISE project. All samples had
 308 artificial defects created for testing imaging algorithms. Verifying the effectiveness of the inversion directly
 309 is not straightforward, as usually destructive material examinations are taken over a different part of the
 310 weld than that used in the measurement. But it is widely accepted in the community that grain orientations
 311 do not vary significantly along the weld, except for arc ends. Nevertheless, in a practical scenario, the best
 312 evaluation of the obtained weld structure information is its application to imaging. Hence, for all numerical
 313 examples, we demonstrate how the result of weld map tomography improves ultrasonic array images and
 314 the signatures of target defects.

315 5.1. Weld with uniform orientations

316 First, we consider a homogeneous Inconel 182 wide chamfer weld from EDF, with approximately uniform
 317 orientations, as used in [16]. Fig. 10 shows a photograph and a schematic diagram of the experimental pulse-
 318 echo setup used for inversion. As opposed to [16], we used contact measurements, with the array (5 MHz
 319 64-element; pitch 0.6 mm, element: 0.5 mm) placed on the surface of the weld, centrally with respect to
 320 the middle defect. This is the simplest test case, where the target of the inversion is a single orientation
 321 angle. This means that the objective function depends only on one variable and that the gradient of the
 322 objective function is a sum of contributions from all cells. The top and bottom surfaces of the model were




Figure 10: Schematic diagram for experimental weld map tomography on an EDF Inconel 182 V-weld with uniform orientations: a) photograph; b) experimental setup.

not exactly parallel, which was accounted for in the forward model underlying inversion. Since the pulse echo setup requires mirroring the computational domain, an arbitrary backwall profile can be defined, provided that the respective nodes belonging to the original and auxiliary domains are connected with zero-cost (zero-arrival-time) edges. This arrangement will be further demonstrated in Section 5.2.

The forward model underpinning inversion discretised the domain into 2×2 mm cells, with each edge of a cell seeded with 10 nodes. For the parent material, we used $\rho_{\text{parent}} = 8260 \text{ kg m}^{-3}$, $c_{11} = 260 \text{ GPa}$, $c_{12} = 100 \text{ GPa}$, and $c_{44} = 80 \text{ GPa}$. The weld material had the same density, and the elasticity matrix with $c_{11} = 265 \text{ GPa}$, $c_{12} = 121 \text{ GPa}$, $c_{13} = 140 \text{ GPa}$, $c_{33} = 247 \text{ GPa}$, $c_{44} = 112 \text{ GPa}$, and $c_{66} = 86 \text{ GPa}$. The latter was determined from the data provided by EDF.

As signalled in the preceding section, extracting time-of-arrival from real experimental data, 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. Fig. 11 shows an example of received signals across all array elements for a single excitation. The solid line represents extracted arrival time, whose small variations affect the inversion significantly. While the discussion on algorithms for determining arrival times is outside this paper's scope, we report using a method based on tracking the zero-crossing between two dominant extrema of the wave packet in this paper. Once the zero-point has been identified for all transducer pairs, we employed interpolation and the symmetry of the arrival time matrix to infer the values at missing points. The final result came from subtracting a constant offset, estimated on the cleanest time trace in the set, capturing the difference between the arrival time extracted using the AIC criterion and the previously chosen zero-crossing. While this procedure may appear complicated, it allowed for obtaining consistently satisfactory results for the examples considered in this paper.

Fig. 12 shows the results of the inversion, based on experimental time traces recorded at CEA. Weld map tomography started from uniform vertical orientation of grains and converged to -8.11 deg after 9 iterations. Determining the exact angular error is not possible for this example, but the metallography performed on an off-cut of the same weld suggests the dominant orientation of -8 deg which confirms that the inversion

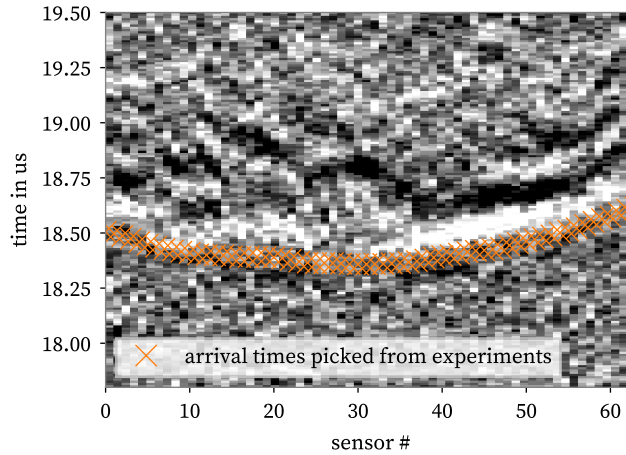


Figure 11: Extracting time-of-arrival as the target for inversion from experimental data (the map combines time traces across all sensors from single-transducer excitation; the markers show the identified times-of-arrival).

348 reached a realistic value. However, quantitative error analysis is possible for arrival times, as depicted in
 349 Fig. 12(b). The mean absolute error was $0.021 \mu\text{s}$, the median absolute error $-0.015 \mu\text{s}$, and the maximum
 350 absolute error $-0.093 \mu\text{s}$. The highest errors are between transducers close to array extremities, suggesting
 351 that the geometry of the sample was not captured correctly. Despite that, the result of the tomography is
 352 close to the reference provided for the mock-up. We attribute this robustness to the fact that only a single
 353 orientation angle is sought for in this example, minimising the impact of locally mismatched definitions.

