

Research Paper

Model-based deep neural network inversion for ultrasonic reconstruction of thick-section welds

Lucas Q. Machado ^a,* , Thomas Blumensath ^a, Vyktintas Samaitis ^c, Michael J.S. Lowe ^b,
Michał K. Kalkowski ^a

^a Institute of Sound and Vibration Research, University of Southampton, SO17 1BJ, United Kingdom

^b Imperial College London, United Kingdom

^c Kaunas University of Technology, Lithuania

ARTICLE INFO

Keywords:

Ultrasound
Material characterisation
Inversion
Weld map
Deep learning

ABSTRACT

The interpretation of ultrasonic wave propagation in inhomogeneous materials remains a challenging task. This is a topic of particular interest in the nuclear industry, where the inspection of inhomogeneous thick-section welds is vital to meet safety standards. Such welds are characterised by columnar grains with varying preferred orientations resulting from a complex solidification process. Consequently, conventional post-processing techniques of ultrasonic data may lead to erroneous conclusions, as the ultrasonic beam is distorted and attenuated, no longer following a straight path. This issue can be resolved if material information is available. We propose a model-based deep neural network (DNN) inversion workflow to infer weld microstructure description from time-of-flight (ToF) maps. Work to date focused on grain orientations; in this contribution, we extend it to weld geometry, array position, and the elastic tensor. The workflow uses numerically generated welds and a fast ray-tracing solver based on the shortest-ray-path (SRP) principle to determine propagation times. The weld orientation configurations (weld maps) use the generalised Ogilvy description, sampled in a parameter space with up to ten variables, such as the dominant grain orientation, the rate of change of the grain orientations, the chamfer angle, the array position, and the coefficients of the transversely isotropic elastic tensor in 2D. Therefore, the characterisation considers both geometric and material properties, encompassing a broad range of welds. We develop, use, and evaluate several DNN-based metamodels to determine weld maps from ultrasonic array data and validate the approach using grain-scale finite-element simulations and experimental ToF measurements. Weld maps resulting from the inversion are used to determine the delay laws for correcting the ultrasonic images of defects, achieving a signal-to-noise ratio improvement of up to 6.4 dB.

1. Introduction

Safety-critical applications, such as nuclear power plants, require the use of materials that can withstand demanding conditions, including high temperatures and high pressures. These materials often exhibit a complex metallographic structure with texture and significant heterogeneity, which pose challenges for ultrasonic inspection and its interpretation. Thick-section welds are a prime example, where columnar grains with varying preferential orientations result in deviation, splitting, and, notably, attenuation of the ultrasonic beam [1,2]. Conventional algorithms for defect detection and imaging operate on the assumption of a homogeneous and isotropic medium, meaning that ultrasound is only understood to travel along straight paths. Employing this assumption to inspect heterogeneous media results in misplaced

defect signatures. More importantly, mismatched delay laws cause the energy associated with a scatterer to spread spatially, resulting in a significant drop in the signal-to-noise ratio [3,4]. While the straight-ray assumption is common in many acquisition systems, material information can be incorporated into ultrasonic imaging through delay laws computed using an appropriate model [5].

The delay laws are computed based on the crystallographic structure of the material and its associated direction-dependent and spatially varying wavespeed. Crystallographic information can be obtained through examination using, e.g., the Electron Backscatter Diffraction (EBSD); numerical welding simulations [6,7]; geometrical models, such as the Ogilvy description [8,9]; and phenomenological models, such as MINA [10]. However, these methods have their shortcomings. EBSD is

* Corresponding author.

E-mail address: L.Queiroz-Machado@soton.ac.uk (L.Q. Machado).

destructive and costly and welding simulation is time-consuming and involves considerable computing resources. On the other hand, weld descriptions, like MINA, provide the grain orientation layout given a set of welding parameters. The grain orientation layout is not a representation of individual grains; it is a grid structure composed of locally homogenised cells, each containing their dominant grain orientation. The associated ray-based model utilises the grain orientation map to compute the times-of-flight (ToF), accounting for the changes of the direction-dependent wavespeed.

While Ogilvy and MINA weld descriptions are efficient, they require parameter tuning and are limited in their ability to capture the complexity of the weld. Alternatively, grain orientations can be characterised directly from ultrasonic measurements by solving the inverse problem: given the experimentally measured times-of-flight, we seek to obtain the associated grain orientations or welding parameters. The predicted values are then employed to compute updated delay laws that can support imaging, using e.g. the Total Focusing Method (TFM) [11].

Several approaches have been developed to perform the inversion. They were usually formulated as optimisation problems, that is, minimising an objective function, [12–15]. In [16], Fan et al. focused on updating the four parameters in the MINA model whose influence on the ToF was most significant. They then proceeded to employ a global optimisation process based on the genetic algorithm, iteratively solving the ray-tracing model for the generated set of input parameters. The technique produced an optimal set of parameters whose computed ToF best matched the experimental (target) ToF.

In [5], Kalkowski et al. developed a two-stage physics-based workflow called weld map tomography. In the first stage, their generalised Ogilvy model was introduced to provide a starting orientation layout, and their cost function was evaluated with respect to the dominant orientation defined in that model. In the second stage, each cell orientation and resulting arrival times were updated independently of the weld description. The second stage was an implementation of time-of-flight tomography, in which the gradient of the objective function was calculated by back-projecting the difference between measured and computed arrival times across a propagation path, considering anisotropy. Although their approach was sensitive to the initial guess, the two-stage setup aided convergence and provided a significant improvement to imaging, as compared to the straight ray path assumption.

Depending on the context, the gradient-based approach can be advantageous because it provides a process that is usually transparent and explainable (not a black box) and maintains mathematical rigour. Although these methods are generally effective at solving the inversion task, their computational cost can be significant, and they tend to be unstable and affected by local minima. Moreover, their application is subject to problem-specific constraints, thus requiring the optimisation procedure to be significantly modified for each new inspected case. Therefore, their use may be limited in contexts where real-time inspection is indispensable.

Emerging data-enabled technologies have shown promising results in tackling some of the limitations of the traditional optimisation framework. The obvious advantage is the potential to support scalable real-time tasks. However, they usually require large datasets and carry the risk of providing unreliable predictions due to a decision-making process that is difficult to interpret (black-box). Still, applying trained models to experimental data can be highly advantageous in obtaining material properties [17]. Motivated to explore a different approach to the weld characterisation problem, this paper investigates whether neural networks can reliably extract features of the ToF maps that enable linking them to grain orientations.

Neural networks have been employed for weld map inversion before. In [18], Singh et al. used deep neural networks (DNNs) to infer grain orientations from ToF maps [18,19]. Synthetic polycrystals were generated using Voronoi tessellations with randomly oriented cells, discretised into a 16×16 grid and smoothed to reduce artefacts. Training

data consisted of orientation maps and ToF maps computed using the anisotropic multi-stencil fast marching method (AMSFMM) [12]. One DNN was trained per output pixel, each mapping the full ToF map to a single orientation. A generative adversarial network (GAN) was trained to upscale the predicted 16×16 maps to 64×64 , substantially improving resolution and accuracy of the DNN predictions.

In their following work [19], Singh et al. employed a similar workflow to reconstruct weld grain orientation maps from ToF maps. The weld maps used for training were generated by randomly perturbing a single reference grain orientation layout derived from the Ogilvy description. Their approach required training one DNN per orientation cell, achieving an average orientation prediction error of 2.61° . The DNN model was also used to predict the grain orientation from ToF data obtained from homogenised FEM simulations of a weld with embedded defects. The weld map obtained was utilised to update the ultrasonic delay laws in the TFM calculation, achieving an improvement of up to 5.3 dB.

Another application of DNNs for microstructure characterisation is explored in [20], where Patel et al. employed a neural network model to compute crystallographic orientation directly from surface wave velocity measurements obtained by Spatially Resolved Acoustic Spectroscopy (SRAS). SRAS usually employs the brute force search algorithm (BFSA) to determine the grain orientations from surface acoustic waves measurements, which is a bottleneck in the process. By employing DNNs, Patel et al. reported similar outcomes as the BFSA but with significantly faster computation times.

Previous studies have demonstrated the strong potential of machine learning for extracting material properties from experimental data. However, several key aspects remain unexplored: (i) previous works assumed known elastic tensor properties, which is rarely the case in practice; (ii) predictions were limited to determining grain orientations only, and did not include information about material or geometry; (iii) previous models were tested only with homogenised numerical data, not experimental measurements.

We propose two inversion workflows: the first predicts grain orientations directly from ToF maps assuming prior knowledge of the elastic constants, and fixed geometry. The second predicts the parameters of a generalised Ogilvy description, but also the elastic tensor and selected geometric parameters. Each inversion workflow involves a single DNN model for the entire domain. This model accepts the full time-of-flight map as an input, and predicts the full grain orientation map (first workflow) or the tensor and weld parameters (second workflow).

We evaluate two network architectures for the weld inversion task: a fully-connected neural network (FCNN) and a convolutional neural network (CNN), the latter to leverage the CNN's ability to learn from structured spatial data [21]. The trained models are evaluated using ToF maps obtained from grain-scale FEM simulations (with a fully controlled microstructure) and experimental measurements. The experimental data come from an industry-relevant austenitic stainless steel weld with drilled artificial defects. Ultrasonic reconstruction of the weld serves as an input for updating the delay laws and correcting the ultrasonic image for defect localisation.

The paper is organised as follows: after a brief review of current approaches to soft computing for material characterisation, we introduce the considered setup in Section 2. Next, in Section 3, we discuss the weld description and the forward model implemented to generate the training, validation, and testing datasets along with a data-based sensitivity analysis of the metamodel. Subsequently, two inverse tasks are presented: (i) obtaining grain orientations from ToF maps (Section 5), and (ii) obtaining weld parameters from ToF maps (Section 6). For each task, we assess the performance of their respective deep learning models by evaluating prediction errors in (a) train-like unseen datasets, (b) k-fold cross-validation estimates, and (c) FE grain-scale data. The paper is concluded with an experimental case study in Section 7, aiming at improving the signal-to-noise ratio and defect signatures of an industry-relevant weld by employing the developed models.