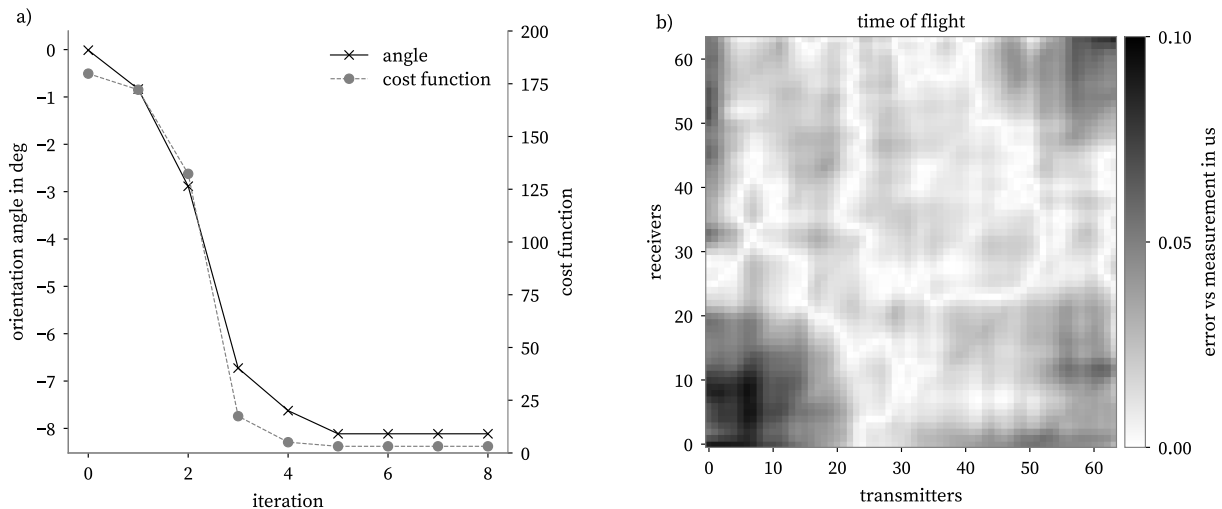


Figure 12: Results of weld map tomography for the uniform Inconel weld: a) cost function and orientation angle at different iterations; b) a map of absolute time-of-flight errors ($|\tau_{\text{original}} - \tau_{\text{inversion}}|$).

The practically relevant verification of the inversion is to perform imaging with the obtained material information. We used the classical total focusing method (TFM) [26] which is a modified, Hilber-envelope-

based version of the beamforming algorithm. The imaging function $I(x, y)$ is defined as

$$I(x, y) = \left| \sum_{i=1}^M \sum_{j=1}^M \text{H} [u_{i,j} (\tau_{i \rightarrow (x,y)} + \tau_{j \rightarrow (x,y)})] \right| \quad (13)$$

where (x, y) is the position in the physical domain, M is the number of elements in the array, $u_{i,j}$ 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 $\text{H}[\cdot]$ is the Hilbert transform.

Fig. 13 shows TFM images calculated using the same dataset as the one 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 normalised to the maximum amplitude of the middle defect (individually). The circular markers indicate true positions of side drilled holes. To quantify the difference between the two images we calculated the signal-to-noise ratio (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)

$$\text{SNR} = 20 \log_{10} \frac{\max(I_{\text{sig}})}{\sqrt{I_{\text{noise}}^2}} \quad (14)$$

The results in Fig. 13 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 the benefits it provides to imaging.

5.2. 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 (approximately 10 and 20 mm from the top surface) provide a reference for imaging. Compared to 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, taken at CEA, 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 Section 5.1. Weld map tomography, as

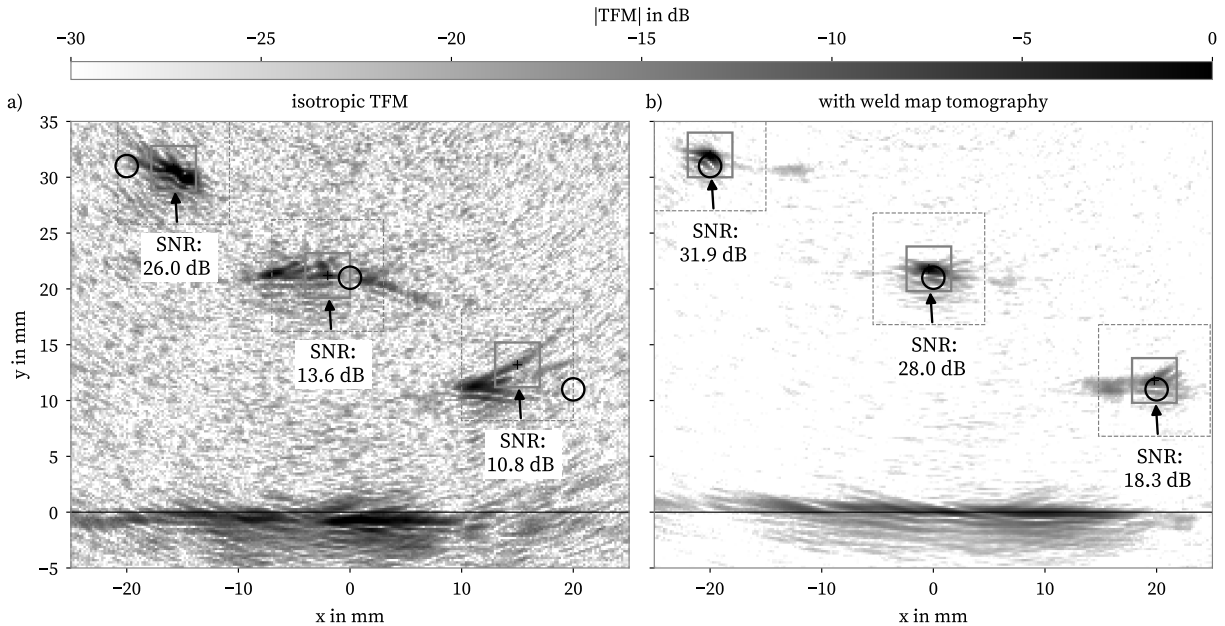


Figure 13: Total focusing method imaging with material information: a) standard isotropic delay laws; b) with delay laws coming from weld map tomography. The SNRs of the defect signatures indicated for comparison.

372 presented in this paper, focuses on determining local grain orientations, that is, material information only.
 373 Thus, the geometry was fixed and assumed known, which is a straightforward task for a mock-up. Since
 374 this paper focuses on updating material properties, keeping the geometry fixed, we omit the simplified
 375 representation of the backwall profile for brevity.

376 The forward model used in the inversion discretised the domain into 2×2 mm cells, seeded with ten
 377 nodes per each side. For the parent material, we used $\rho_{\text{parent}} = 7800 \text{ kg m}^{-3}$, $c_{11} = 276 \text{ GPa}$, $c_{12} = 102 \text{ GPa}$,
 378 and $c_{44} = 87 \text{ GPa}$. The weld material had the same density, and the elasticity matrix with $c_{11} = 245 \text{ GPa}$,
 379 $c_{12} = 110 \text{ GPa}$, $c_{13} = 145 \text{ GPa}$, $c_{22} = 253 \text{ GPa}$, $c_{33} = 221 \text{ GPa}$, $c_{44} = 110 \text{ GPa}$, and $c_{66} = 75 \text{ GPa}$ [17]. We
 380 exercised the two-stage inversion, identical to that in Section 4, where the first stage looks for the optimal
 381 dominant orientation angle for the generalised Ogilvy model, and the second stage updates local orientations
 382 to further decrease the cost function until convergence.

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

387 The arrival time errors are higher than in the previous example, with the mean absolute error of $0.034 \mu\text{s}$,
 388 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
 389 the ‘white spots’ in Fig. 14(b) indicate that for these transducer pairs it was not possible to extract arrival

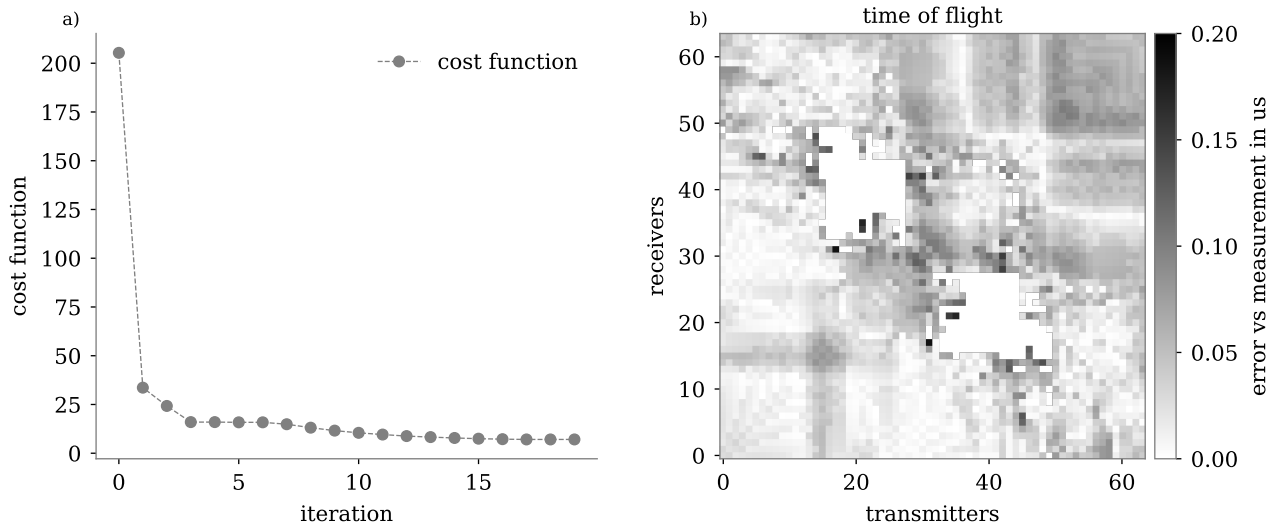


Figure 14: Results of weld map tomography for the stainless steel weld with a non-planar backwall: a) cost function and orientation angle at different iterations; b) a map of absolute time-of-flight errors ($|\tau_{\text{original}} - \tau_{\text{inversion}}|$).

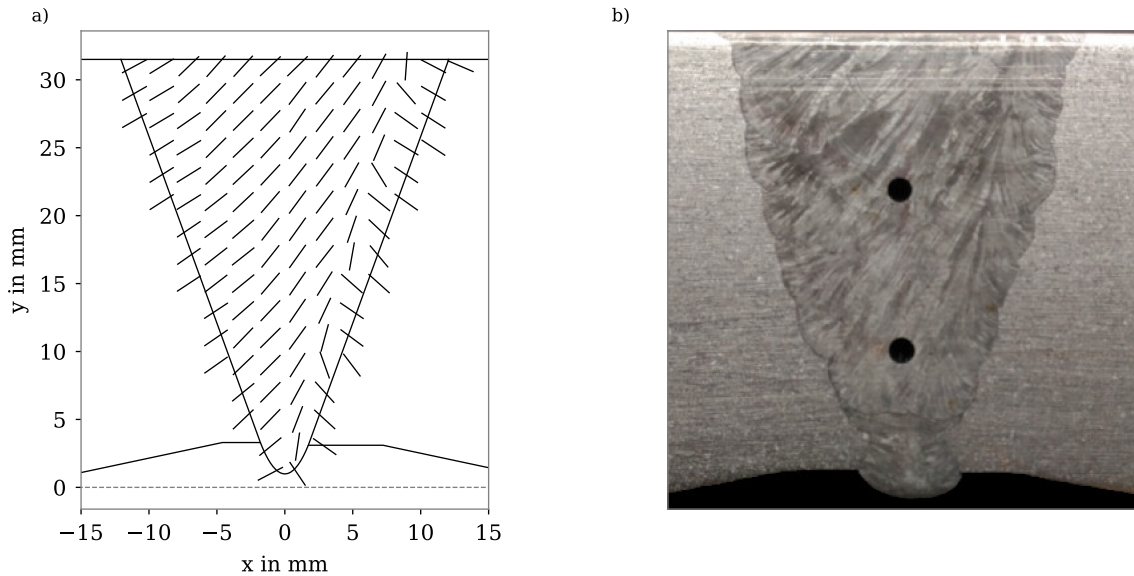


Figure 15: Results of weld map tomography for the stainless steel weld with a non-planar backwall: a) updated local grain orientations; b) the metallography image of the mock-up.