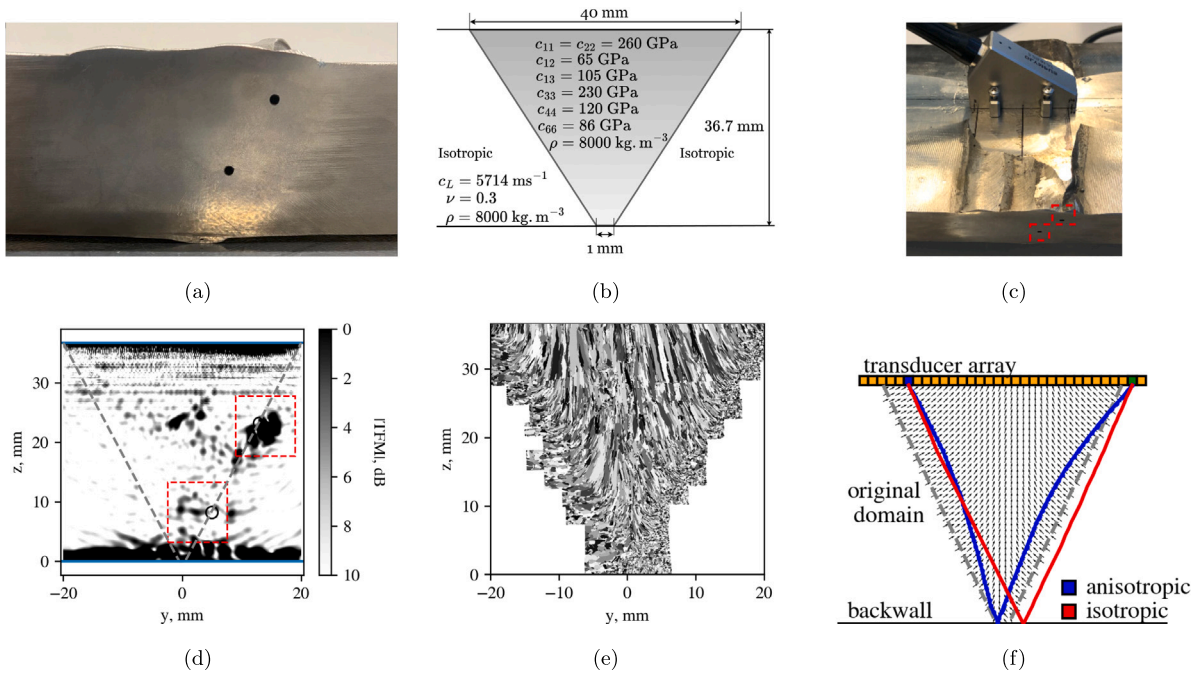


Fig. 1. Weld mock-up: (a) weld cross-section with artificial defects drilled through the right chamfer; (b) weld dimensions and material properties; (c) 64-element array transducer positioned centrally over the weld; (d) total focusing method imaging with isotropic assumption; (e) EBSD imaging of the weld; (f) ultrasonic beam paths assuming isotropy (red) and anisotropy (blue). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

2. Weld characterisation setup

In this section, we analyse the austenitic stainless steel mock-up weld [5] presented in Fig. 1(a) using the traditional ultrasound imaging workflow, where the weld is assumed homogeneous and isotropic. Later in Section 7, we will show how this workflow can be improved by accounting for the weld's heterogeneity and anisotropy through an inversion process based on neural networks. The weld has two holes side-drilled through the weld's right chamfer to simulate target defects (see Figs. 1(a) and 1(c)). Fig. 1(b) presents the weld dimensions and material properties - the latter are generally unknown. Fig. 1(c) shows the experimental setup for the characterisation of the weld.

The measurements were carried out with a 2.25 MHz 64-element array with a pitch of 0.75 mm. The array was in direct contact with the top of the sample (using a standard gel couplant), with the weld cap machined into a flat surface, as shown in Fig. 1(c). The ultrasonic data were acquired using full matrix capture (FMC), in which signals are recorded across all array elements for successive excitations of each individual element. The data was subsequently post-processed using the total focusing method (TFM) [11] for image reconstruction as per Eq. (1), which assumes an isotropic medium:

$$I(y, z) = \left| \sum_{tx, rx} h_{tx, rx} \frac{\sqrt{(y_{rx} - y)^2 + z^2} + \sqrt{(y_{tx} - y)^2 + z^2}}{c_L} \right|, \quad (1)$$

where h is the analytic signal, whose real and imaginary parts are related by the Hilbert transform of the acquired time series transmitted by the element tx , and received by the element rx . The fraction that follows defines the delay laws: the numerator computes the propagation distance from the transmitting element at position $(y_{tx}, 0)$ to each grid reflector point in the domain (y, z) , and then to the receiving element at position $(y_{rx}, 0)$. The denominator c_L is the constant wave velocity in the medium.

The isotropic assumption is clearly indicated in the calculation of the delay laws, as it presumes straight beam paths (numerator) and a constant wave velocity throughout the medium (denominator).

Employing this approach yields the TFM image in Fig. 1(d), and we observe that the ultrasonic energy is scattered, resulting in an image with elevated structural noise and attenuation. Note how the top chamfer defect is unfocused, and the bottom one is not visible.

As discussed in Section 1, the TFM image can be corrected if the material properties of the weld are incorporated into the TFM computation. Fig. 1(e) shows the EBSD imaging of the weld, and reveals its microstructure: a collection of elongated grains with spatially varying dominant orientations. These columnar grains with preferential orientations are responsible for distorting the ultrasonic beam. Fig. 1(f) illustrates the impact of the columnar grains on the beam path (blue) in comparison with the isotropic beam path (red). The discrepancy between beam paths is significant, demonstrating the need to account for the material information when computing the delay laws. Since EBSD is a destructive procedure, we need to extract that information by other means. As discussed in Section 1, we explore ultrasonic reconstruction of grain orientations using two model-based deep neural network (DNN) inversion workflows. The models supporting this inversion are discussed next.

3. Modelling anisotropic representations of the weld

As with any deep neural network training process, a substantial volume of data is essential to ensure model robustness and generalisation. This section presents the forward modelling strategies for data generation, and the considerations for error quantification: first, for efficiency, we use the shortest ray path (SRP) model [5,9,22] to generate ToF maps for training, validation, and testing; the SRP model is based on the generalised Ogilvy description of weld orientations - a simple but versatile model for initial studies. While we employ the Ogilvy description, the process is general and well suited for use with other grain orientation models, including, e.g., MINA. To support testing while avoiding the inverse crime [23], we also present a time-domain grain-scale finite element model to mimic experimental measurements, including all the rich physics of wave scattering in heterogeneous anisotropic media. Lastly, this section concludes with a discussion about the data generation procedure and their impact on training.




Fig. 2. (a) Single v-weld as per generalised Ogilvy model [5].

3.1. Ogilvy model and weld configuration

Fig. 2 illustrates the V-weld configuration examined in this study, along with its corresponding geometric parameters. These parameters are a generalisation of the Ogilvy map [9], which allows for representing anisotropic heterogeneous materials with an equivalent locally homogenised medium. In addition to these geometric characteristics, we also incorporate the four (five in 3D) independent material parameters of the transversely isotropic tensor of the weld material in 2D, as well as the longitudinal wave speed (c_L) in the isotropic parent material.

The single-V weld depicted in Fig. 2 follows the generalised Ogilvy map configuration presented in [5]. The parameter δ determines the relative position of the transducer array with respect to the weld centreline. The grain orientations θ at each location in the y - z plane are calculated as:

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

where T_L and T_R represent the tangents of the orientation angles on the left and right chamfers, respectively. α_L and α_R are the left and right chamfer angles, while d_L and d_R denote the lengths of the left and right weld beads. η_L and η_R correspond to the rates of change of the orientation angle when moving horizontally from the chamfer towards the centreline. The parameter a is the weld thickness, and ζ is the dominant orientation, reflecting the preferred grain growth direction. See [5,9] for more details on the derivation of Eq. (2).

For training purposes, we make no distinction between the left and right parameters (i.e., $\eta_L = \eta_R$, $T_L = T_R$, $\alpha_L = \alpha_R$, and $d_L = d_R$). It does not necessarily impose a symmetric grain layout as ζ shifts the angles along a dominant orientation. Table 1 presents the ten variables of interest for training.

3.2. Shortest ray path model

The shortest ray path model processes the grain orientation layout defined by the Ogilvy model (Eq. (2)) to compute time-of-flight maps. It accounts for directional variations in wave speed due to material anisotropy, as defined by the grain orientation layout, and employs

Table 1

Weld parameters of interest for training.

Parameter	Symbol	Unit
Chamfer angle	α	$^\circ$
Array offset	δ	mm
Dominant orientation	ζ	$^\circ$
Angle orientation tangent	T	–
Orientation change rate	η	–
Longitudinal wave speed	c_L	m s^{-1}
Independent elastic constants	$C_{22}, C_{33}, C_{23}, C_{44}$	GPa

Fermat's principle [2,14,22] to compute the propagation path between two points that minimises travel time.

The effect of heterogeneity becomes important when there is material orientation deviation over a significant spatial distance; hence, the precision at the local scale is not as important as one might think, and there is no significant advantage in pursuing locally accurate representations [5,16]. As demonstrated in [3,5,24], grains with similar orientations (Fig. 3(a)) can be described by a macroscopic equivalent representation of the grain topology, resulting in a map of homogenised cells (Fig. 3(b)), often referred to as weld maps [16]. The behaviour of the homogenised cell is defined by a transversely isotropic material with its soft axis aligned with the dominant grain orientation. In [5], Kalkowski et al. demonstrate that SRP and FE ToF computations converge when a 2 mm grid cell is employed, yielding ToF errors well below 0.05 μs , in line with the accuracy target reported in [22]. Therefore, a 2 mm grid size is employed in all SRP computations.