390 times from the measured data. While this reduction in the available data is undesired for the inversion, this
 391 case confirms that weld map tomography does not require the full dataset to work. We also observe that
 392 the maximum errors occurred for transducer pairs close to the undefined arrival time regions, bound by high
 393 uncertainty. The mean error is still low and represents a practically promising performance. In Fig. 15 we
 394 offer a visual comparison of the updated orientations with the metallography of the mock-up, confirming

395 the effectiveness of the inversion.

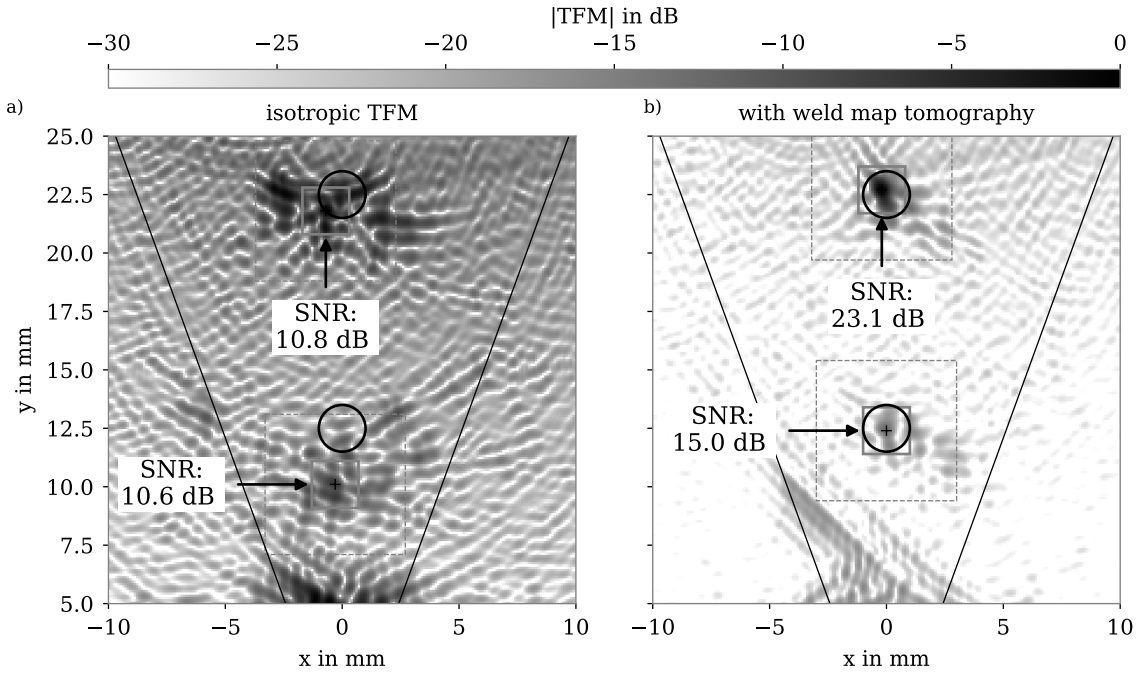


Figure 16: Total focusing method imaging with material information for the stainless steel weld with a non-planar backwall: a) standard isotropic delay laws; b) with delay laws coming from weld map tomography. The SNRs of the defect signatures indicated for comparison.

396 The final verification involves calculating the TFM image using the same recorded dataset with and
 397 without material information provided by weld map tomography, as shown in Fig. 16. Updated material
 398 information reduces the noise in the image and improves the prominence and focus of defect signatures.
 399 While the noise is not eliminated entirely, the SNRs of the target defects improve by 12.3 and 4.4 dB,
 400 respectively. We acknowledge artefacts close to the root of the weld; however, their origin, and ways for
 401 reducing them, remain the subject of future work. Nevertheless, our proposed inversion technique output
 402 local orientations visually agreeing well with the metallography image and providing good enhancement to
 403 the ultrasonic image of target defects, despite the complex backwall geometry.

404 5.3. ADVISE mock-up

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

409 The mock-up was manufactured by joining two stainless steel pipes (1.4550) with wall thickness of 39 mm
 410 (outside diameter: 517 mm) - Fig. 17. The circumferential weld, in the flat position, was manufactured




Figure 17: The ADVISE WP2 mock-up: a photograph and schematic drawings of a cut out section.

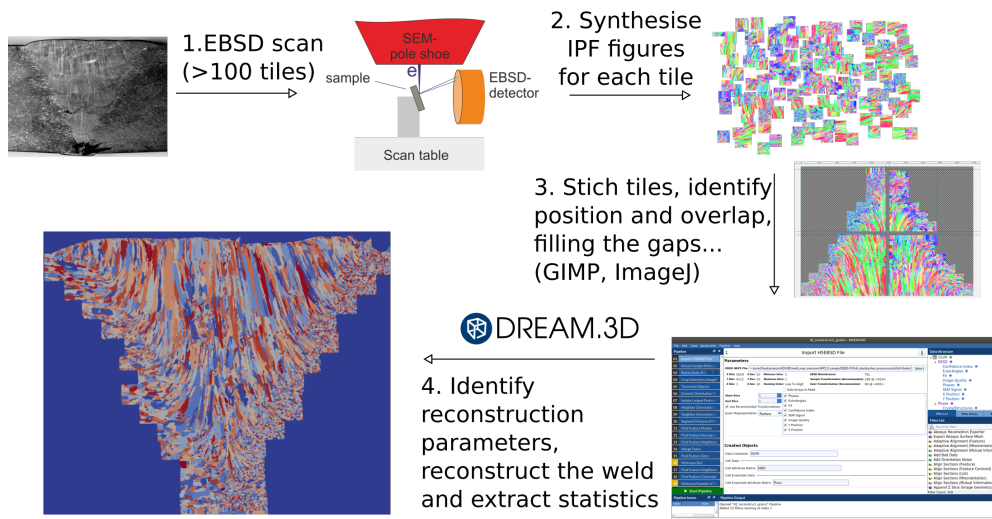


Figure 18: The workflow for the reconstruction of the microstructure from EBSD measurements.

411 using the shielded metal arc method with ESAB OK 61.30 filler (1.4316). The cylindrical specimen was
 412 then sectioned into several parts along the circumference (Fig. 17), which served other research tasks.

413 One section of the mock-up was used for material examinations, both traditional metallography and
 414 EBSD. For the EBSD measurement, the weld was further sectioned into four parts and each was covered by
 415 a grid of smaller measurement areas in the microscope. The obtained dataset, comprising over 100 ‘tiles’ was
 416 stitched and analysed using Dream.3D [27], an open-source package for microstructure synthesis, to deliver
 417 a numerical microstructure, that can be visualised, analysed, and subsequently used in further modelling or
 418 the verification of experimental inversion.

419 To facilitate measurements, the top and bottom surfaces of the mock-up were machined to achieve flat
 420 surface (Fig. 19). This modification helped ensure good contact between the transducer and the sample and
 421 allowed focusing on the inversion, rather than the transducer issues. Material evaluations at the University

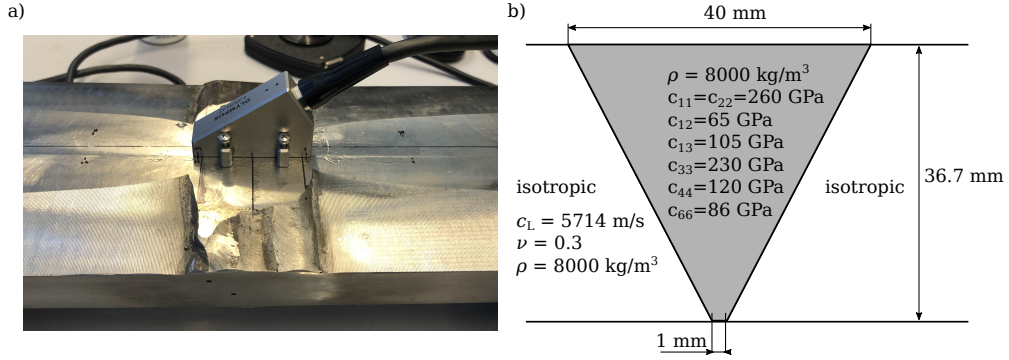


Figure 19: ADVISE WP2 mock-up and the transducer used in weld map tomography: a) a photograph; b) dimensions with material properties.

422 of Stuttgart indicated several natural defects within the weld and additional targets (holes) were drilled
 423 along the chamfer at CEA.

424 Array measurements using a 64-element, 0.75 mm 2.25 MHz transducer (in contact with the sample,
 425 directly on top of the weld, coupled through silicone gel) were taken at KTU using the Gekko system
 426 from Eddyfi. The acquisitions were processed using the semi-manual time-of-arrival extraction procedure
 427 described in Section 5.1, yielding the target for the weld map tomography algorithm.

428 The forward model used in inversion assumed known material properties of the sample and known
 429 geometry. We used the EBSD dataset to execute homogenisation and determine the macroscopic transversely
 430 isotropic elastic tensor based on the grain structure and the crystal elastic tensor. These parameters are
 431 listed in Fig. 19.




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

432 Two-stage weld map tomography, analogous to that in Section 4, converged after 26 iterations, with
 433 the first stage (updating the dominant orientation of the generalised Ogilvy map) taking four iterations.
 434 Fig. 20 shows the cost function and the final map of absolute time of flight errors (mean absolute error:
 435 0.044 μs , median absolute error: 0.036 μs , and maximum absolute error: 0.193 μs). To further illustrate the

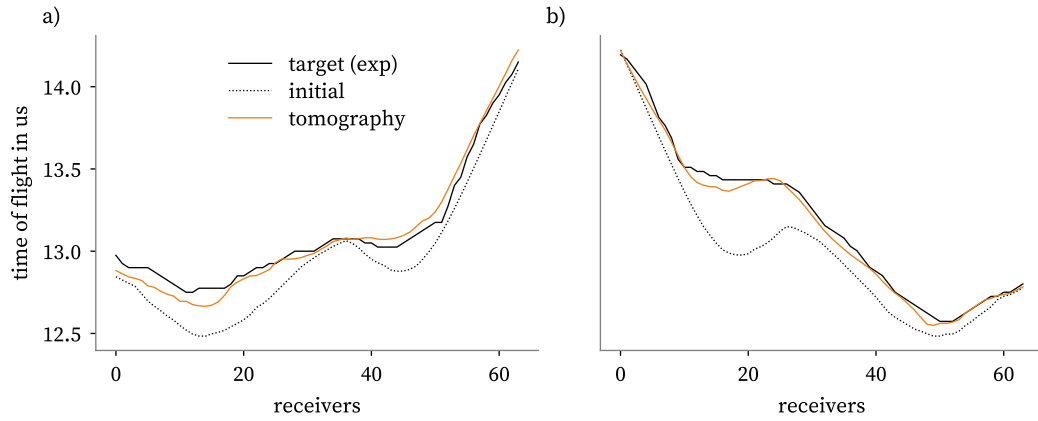


Figure 21: Evolution of the time-of-flight between the initial guess and the converged weld map tomography output vs the target: a) source at transducer #13; b) source at transducer #53.

436 evolution of the time-of-flight map during inversion, we show the initial, the final, and the target results
 437 for two sources in Fig. 21. In the illustrated cases, initial time-of-flight errors were up to $0.5\ \mu\text{s}$ for some
 438 receivers, while reducing to less than $0.1\ \mu\text{s}$ in the final result.