The inversion uses the pulse-echo configuration, with a 64-element array. Although this arrangement limits the range of viewing angles [25] - and consequently limits the amount of material information that can be recovered - the pulse-echo setup remains a practical choice given the area available around the weld [5,26]. The transducer array works as transmitter and receiver; for computation purposes, the backwall reflection is modelled by mirroring the domain around the backwall of the weld, as illustrated in Fig. 3(c). This approach is adopted to comply with the Fermat principle and allow the computation of the wave propagation path back to the transducer. The solid blue line indicates the incident wave path, while the solid green line indicates the reflected wave path from the backwall. The SRP model is used to generate ToF data in a controlled setting, where the causal relationship between weld parameters and ToFs is explicitly defined. This fast and deterministic workflow underpins the training of deep neural networks.

3.3. Finite element modelling

Grain-scale simulation is performed using the finite element method. Contrary to homogenisation (Fig. 3(b)) adopted in the SRP model, the FE models grain topology directly (Fig. 3(a)). It acts as a virtual experiment by considering the full physics of elastic wave propagation through a textured, heterogeneous, anisotropic medium. We implement a numerical workflow analogous to a full matrix capture [11], simulating the action of a 64-element array with a pitch of 0.75 mm. Each element sequentially applies a three-cycle toneburst with a central frequency of 5 MHz from the top boundary of the weld.

Fig. 4(a) presents the grain orientation layout of the FE model, constructed using an equivalent homogenised configuration defined by the weld parameters presented in Table 2. Each cell in this layout was assigned a standard Poisson-Voronoi tessellation comprising 40 grains. The microstructure of each cell was elongated by a factor of 10 and rotated to match the grain orientations (see Fig. 3(a)). The weld model shown in Fig. 4(a) was obtained by aggregating these individual grain-populated cells at their respective locations. The grain microstructure was generated using Neper, an open-source software for polycrystal generation and meshing [27].

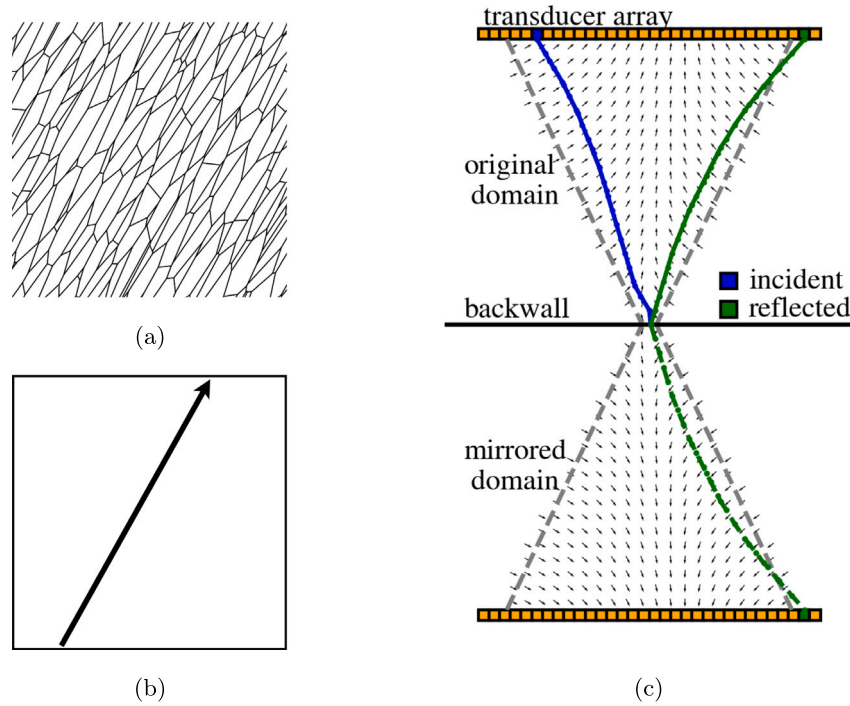


Fig. 3. SRP model: (a) grain-scale representation of the grid cell; (b) homogenised representation of the grid cell, which is employed in the SRP computations, and (c) the pulse-echo set-up and computation strategy with mirrored domain. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 2
V-weld parameters for FE model.

Parameters	α	δ	ζ	T	η	C_{22}	C_{33}	C_{23}	C_{44}	c_L
Units	$^\circ$	mm	$^\circ$	–	–	GPa	GPa	GPa	GPa	m s^{-1}
FEM model	27.5	0	-7.3	0.8	1.1	241.3	228.5	91.2	94.7	5668.1

The weld model was discretised with a regular quad mesh (CPE4-type linear elements) with an element size of $25\ \mu\text{m}$, totalling over 9 million degrees of freedom. The model satisfied the 10 elements per wavelength and the 10 elements per linear grain dimension criteria suggested by published studies [28–31]. Absorbing boundary conditions were applied to the left, right, and upper edges. The simulations were performed in Pogo [32], a GPU-based code for fast modelling of ultrasound. Each element excitation requires an independent FEM computation, and the full 64-element FMC computation was completed in approximately 12 min (see Table 7 for workstation specs). The inversion process developed in this paper uses the ToF map which first needs to be manually extracted from time traces. The accuracy of FMC data (either simulated or measured) post-processing directly affects the ToF maps, and consequently, the performance of the DNN model (and any ToF map-based inversion).

The ToF extraction process is often time-consuming owing to backscattering and grain-induced attenuation. However, the numerical setup allows for using auxiliary simulation of a model without a backwall, which can then be subtracted from data computed with backwall. This operation eliminates backscatter preceding the first arrival and simplifies post processing 4(b). Thus, a second set of simulations is executed by adding absorbing boundaries to the model depicted in Fig. 4(a).

3.4. Generation of training data

The weld model and its resulting time-of-flight maps are governed by a set ten independent parameters: the weld chamfer angle (α), the

transducer array offset (δ), the dominant grain dominant orientation (ζ), the tangent of the orientation angles on the chamfers (T), the rate of change of the orientation angle (η), the longitudinal wavespeed in the parent material (c_L), and the transversely isotropic elastic constants of the weld material C_{22} , C_{33} , C_{23} , and C_{44} (see Table 1).

Before we attempt the inverse problem, which is ill-posed and challenging, we seek to understand how sensitive the time-of-flight maps are to perturbations in the input. This helps to comprehend the role of each parameter, supports the interpretation of the subsequent DNN model predictions, and assists in the evaluation of the model performance by accounting for the associated errors and confidence intervals. Table 3 presents the 10 weld parameters and their bounds, adopted to generate the training dataset through the Ogilvy-SRP model. The sensitivity of the forward model is evaluated using the Sobol sampling method, which generates low-discrepancy sequences to ensure a uniform distribution and mitigate clustering effects [33].

Identifying parameters with low sensitivity informs possible dimensionality reduction in subsequent modelling, and a lower number of observations may be required to represent the featured space [33]. However, the sensitivity computations can be significantly affected by parameter sampling ranges [34]. Table 3 shows the bounds of the weld parameters spanning relatively broad ranges which need to be considered during interpretation.

Model sensitivity is analysed using Sobol indices [33] (implemented in the SALib python library [35,36]). The first- (S_i) and second-order (S_{ij}) Sobol indices are defined by Eqs. (3) and (4), respectively [33,34]:

$$S_i = \frac{D_i}{D}, \quad (3)$$

$$S_{ij} = \frac{D_{ij}}{D}, \quad (4)$$

where D_i is the first-order variance of each parameter (p_i), and D_{ij} is the second-order variance of each parameter interactions with respect to each other (p_i, p_j). The denominator D is the total variance of $f(\mathbf{p})$, which accounts for all higher-order interactions:

$$D(f) = \sum_i D_i + \sum_{i < j} D_{ij} + \sum_{i < j < k} D_{ijk} + D_{12\dots p}. \quad (5)$$

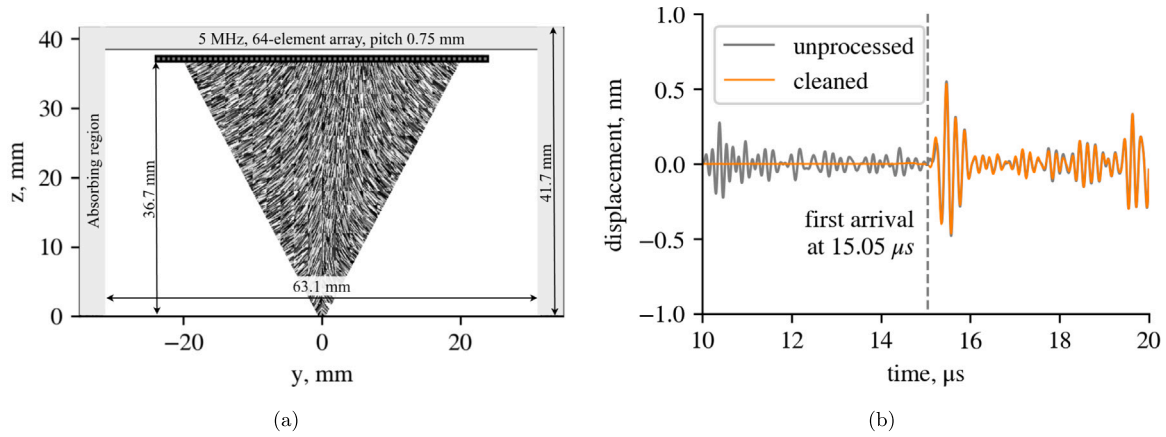


Fig. 4. FE modelling: (a) grain-scale model, and (b) signal processing strategy, showing the unprocessed and cleaned simulated time traces transmitted from transducer 4 and received by transducer 61.

Table 3
Weld parameter space.