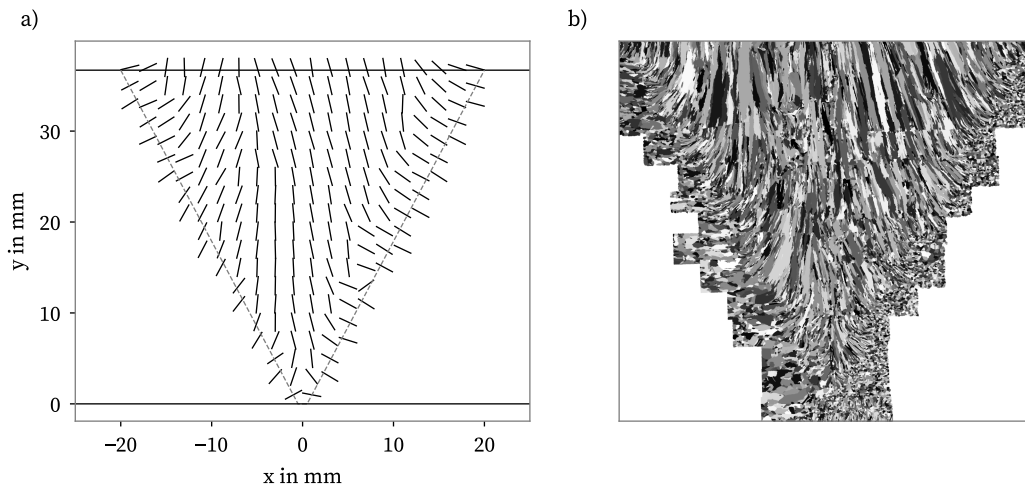


Figure 22: Results of weld map tomography for the stainless steel ADVISE mock-up: a) updated local grain orientations; b) microstructure reconstruction based on EBSD measurements.

439 Fig. 22 compares grain orientations determined using weld map tomography and the reconstruction of
 440 the EBSD measurement taken at the University of Stuttgart. This visual comparison gives a convincing
 441 indication of the effectiveness of the method. We note that the errors in orientation angles are mainly
 442 concentrated around the chamfer, where the EBSD data had low confidence index. Nevertheless, as in
 443 previously presented examples, the final confirmation of the efficacy of weld map tomography is illustrated
 444 by applying its output to imaging.

445 The TFM imaging used the dataset as that produced the time-of-flight target for inversion. Fig. 23
 446 compares images calculated: a) using the conventional isotropic assumption and b) delay laws based on the
 447 output of weld map tomography. The effect of the proposed inversion technique is clear and convincing.
 448 First, the sample had two known target defects (side-drilled holes) along the right chamfer - Fig. 23 shows
 449 that only one is visible in the conventional TFM image. Including material information coming from inversion
 450 raised the bottom chamfer defect from the background noise. The SNRs of the chamfer defects improved
 451 by around 4 and 8 dB. We note that if a defect signature is not present in the conventional image, the SNR
 452 value has no meaning (as there is no ‘signal’); nevertheless, it is calculated in the image for completeness.
 453 Additionally to the intentionally introduced side-drilled holes, two natural defect signatures emerged in
 454 Fig. 23b) - one close to the centreline and another in the vicinity of the chamfer. Post-manufacture X-ray
 455 scans of the weld indicated the presence of voids in some parts of the weld and it is confirmed with ultrasonic
 456 imaging. More importantly, the absence of these natural defects’ signatures in the conventional image proves
 457 the effectiveness of weld map tomography in supporting imaging and is a compelling proof of the importance
 458 of accounting for material information.

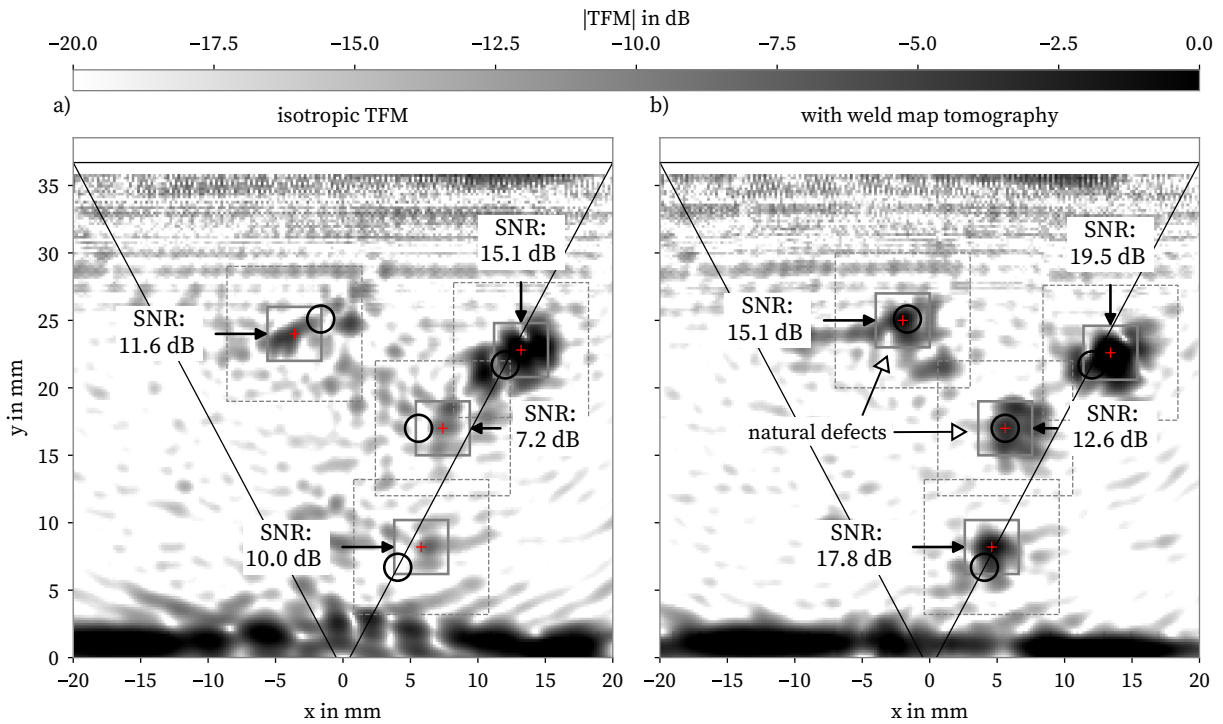


Figure 23: Total focusing method imaging with material information for the ADVISE mock-up: a) standard isotropic delay laws; b) with delay laws coming from weld map tomography. The SNRs of the defect signatures indicated for comparison.

459 6. Conclusions

460 This paper proposed a method for determining local grain orientations in complex, thick welds from
461 ultrasonic array data. The reconstruction followed the ray tomography principle, accounting for the hetero-
462 geneous anisotropy of the welds of interest. The inversion used a shortest-ray-path model as the forward
463 solver and the paper detailed the principles and construction of the model. We discussed the importance
464 of the a priori information about the weld and the advantages of using simple geometrical descriptions
465 for initial stages of inversion. The algorithm was first showcased using a numerical example based on an
466 industrial inspection scenario, with grain-scale time-domain finite element simulations acting as a virtual
467 experiment. The results of the inversion showed good agreement with the numerically generated microstruc-
468 ture. Encouraged by this result, we demonstrated weld map tomography on three mock-ups - two supplied
469 by the industry, and one manufactured specifically for this research. The outcome of the method compared
470 favourably with destructive material examinations (both metalography and EBSD) and provided the input
471 for updating the delay laws for imaging. In all examples, TFM images using material information from
472 inversion showed significant improvement of both the location of defect signatures and SNR, confirming the
473 efficacy of the proposed algorithm.

474 Acknowledgement

475 This research was a part of the ADVISE project (www.advise-h2020.eu) funded from the Euratom
476 research and training programme 2014-2018 under grant agreement no. 755500.

477 Appendix A. Generalised Ogilvy map

478 The geometric description of grain orientations proposed by Ogilvy [3] is a simple but powerful method.
479 Throughout recent decades it became the gold standard for modelling local orientations in austenitic stainless
480 steel welds. This paper uses a slightly extended version of the Ogilvy map, accounting for the dominant
481 orientation angle which is relatively common in published examinations (see e.g. Fig. 10a)). What is
482 presented below is conceptually analogous to an approach used in [17].

483 Let us start from the classical Ogilvy map [3] for completeness. We consider a simple V-weld, of thickness
484 a , left and right weld angles α_L and α_R , and weld bead lengths (positive) d_L and d_R . For a domain
485 defined in $y - z$ coordinates with z being the vertical, the Ogilvy model requires four parameters: T_L
486 and T_R , representing the tangent of the orientation angles at the left and right chamfers, and η_L and
487 η_R , representing the rate of change of the angle while moving horizontally from the chamfer towards the
488 centreline (where the grains are always vertical). The setup, the coordinate system, and the parameters are
489 illustrated in Fig. A.1. Our convention assumes the lengths, weld angles, and η parameters to be always




Figure A.1: Schematic diagram for: a) the classical Ogilvy model; b) auxiliary rotation of the coordinate system to account for dominant orientation ζ .

490 positive. Orientation tangent T_L is expected to be positive, while T_R - to be negative. It is common to
 491 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 generalised 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 (see Fig. A.1b)). 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:

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

492 When applying the dominant orientation parameter ζ to an originally symmetric weld (with $T_L = -T_R$),
 493 one should adjust these parameters to retain relevance to realistic cases. A rule-of-thumb adjustment based
 494 on comparing such maps to metallography images suggests multiplying initial T_L and T_R parameters by
 495 $(1 + 0.01\zeta)$ and $(1 - 0.01\zeta)$ respectively.

496 Appendix B. Orientation and material properties conventions

497 The inversion algorithm and the SRP forward model proposed in this paper rely on a macroscopic
498 representation of the weld, in which local properties are governed by a transversely isotropic material rotated
499 according to the orientation angle. The rotation is such that the ‘soft’ axis of the material aligns with the
500 orientation direction. The SRP forward model and weld map tomography are set up in the y - z domain,
501 hence the ‘soft’ axis will be the 33 material direction. If an analogous calculation was performed in the x - y
502 domain, e.g. in a 2D finite element simulation, one would need to appropriately rotate the material stiffness
503 tensor, so that the ‘soft’ axis will appear as the 22 material direction.