Parameters	α	δ	ζ	T	η	c_L	C_{22}	C_{33}	C_{23}	C_{44}
Unit	°	mm	°	–	–	m s^{-1}	GPa	GPa	GPa	GPa
Bounds	[23, 32]	[–8, 8]	[–10, 10]	[0.8, 1.3]	[0.8, 1.3]	[5650, 5900]	[240, 280]	[215, 245]	[91, 115]	[74, 98]
τ	2	1	2	0.1	0.1	20	2	2	2	2
Sample 1	27.949	4.957	–8.172	1.195	1.097	5804.926	259.475	235.837	92.222	90.667
$\Delta p = +1\%$	0.279	0.050	0.082	0.012	0.011	58.049	2.595	2.358	0.922	0.907

The first-order sensitivity index S_i measures the individual influence of each p_i on the output $f(\mathbf{p})$, while the second-order sensitivity index S_{ij} accounts for the contributions of the interactions between parameters [34,37]. The results of the sensitivity analysis are presented in Figs. 5(a) and 5(b). Since the output of $f(\mathbf{p})$ is a 64×64 image, the indices are calculated for each pixel and averaged. The average values and their associated standard deviation ranges are also displayed.

Fig. 5(a) shows that C_{44} has the highest first-order index ($S_{C_{44}} = 0.224$), while the chamfer angle α has the lowest first-order index ($S_\alpha = 0.021$). Since the Sobol indices are unitless and normalised from 0 to 1 ($\sum_{s=1}^n \sum_{i_1 < \dots < i_s} S_{i_1 \dots i_s} = 1$), the Sobol index $S_{C_{44}} = 0.224$ indicates that 22.4% of the total variance in the model output is associated with variations in C_{44} alone, and $S_\alpha = 0.021$ indicates that 2.1% of the total variance in the model output is associated with variations in α alone [37] - disregarding interactions with other parameters (i.e., the higher-order indices). The sensitivity analysis also indicates that material parameters ($c_L, C_{22}, C_{33}, C_{23}$, and C_{44}) take precedence over the geometric and Ogilvy parameters ($\alpha, \zeta, \delta, \eta$, and T). Thus, under the featured parametric space and sampling range, higher accuracy is required from the inverse model in determining the material parameters (higher output sensitivity) than in determining geometric parameters (lower output sensitivity).

Fig. 5(b) presents the second-order Sobol indices, capturing the effect of the interaction between p_i and p_j on the variation D . The colour map indicates the mean S_{ij} values for each $p_i - p_j$ combination. At each pixel, there are two dots indicating the associated standard deviation. While the material parameters exhibit high significance by themselves (see Fig. 5(a)), their combined effect is of less importance. Conversely, the geometric parameters, while not playing a dominant role individually, when combined with other geometric and material parameters, lead to high output variance.

In addition to the global Sobol sensitivity analysis, we conducted a local analysis, evaluating the effect of small changes of the inputs \mathbf{p} around a given p_i on model output (valid only for the presented case, assuming a linear behaviour in the neighbourhood of p_i). Table 3 presents a reference sample and its associated perturbation Δp . We are interested in analysing the inputs individually, so we apply one

perturbation at a time, resulting in eleven output responses (reference plus ten perturbations). The perturbed ToF results are subtracted from the reference ToF result to assess the influence of changes in each parameter p_i on the ToF map.

Fig. 5(c) show the results in absolute μs values for each parameter, i.e. indicates the mean ToF discrepancy in μs due to a 1% change in p_i . Here, $i = c_L$ occupies the top rank, where a 1% change in its value leads to a mean $\Delta\text{ToF} = 0.076 \mu\text{s}$. Conversely, the 1% change in the dominant orientation ζ leads to an insignificant change in ΔToF . Except for the parameter α , the model output is more locally sensitive to changes in material parameters than to changes in geometric parameters, as observed in the first-order Sobol index results.

Another perspective is added to the interpretation of the results by considering the range and tolerance associated with each parameter. While Sobol analysis is scale-invariant, the indices are affected by the initial ranges assigned to each input [34]. Consequently, parameters with wider bounds are likely to contribute more to the output variance, resulting in a higher S_i , as observed for c_L and the elastic constants C_{ij} in Fig. 5(a). This can lead to erroneous conclusions about parameters with lower indices, such as ζ and α , which appear to be less important to the model in Fig. 5(a). Adjusting the bounds to better reflect parameter relevance would break the link to the practical knowledge about the problem. More importantly, it will ignore the fact that Sobol indices quantify the effect of the parameter uncertainty on the changes in the model's output, rather than the output itself [38]. To emphasise that, we conducted local sensitivity analysis for the tolerances (τ) associated with each parameter, as shown in Table 3. These tolerances are chosen based on the practical constraints of typical thick-section weld configurations.

Fig. 5(d) presents the local sensitivity analysis when the model is perturbed by $\Delta p = 2\tau$, that is, when it accounts for the tolerances established for each parameter in Table 3. From this perspective, for example, varying δ by $2\tau_\delta = 2 \text{ mm}$ is more impactful than varying c_L by $2\tau_{c_L} = 40 \text{ m s}^{-1}$. In fact, all parameters exhibit comparable levels of relevance: δ climbed from 7th to 1st in its impact on ΔToF , while ζ , although climbing just one position, still displays significant relevance. Furthermore, the results indicate high variance throughout the 64×64

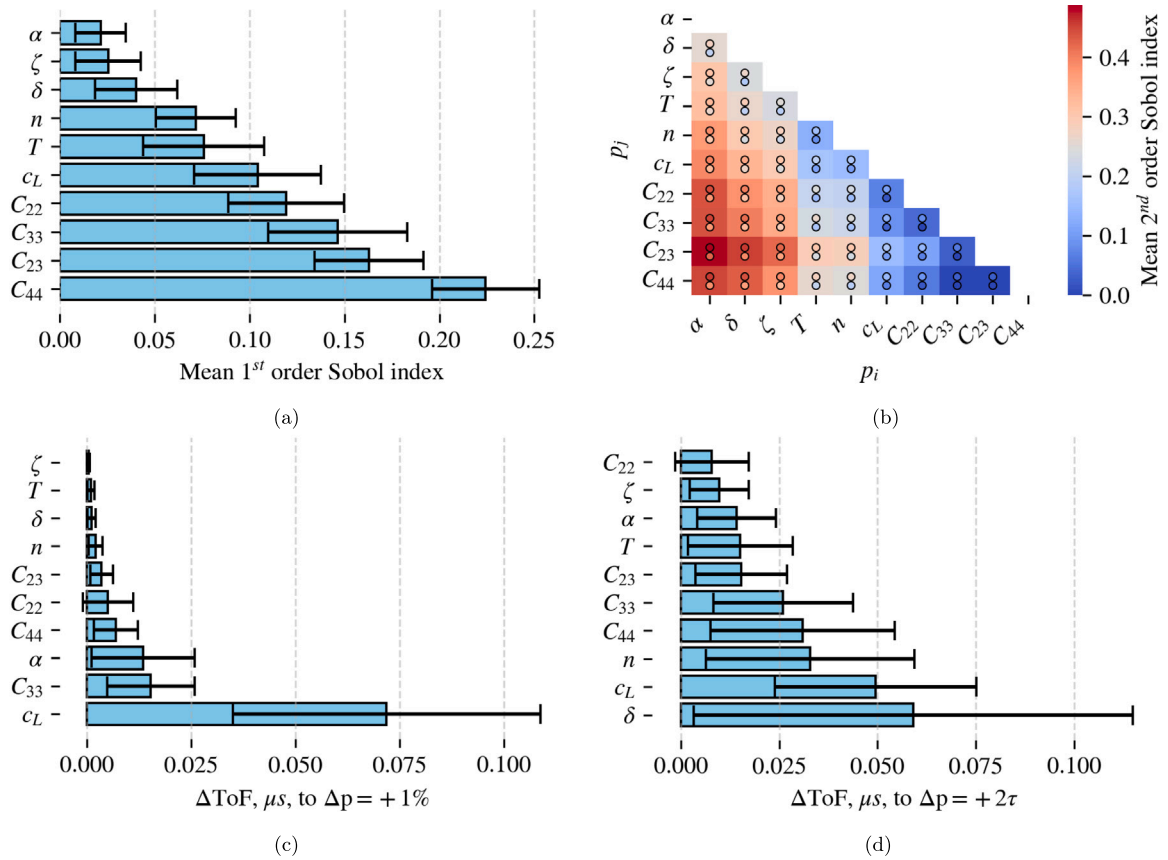


Fig. 5. Global sensitivity analyses: (a) first-order, and (b) second-order Sobol indices. First-order local sensitivity analyses: (c) local parameter step of 1%, and (d) local parameter step of $2r$ (see Table 3). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

ToF map, i.e., c_L is more relevant for paths passing predominantly through the isotropic region, while C_{22} is more relevant for paths crossing the weld region. This confirms that the weld parameters' sensitivity depends on the wave propagation path captured by each transmitter-receiver pair. Building on the insights from the sensitivity analysis, and the relationship between the inputs and outputs, we will evaluate the inversion models by their ability to make predictions within the tolerances specified in Table 3. Next, we describe the training strategies and neural network architectures used for inversion.

4. Deep neural network models

This section describes the deep neural networks used for the inversion tasks in Sections 5, 6, and 7. We use both the fully connected neural network (FCNN), and the convolutional neural networks (CNN), for the tasks of: (a) predicting grain orientation from time-of-flight data (Section 5), and (b) predicting the Ogilvy weld parameters and materials constants from the time-of-flight data (Section 6). While the FCNN is generally used for regression tasks, the CNN is known for its efficacy with structured spatial data. Our analysis focuses on their performance for both of the inversion tasks, and the discrepancy between predicted and true values are computed using the mean absolute error (MAE) and maximum absolute error (MaxAE) metrics.

Prior to training, a pre-processing step was conducted to prepare the data: in the task (a) of predicting grain orientation from ToF maps using the FCNN, both the input and target variables were standardised using the `StandardScaler` from the `scikit-learn` preprocessing module. When using the CNN, only the target data (grain orientation) were standardised. The transformed data were standardised to have zero mean and unit variance, which improved numerical stability during

training [39]. After training, the data is transformed back to their original representation.