504 References

- 505 [1] G. D. Connolly, M. J. S. Lowe, J. A. G. Temple, S. I. Rokhlin, The application of Fermat’s principle for imaging
506 anisotropic and inhomogeneous media with application to austenitic steel weld inspection, *Proceedings of the Royal*
507 *Society A: Mathematical, Physical and Engineering Sciences* 465 (2009) 3401–3423. doi:10.1098/rspa.2009.0272.
- 508 [2] J. Moysan, A. Apfel, G. Corneloup, B. Chassignole, Modelling the grain orientation of austenitic stainless steel multipass
509 welds to improve ultrasonic assessment of structural integrity, *International Journal of Pressure Vessels and Piping* 80
510 (2003) 77–85. doi:10.1016/S0308-0161(03)00024-3.
- 511 [3] J. A. Ogilvy, Computerized ultrasonic ray tracing in austenitic steel, *NDT International* 18 (1985) 67–77. doi:10.1016/
512 0308-9126(85)90100-2.
- 513 [4] B. Chassignole, R. El Guerjouma, M.-A. Ploix, T. Fouquet, Ultrasonic and structural characterization of anisotropic
514 austenitic stainless steel welds: Towards a higher reliability in ultrasonic non-destructive testing, *NDT & E International*
515 43 (2010) 273–282. doi:10.1016/j.ndteint.2009.12.005.
- 516 [5] A. J. Hunter, B. W. Drinkwater, J. Zhang, P. D. Wilcox, D. O. Thompson, D. E. Chimenti, A study into the effects of
517 an austenitic weld on ultrasonic array imaging performance, 2011, pp. 1063–1070. doi:10.1063/1.3592054.
- 518 [6] A. F. Mark, Z. Fan, F. Azough, M. J. S. Lowe, P. J. Withers, Investigation of the elastic/crystallographic anisotropy of
519 welds for improved ultrasonic inspections, *Materials Characterization* 98 (2014) 47–53. doi:10.1016/j.matchar.2014.09.
520 012.
- 521 [7] M. Spies, Modeling of transducer fields in inhomogeneous anisotropic materials using Gaussian beam superposition, *NDT*
522 *& E International* 33 (2000) 155–162. doi:10.1016/S0963-8695(99)00036-5.
- 523 [8] K. J. Langenberg, R. Hannemann, T. Kaczorowski, R. Marklein, B. Koehler, C. Schurig, F. Walte, Application of
524 modeling techniques for ultrasonic austenitic weld inspection, *NDT & E International* 33 (2000) 465–480. doi:10.1016/
525 S0963-8695(00)00018-9.
- 526 [9] C. Gueudre, L. Le Marrec, J. Moysan, B. Chassignole, Direct model optimisation for data inversion. Application to
527 ultrasonic characterisation of heterogeneous welds, *NDT & E International* 42 (2009) 47–55. doi:10.1016/j.ndteint.
528 2008.07.003.
- 529 [10] Z. Fan, A. F. Mark, M. J. S. Lowe, P. J. Withers, Nonintrusive estimation of anisotropic stiffness maps of heterogeneous
530 steel welds for the improvement of ultrasonic array inspection, *IEEE Transactions on Ultrasonics, Ferroelectrics, and*
531 *Frequency Control* 62 (2015) 1530–1543. doi:10.1109/TUFFC.2015.007018.
- 532 [11] J. Zhang, A. Hunter, B. W. Drinkwater, P. D. Wilcox, Monte carlo inversion of ultrasonic array data to map anisotropic
533 weld properties, *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 59 (2012) 2487–2497. doi:10.
534 1109/TUFFC.2012.2481.
- 535 [12] K. M. M. Tant, E. Galetti, A. J. Mulholland, A. Curtis, A. Gachagan, A transdimensional Bayesian approach to ultrasonic
536 travel-time tomography for non-destructive testing, *Inverse Problems* 34 (2018) 095002. doi:10.1088/1361-6420/aaca8f.
- 537 [13] K. M. M. Tant, E. Galetti, A. J. Mulholland, A. Curtis, A. Gachagan, Effective grain orientation mapping of complex
538 and locally anisotropic media for improved imaging in ultrasonic non-destructive testing, *Inverse Problems in Science and*
539 *Engineering* 28 (2020) 1694–1718. doi:10.1080/17415977.2020.1762596.
- 540 [14] J. Singh, K. Tant, A. Mulholland, C. MacLeod, Deep Learning Based Inversion of Locally Anisotropic Weld Properties
541 from Ultrasonic Array Data, *Applied Sciences* 12 (2022) 532. doi:10.3390/app12020532.
- 542 [15] J. Singh, K. Tant, A. Curtis, A. Mulholland, Real-time super-resolution mapping of locally anisotropic grain orientations
543 for ultrasonic non-destructive evaluation of crystalline material, *Neural Computing and Applications* 34 (2022) 4993–5010.
544 doi:10.1007/s00521-021-06670-8.
- 545 [16] C. Ménard, S. Robert, R. Miorelli, D. Lesselier, Optimization algorithms for ultrasonic array imaging in homogeneous
546 anisotropic steel components with unknown properties, *NDT & E International* 116 (2020) 102327. doi:10.1016/j.
547 ndteint.2020.102327.
- 548 [17] C. Menard, S. Robert, D. Lesselier, Ultrasonic Array Imaging of Nuclear Austenitic V-Shape Welds with Inhomogeneous
549 and Unknown Anisotropic Properties, *Applied Sciences* 11 (2021) 6505. doi:10.3390/app11146505.

- 550 [18] A. Hormati, I. Jovanović, O. Roy, M. Vetterli, Robust ultrasound travel-time tomography using the bent ray model,
551 in: Medical Imaging 2010: Ultrasonic Imaging, Tomography, and Therapy, volume 7629, SPIE, 2010, pp. 165–176.
552 doi:10.1117/12.844693.
- 553 [19] O. Nowers, D. J. Duxbury, J. Zhang, B. W. Drinkwater, Novel ray-tracing algorithms in NDE: Application of Dijkstra
554 and A* algorithms to the inspection of an anisotropic weld, NDT & E International 61 (2014) 58–66. doi:10.1016/j.
555 ndteint.2013.08.002.
- 556 [20] O. Nowers, D. J. Duxbury, B. W. Drinkwater, Ultrasonic array imaging through an anisotropic austenitic steel weld using
557 an efficient ray-tracing algorithm, NDT & E International 79 (2016) 98–108. doi:10.1016/j.ndteint.2015.12.009.
- 558 [21] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson,
559 W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones,
560 R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman,
561 I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, SciPy 1.0 Con-
562 tributors, SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, Nature Methods 17 (2020) 261–272.
563 doi:10.1038/s41592-019-0686-2.
- 564 [22] P. Huthwaite, Quantitative Imaging with Mechanical Waves, Ph.D. thesis, Imperial College London, 2012.
- 565 [23] P. Huthwaite, Accelerated finite element elastodynamic simulations using the GPU, Journal of Computational Physics
566 257 (2014) 687–707. doi:10.1016/j.jcp.2013.10.017.
- 567 [24] R. Quey, P. R. Dawson, F. Barbe, Large-scale 3D random polycrystals for the finite element method: Generation, meshing
568 and remeshing, Computer Methods in Applied Mechanics and Engineering 200 (2011) 1729–1745. doi:10.1016/j.cma.
569 2011.01.002.
- 570 [25] A. St-Onge, Akaike information criterion applied to detecting first arrival times on microseismic data, in: SEG Technical
571 Program Expanded Abstracts 2011, SEG Technical Program Expanded Abstracts, Society of Exploration Geophysicists,
572 2011, pp. 1658–1662. doi:10.1190/1.3627522.
- 573 [26] C. Holmes, B. W. Drinkwater, P. D. Wilcox, Post-processing of the full matrix of ultrasonic transmit–receive array data
574 for non-destructive evaluation, NDT & E International 38 (2005) 701–711. doi:10.1016/j.ndteint.2005.04.002.
- 575 [27] M. A. Groeber, M. A. Jackson, DREAM.3D: A Digital Representation Environment for the Analysis of Microstructure
576 in 3D, Integrating Materials and Manufacturing Innovation 3 (2014) 56–72. doi:10.1186/2193-9772-3-5.