In the task (b) of predicting the weld parameters from the ToF maps, only the weld parameters were scaled since they are defined over different ranges. This scaling procedure was applied according to the bounds established for each parameter (see Table 3), as follows:

$$\hat{p}_i = \frac{p_i - b_i^l}{b_i^u - b_i^l}, \quad (6)$$

and were unscaled after training using:

$$p_i = \hat{p}_i \cdot (b_i^u - b_i^l) + b_i^l, \quad (7)$$

where p_i are the original parameter values, \hat{p}_i are the scaled values, and b_i^u and b_i^l are the upper and lower bounds for the i_{th} parameter, respectively.

4.1. Fully connected neural network

The fully connected neural network utilised in this study comprises an input layer, multiple hidden layers with non-linear activation functions, and an output layer with a linear activation function. Each hidden layer consists of a defined number of neurons (or units), which directly influence the model's capacity for learning and generalisation. The number of dense layers, neurons, and the activation function were determined through a hyperparameter optimisation procedure employing Bayesian optimisation, as implemented in the `KerasTuner` `BayesianOptimization` tuner [40,41]. A detailed description of the hyperparameter search space is listed in Table 4.

The best-performing hyperparameter configuration was used as an initial baseline to build the DNN models. Manual tuning was subsequently introduced to the DNN model based on empirical convergence

Table 4
Bayesian hyperparameter search space and selection criterion for the fully-connected models.

Hyperparameter	Type/Sampling	Search space (bounds/options)
Number of hidden layers, n_{hidden}	Integer	0 to 20
Neurons in first dense layer	Integer	16 to 400
Neurons per hidden dense layer	Integer	16 to 400
Activation function	Categorical	{ReLU, GELU}
Learning rate (Adam)	Continuous (log-uniform)	10^{-4} to 10^{-2}
Batch size	Integer	16 to 100
Loss function	Fixed	Mean squared error (MSE)
Optimiser	Fixed	Adam
Validation split	Fixed	0.2
Early stopping	Fixed	Patience = 20 (monitor: val_loss; restore best weights)
Selection criterion	Fixed	Minimum validation MAE across all trials

Table 5
Fully Connected Neural Network (FCNN) architectures and hyperparameter configurations for extracting grain orientations and weld parameters from ToF maps.

Component	ToF to grain orientation (see Section 5.1)	ToF to weld parameters (see Section 6.1)
Input layer	4096 input features (flattened 64×64 ToF map)	4096 input features (flattened 64×64 ToF map)
Initial dense layer	63 neurons, GELU activation	400 neurons, GELU activation
Hidden layers	2 Dense layers, each with 204 neurons and GELU activation	4 Dense layers, each with 158 neurons and GELU activation
Output layer	380 units (19×20 grain orientation map), linear activation	10 units (weld parameters), linear activation
Optimiser	Adam	Adam
Learning rate	4.95×10^{-4}	1×10^{-5}
Loss function	Mean Absolute Error (MAE)	Mean Absolute Error (MAE)
Batch size	72	49
Epochs (max)	5000	5000
Validation split	0.2 (20% of training data)	0.2 (20% of training data)
Callback	Early stopping (patience 100, restore best)	Early stopping (patience 20, restore best)
Regularisation	None	L2 kernel regularisation (0.005)

behaviour. For example, although the learning rate was bounded to the range 10^{-4} – 10^{-2} , a lower learning rate of 10^{-5} was applied to the second task (ToF to weld parameters). The final FCNN configuration for each regression task is described in Table 5.

Table 5 shows the (FCNN) architectures and hyperparameter configurations used in Sections 5.1 and 6.1 for extracting (a) grain orientation maps, and (b) weld parameters from ToF maps, respectively. While both receive ToF maps with 4096 features as input, the grain orientation map requires 380 output nodes, and the weld parameter vector requires 10 output nodes. However, Table 5 shows that the first inversion task requires a simpler network, with only 3 dense layers, whereas the second inversion task requires two additional dense layers, a significantly lower learning rate, and an L_2 kernel regularisation to prevent overfitting by penalising larger weights.

The higher complexity of the network for task (b) might be counter-intuitive at first, as extracting grain orientations required an output with 380 units, whereas extracting weld parameters required only 10. However, the latter had no knowledge of material parameters, adding significant complexity to the inversion. Also, image outputs provide dense supervision, where each pixel contributes to the loss function. In contrast, the weld parameters provide sparse supervision, which gives significantly less information about the underlying physical or latent space. Consequently, the inversion to weld parameters is more prone to ill-posedness due to fewer constraints and requires a more complex network.

Fig. 6 illustrates the ToF to grain orientation regression. Both the ToF and the grain orientation maps are flattened to accommodate the FCNN architecture and then resized back to their original shapes. This flattening step is characteristic of FCNNs, which treat each pixel of the input image as an individual neuron. In the FCNN architecture, each neuron is connected to every neuron of the subsequent layer, each treated as an independent observation [42].

The fully connected network offers flexibility in handling various data types, but any information associated with the input's relative

position is lost. Consequently, FCNNs have been deemed inefficient for spatially structured data, such as images [21]. Here, we deal with time-of-flight and grain orientation maps, which contain spatial information and can be represented as images. On a second note, we also deal with parametric representations, which map n-dimensional vectors to grain orientations and time-of-flight maps. Thus, we explore how this architecture performs in both tasks.

4.2. Convolutional neural network

Convolutional neural networks differ from fully connected neural networks in that they preserve spatial information through kernels. CNNs offer strong performance for object detection and image classification [43], although they have also been employed for regression tasks [44,45]. In our workflow, we aim to obtain grain orientation maps and 10-dimensional weld description vectors from time-of-flight maps. In the first inversion, we perform image-to-image regression, while in the second, image-to-vector regression.

FCNNs are strong candidates for regression tasks, while CNNs are preferred for image processing. We explore both individually in the image-to-image task, whereas for the image-to-parameter regression, a fully connected regression block is added to the convolutional architecture, as illustrated in Fig. 7. In the first task, we explore an architecture consisting of a number of Conv2D \rightarrow Conv2D \rightarrow MaxPool blocks with progressively increasing number of filters, followed by a set of deconvolution layers for gradual upsampling. Table 6 shows the CNN architectures and hyperparameter configurations used in Sections 5.2 and 6.2 for extracting (a) grain orientation maps, and (b) weld parameters from ToF maps, respectively.

In the image-to-image task, we employ a fully convolutional network comprising four convolutional blocks, one deconvolutional block, and a final convolutional output layer. Each of the first four convolutional blocks stacks two convolutional layers followed by a max-pooling layer. The convolutional layers are responsible for feature extraction




Fig. 6. Fully connected neural network layout for obtaining grain orientations from ToF data.




Fig. 7. Convolutional neural network layout with fully connected regression block for obtaining weld parameters from ToF data.

Table 6

CNN architectures and hyperparameter configurations for extracting grain orientations and weld parameters from ToF maps.

Component	ToF to grain orientation (see Section 5.2)	ToF to weld parameters (see Section 6.2)
Input	64×64 greyscale image ($64 \times 64 \times 1$)	64×64 greyscale image ($64 \times 64 \times 1$)
Encoder	Four convolutional blocks, each comprising two Conv2D layers (3×3 kernels, GELU activation) followed by MaxPooling2D (2×2). The number of filters increases progressively: 32, 64, 128, and 256. The final block uses dilation rate = 2.	Three convolutional blocks with Conv2D layers (3×3 kernels, GELU activation) followed by MaxPooling2D (2×2). Filter sizes: 32, 64, and 128. The final block uses dilation rate = 2.
Decoder	Four transposed convolutional layers (Conv2DTranspose) with decreasing filter sizes: 256, 128, 64, and 32. Each uses 3×3 kernels, GELU activation, and stride = 2 for upsampling.	None
Dense layers	None	Flatten layer followed by a dense layer with 128 neurons, and 4 dense layers, each with 100 neurons. GELU activation adopted.
Output layer	Conv2D(1), linear + resize to 19×20	10 units (regression target), linear activation
Optimiser	Adam	Adam
Learning rate	0.00005	0.00005
Loss function	Huber Loss	MAE
Batch size	64	16
Epochs (max)	5000	5000
Validation split	0.2	0.2
Callbacks	Early stopping (patience = 20, restore best)	Early stopping (patience = 20, restore best)
Regularisation	None	L2 (0.005) on first dense layer

by employing filters with a receptive field size of 3×3 and a stride of 1, while the max-pooling layers are responsible for the spatial size reduction with a stride of 2. The final convolutional block uses a dilation rate of 2, which increases the receptive field, by spacing the kernel elements to capture larger context [46].

The fourth block stacks three deconvolutional layers that upsample the map and reconstruct it into a locally homogenised grain orientation map. The deconvolutional layers employ a receptive field size of 3×3 with a stride of 2. The network is finalised with a convolutional layer with a linear activation function, followed by a resizing operation to match the grain orientation domain layout.

In the image-to-vector regression task, we employ a convolutional encoder followed by a fully connected regressor. The encoder consists of three convolutional blocks, each comprising a 3×3 convolutional

layer with a stride of 1, followed by a max-pooling layer with a stride of 2. These blocks progressively extract hierarchical spatial features from the input ToF maps. Instead of a deconvolutional decoder, the convolutional stack is followed by a fully connected network. This regressor begins with a transition dense layer containing 128 neurons, followed by four dense layers with 100 neurons each. The dense layers use the GELU activation, selected after iterations with other activation functions such as ReLU and Swish. The final output layer consists of 10 units, each representing a predicted weld parameter.

5. Predicting grain orientation from time-of-flight data

First, we discuss training and performance of networks estimating weld maps from ToF data as illustrated in Fig. 2. To ensure a uniform




Fig. 8. FCNN model performance: (a) grain orientation MAE map across 14,120 test samples, and (b) training and validation losses.

grid for the grain orientation map, the offset (δ) and chamfer angles (α) are kept constant, with values $\delta = 0$ mm and $\alpha = 27.5^\circ$. The training dataset is generated by computing the Ogilvy model across eight weld parameters, using the sample space defined in Table 3.

The whole dataset consisted of 45,183 training samples, 11,295 validation samples and 14,120 testing samples, totalling 70,598 observations. Additionally, we conduct a 5-fold cross-validation and test the trained model on a numerical dataset obtained from the grain-scale FE model introduced in Section 3.3.

5.1. Method 1: Fully connected neural network

The fully-connected neural network presented in Section 4.1 receives a flattened 64×64 ToF map and outputs a flattened 19×20 grain orientation map (weld map). Fig. 8(a) shows the absolute mean error of grain orientations, averaged over 14,120 test samples. Each pixel indicates the mean error at that location in their original physical units ($^\circ$). The model performed to a high degree of accuracy, exhibiting a mean absolute error and standard deviation of $0.12 \pm 0.06^\circ$ and maximum absolute error of 0.38° . Fig. 8(b) presents the training and validation losses based on the MAE metric, which indicate good convergence and no overfitting. These are promising results compared to the average error value of 2.61° reported in the literature [19]. The maximum error of 0.38° in Fig. 8(a) is located at the root of the weld, which is a reasonable observation due to the limited ultrasonic beam coverage in that region [47,48].

5.1.1. 5-Fold cross-validation

A 5-fold cross-validation is conducted to assess our model against biased predictions [49]. Here, the testing dataset is randomly distributed into $k = 5$ mutually exclusive subsets (folds). Each model is trained using $k - 1$ folds, with one fold designated for testing. Additionally, each model trained through the k-fold process is used to predict the 14,120 unseen testing samples, set apart before the k-split, presented in Fig. 9.

Fig. 9(a) shows the grain orientation error map for the third fold, which yielded the poorest performance with a mean absolute error and standard deviation of $0.21 \pm 0.12^\circ$ and a maximum absolute error of 0.68° . Yet, Fig. 9(b) shows that all models exhibited similar performance, comparable to the results obtained from the model trained on the full training set presented in Fig. 8.

5.1.2. Validation against finite element simulation

This section evaluates the DNN model using data from FE grain-scale simulation, representing a virtual FMC experiment. See Sections 3.1 and 3.3 for details on the weld configuration and the FE model. Fig. 10(a) presents the absolute ToF error map in μs with

respect to the ToF map obtained from the equivalent SRP model. The central region displays significant discrepancies, reaching a difference of $0.089 \mu\text{s}$. However, note that the FMC post-processing introduces noise and uncertainty, and that the grain-scale model generates more information and complexity than the SRP model, thereby affecting the inversion performance.

Despite all the inherent differences between the SRP and FEM models, the predicted weld map closely aligns with the true one, as confirmed by the absolute orientation error map presented in Fig. 10(b). Note that the mean absolute error and standard deviation are $1.0 \pm 1.2^\circ$, with a maximum error of 6.67° located at the root of the weld. The root of the weld experiences lower wave coverage and higher attenuation, thus making its visualisation more challenging and less precise. Consequently, regions with lower beam coverage are expected to display higher prediction errors.

An attempt is made to improve the predictions from cross-source data inputs (e.g., FE and experimental) by augmenting the training data with a random noise sampled uniformly from $[-0.05, 0.05] \mu\text{s}$. The model is retrained as represented by the shaded region in Fig. 10(c), reaching lower training and validation loss values. Fig. 10(d) shows that the mean absolute error and standard deviation, and maximum absolute error display a slight improvement reaching $0.82 \pm 0.86^\circ$, and 4.54° , respectively.

5.2. Method 2: Convolutional neural network

As with the FCNN architecture in Section 5.1, this section evaluates the CNN through three key steps: training on simulated data, assessing generalisation via k-fold cross-validation, and testing its ability to predict a grain orientation map from a ToF input generated by FE grain-scale simulation.

Fig. 11 presents the performance of the model on simulated data. The grain orientation error map in Fig. 11(a) displays a mean absolute error and standard deviation of $0.27 \pm 0.17^\circ$ and a maximum absolute error of 0.75° . The training and validation losses in Fig. 11(b) indicate the model's convergence and no overfitting.

The FCNN model (Fig. 8) outperformed the CNN model (Fig. 11(a)), although the difference was minor. The error maps from both models displayed larger discrepancies at the root of the weld, which is associated with low ray coverage. However, the CNN prediction errors are slightly more pronounced and distributed throughout the weld layout. A probable cause is CNNs' processing of the image, considering the spatial structure via pooling layers, whereas FCNNs process each pixel individually. While pooling layers help address overfitting and increase computational efficiency [50], they may lead to a loss of detail in regions with sharp transitions [51].

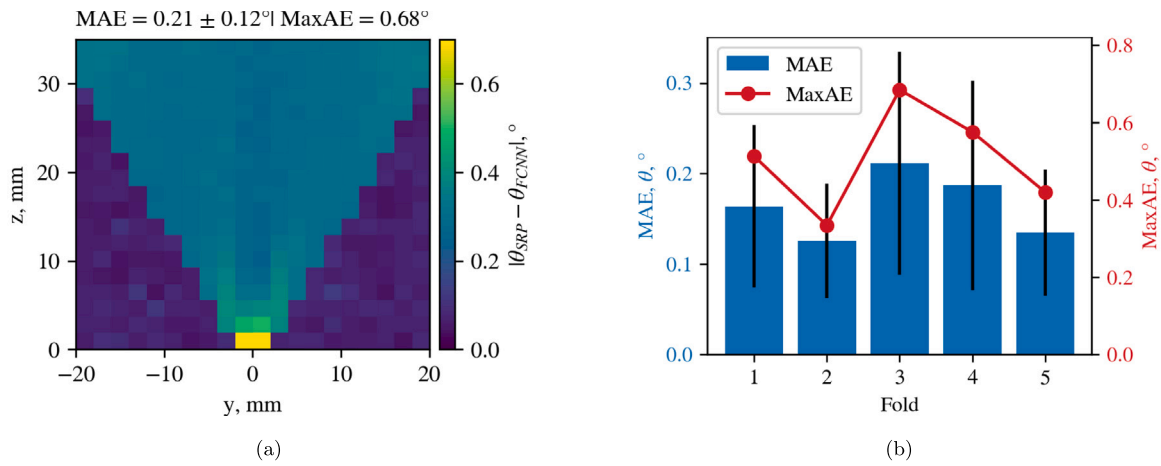


Fig. 9. 5-fold cross-validation error analysis: (a) grain orientation MAE map as predicted by the 3rd fold, and (b) grain orientation MAE and MaxAE per fold.

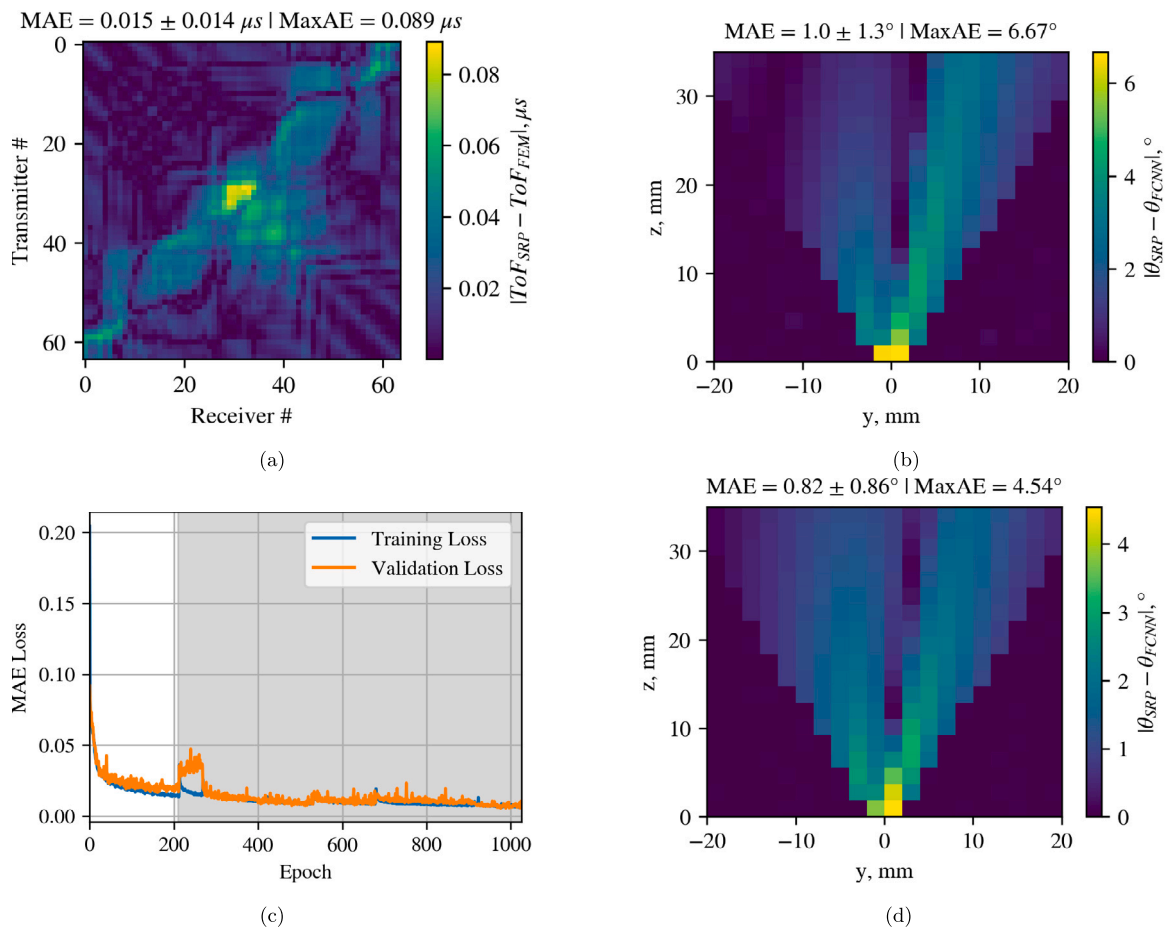


Fig. 10. FCNN performance on grain-scale FE data: (a) ToF MAE between SRP and FEM results. (b) MAE between grain orientations from reference dataset and FCCN prediction. (c) Training and validation MAE loss, with shaded region (epochs > 200) indicating further training performed using the augmented dataset. (d) MAE between grain orientations from reference dataset and FCCN model retrained with augmented dataset.

5.2.1. 5-fold cross-validation

The 5-fold validation is performed for the CNN architecture, yielding the results shown in Fig. 12. Fig. 12(a) displays the grain orientation MAE map generated by the third model, which demonstrated the weakest performance among all folds. This model yielded a mean absolute error and standard deviation of $0.44 \pm 0.29^\circ$ and a maximum absolute error of 1.58° , as depicted in the bar plot in Fig. 12(b).

Observe that the maximum absolute error of 1.58° is notably higher than the one observed for the model trained on the entire dataset. Thus, in addition to the slightly poorer predictions, the CNN model exhibits a lower generalisation and higher bias than the FCNN model for this regression task. Although the CNN’s MAE and MaxAE have been outperformed by the FCNN, the values remain within an acceptable range for inversion from ToF maps.

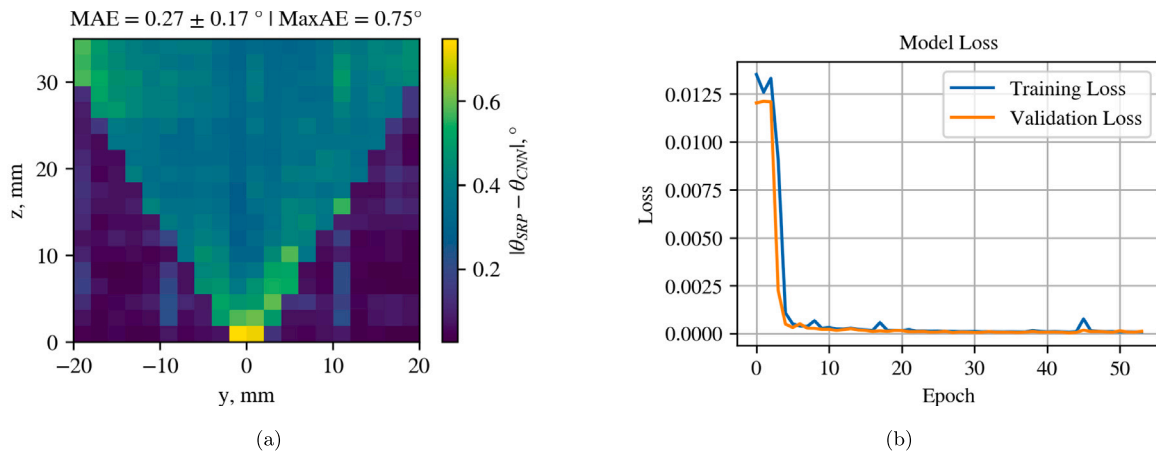


Fig. 11. CNN model performance: (a) grain orientation MAE map across 14,120 test samples, and (b) training and validation losses.

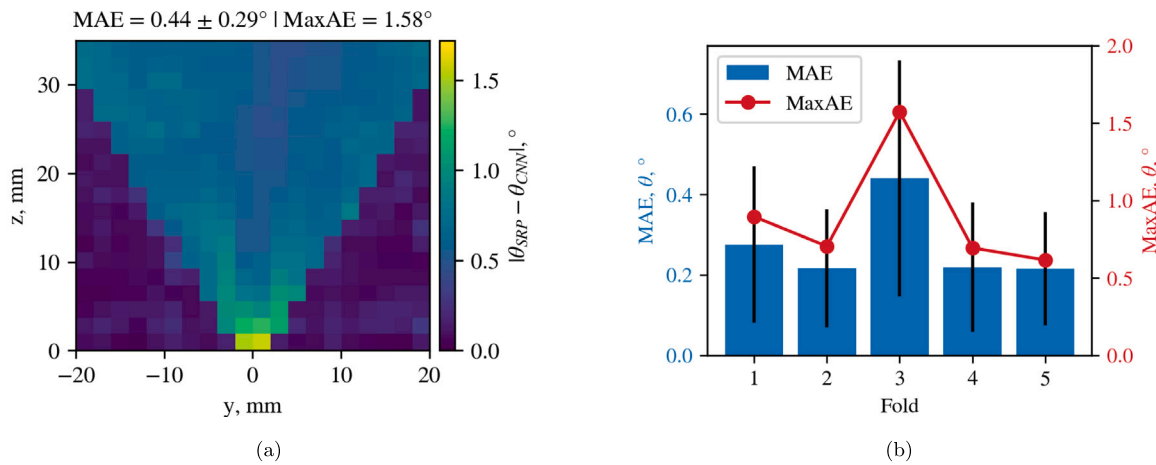


Fig. 12. 5-fold cross-validation error analysis: (a) grain orientation MAE map predicted by the third model, and (b) per-fold metrics: mean absolute error (left) and maximum absolute error (right).

5.2.2. Finite element validation

Fig. 13 shows the grain orientation MAE map predicted from the ToF data obtained via the FE grain-scale simulation. Although the FCNN outperformed the CNN on the testing dataset (compare Figs. 8 and 11), here the FCNN is outperformed by the CNN in predicting grain orientations. However, the performance gap is negligible: CNN yields a mean absolute error and standard deviation of $0.53 \pm 0.59^\circ$ and a maximum absolute error of 3.79° , compared to 0.82° MAE and 4.54° MaxAE for the FCNN.

Even though the performance of both networks is comparable, their error distribution is pronouncedly different. While both predictions yield higher errors at the root of the weld, the FCNN predictions exhibit a smoother transition with decreasing error levels for cells further away from the root. Conversely, the CNN seems to lack the ability to generalise the grain distribution, so the error map displays sharper transitions. This observation may be a consequence of the intrinsic CNN's characteristic to identify local features rather than global trends [52].

6. Predicting Ogilvy parameters, elastic tensor, and geometry from time-of-flight data maps

Section 5 demonstrated the ability of an FCNN and a CNN to predict grain orientations from ToF maps. However, because the network output layer has a fixed shape, the resulting grain orientation map is constrained to a predetermined layout of the weld. This is not an issue for the input layer, which is dictated by the phased array configuration,

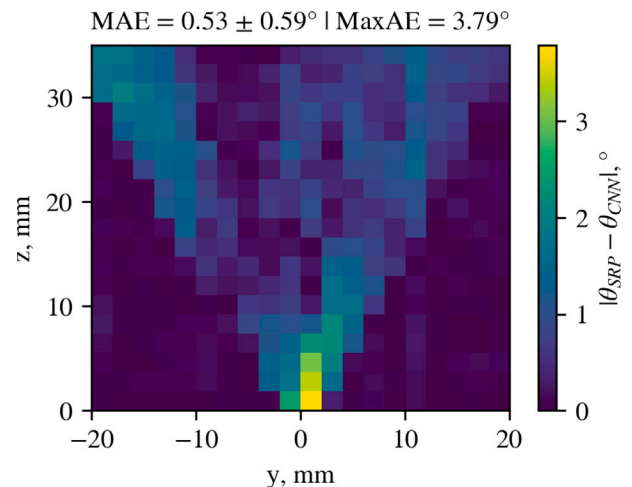


Fig. 13. Grain orientation MAE map predicted from the time-of-flight (ToF) data obtained via grain-scale FE simulation.

and the 64-element array is a standard. However, the shape of the weld map may vary depending on the relative position of the transducer and the weld and the weld's chamfer angle.

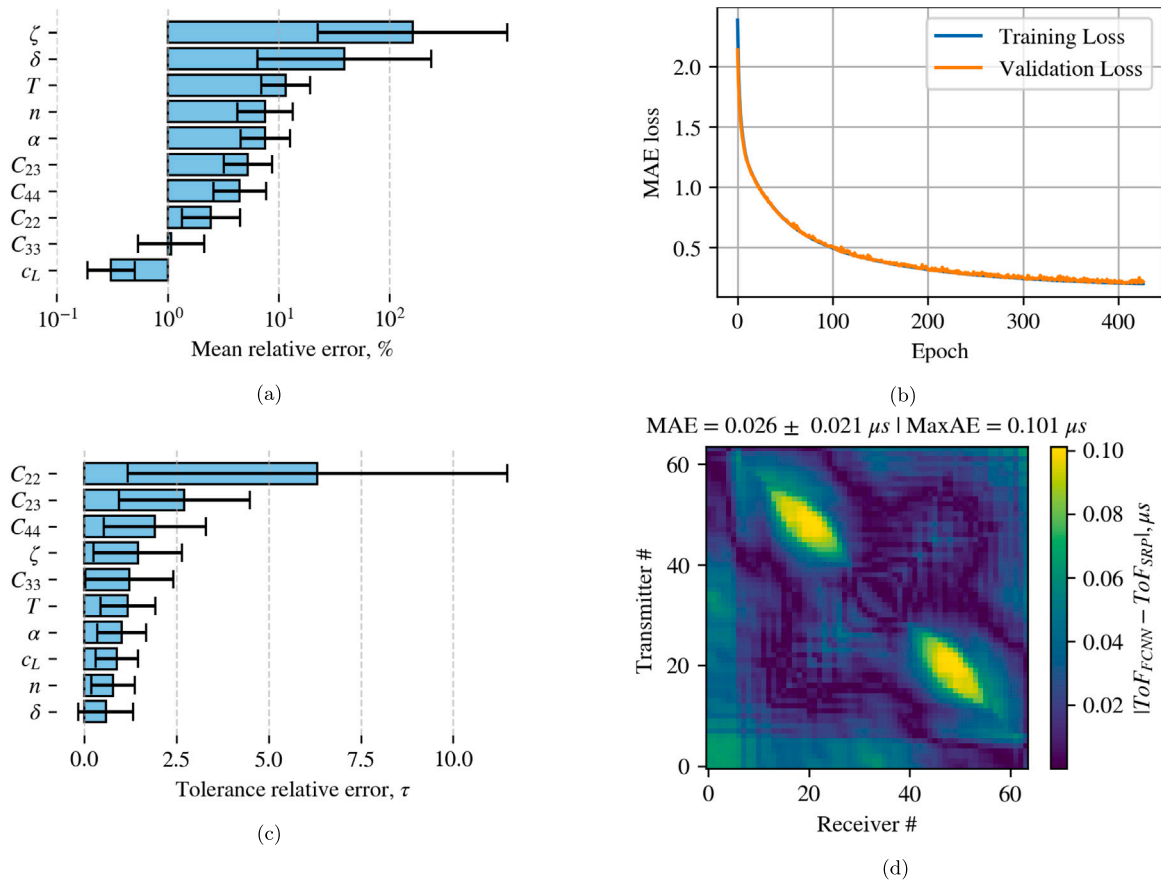


Fig. 14. Error analysis of the FCNN's predictions: (a) mean relative error between predicted and reference parameters; (b) MAE training and validation losses; (c) tolerance relative error between predicted and reference parameters; and (d) absolute error between the ToFs computed from the predicted and reference parameters. Note that (a) and (b) were computed for 18,023 testing observations, while (d) was computed for a single case.

To increase the robustness of the model and remove the constraints imposed on the weld map, we take a step further in the inversion and aim to predict the Ogilvy weld parameters, the transversely isotropic tensor, the array position, and weld geometry. This requires us to expand the parameter space from 8 parameters (Table 2) to 10 parameters by adding the chamfer angle (α) and the offset (δ) parameters, as illustrated in Fig. 2 and displayed in Table 1. The models are trained and evaluated on a dataset comprising 57,671 training samples, 14,418 validation samples and 18,023 test samples, totalling 90,112 observations.

6.1. Method 1: Fully connected neural network

The fully connected neural network receives a flattened 64×64 ToF map and outputs a 10-dimensional parameter vector. The network is composed of 7 dense layers: the first with 256 nodes and a L_2 regulariser of 0.005; the next five with 128 nodes; and the regression layer with 10 nodes to account for the 10 parameters. The first six dense layers employ the *GELU* activation function, while the regression layer utilises a linear activation function. We adopted the Adam optimiser with a learning rate of $1e-5$, and employed the MAE loss function. More details about the FCNN for this task are available in Section 4.1.

Fig. 14 presents the network performance. Fig. 14(a) shows the mean relative error given by the model when predicting the weld parameters. Fig. 14(b) presents the MAE of the training and validation losses, with good convergence and, thus, no sign of overfitting, which opposes the high error values observed in Fig. 14(a), especially for the δ with a $(39 \pm 199)\%$ relative error, and ζ with $(162 \pm 994)\%$, both with high standard deviations, meaning the model was unable to determine them with confidence.

However, the results displayed in Fig. 14(a) are better understood in light of the sensitivity study presented in Section 3.4. The sensitivity analysis led us to conclude that the model is more sensitive to the material parameters ($c_L, C_{22}, C_{33}, C_{23}$, and C_{44}) than to the geometric parameters ($\alpha, \zeta, \delta, \eta$, and T). In fact, the Sobol indices indicated that less than 5% of variance in the output was individually attributed to δ (4.0%) and ζ (2.5%). Additionally, the local sensitivity analysis yielded similar outcomes, with the model response to 1% perturbation in δ and ζ resulting in insignificant variations of the ToF map. Since small variations incurred by these two parameters go nearly undetectable by the network during training, it is more relevant to assess them using their associated tolerances, as presented in Fig. 14(c).

The error representation in Fig. 14(c) better describes how far the model is from providing a reliable prediction. In fact, only δ, η , and c_L exhibit $\tau < 1$. Another route to verify the performance of the inversion is achieved by passing the CNN-predicted parameters to the forward model to compute the associated ToF map. Fig. 14(d) presents the absolute error between the ToF computed from the predicted parameters and the expected ToF obtained from the reference parameters for a randomly selected observation. The result yielded a mean absolute error and standard deviation of $0.026 \pm 0.021 \mu s$ and a maximum absolute error of $0.101 \mu s$.

While the error levels in Fig. 14(d) are generally acceptable from a practical perspective, the performance of the FCNN in this context is certainly insufficient. Therefore, this model is not assessed through the 5-fold cross-validation or used to predict the parameters of the FEM ToF. Instead, in the next section, we explore how to improve the inversion using CNNs.

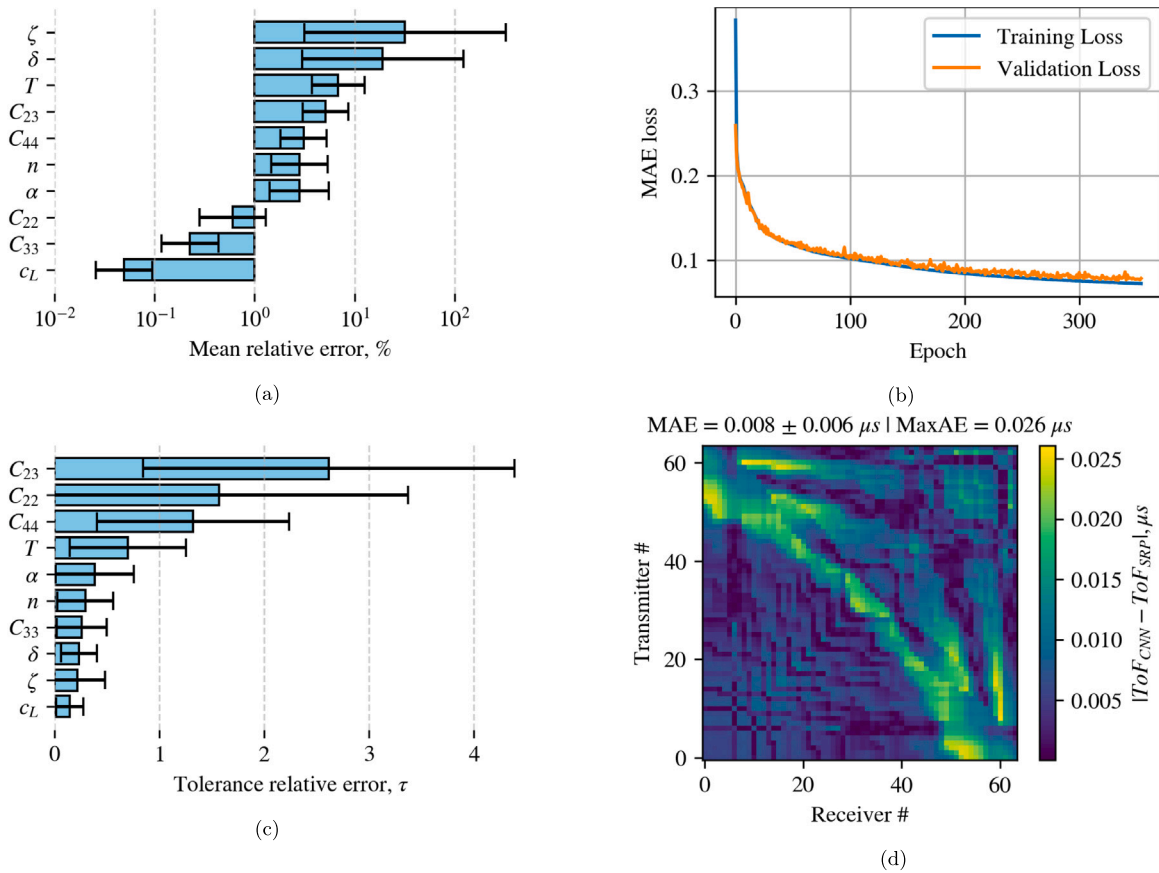


Fig. 15. Error analysis of the CNN's predictions. (a) mean relative error between predicted and reference parameters; (b) MAE training and validations losses; (c) tolerance relative error between predicted and reference parameters; and (d) absolute error between the ToFs computed from the predicted and reference parameters. Note that (a) and (b) were computed for 18,023 testing observations, while (d) was computed for a single case.

6.2. Method 2: Convolutional neural network

In this task, we are mapping an image to ten independent parameters, which results in structured data being turned into an input vector. As described in Section 4.2, this inversion procedure combines convolutional and fully connected layers, leveraging their respective strengths. Convolutional layers extract spatial features from the image, which are then downsampled using max-pooling layers. Finally, the flattened image is processed through fully connected layers to create a 10-dimensional output vector. A detailed description of the network is presented in Table 6.

Fig. 15 presents the performance of the network in learning the weld description parameters from the ToF map. We present the mean relative error in Fig. 15(a), which is significantly lower than the values displayed in Fig. 14(a) for the FCNN. This is also corroborated by the lower training and validation MAE loss to which the model converged during training, as shown in Fig. 15(b). However, the relative error values remains high in Fig. 15(a), particularly for δ ($19 \pm 103\%$), and ζ ($32 \pm 295\%$), both with high standard deviations. This indicates that the model cannot make precise predictions.

Analogously to Section 3.4, a more insightful way of looking at the data is to account for their associated tolerances. Fig. 15(c) shows that δ and ζ are among the parameters with the highest accuracy, i.e., $\tau \ll 1$ (see Table 3 for tolerance specifications). The material parameters C_{23} , C_{22} , and C_{44} are the only ones whose average tolerance-relative values are $\tau > 1$. Parent wavespeed c_L and C_{33} exhibit $\tau \ll 1$, which is supported by the fact that these are more physics-relevant for the model, that is, they are both directly governing wave propagation in the considered materials. Furthermore, the grain orientation parameters α , ζ , δ , η , and T exhibit low mean tolerance-relative error ($\tau < 1$).

The performance of the inversion is notably improved by the CNN, as evidenced in Fig. 15(d). Here, the predicted parameters are passed to the forward model to compute a ToF map. Fig. 15(d) shows the absolute error between the ToF map computed from the predicted parameters and the expected ToF obtained from the reference parameters for a randomly selected observation. The result yielded a mean absolute error and standard deviation of $0.008 \pm 0.006 \mu\text{s}$ and a maximum absolute error of $0.026 \mu\text{s}$, which are well below the error values observed for FCNN, and within reasonable tolerances for inversion tasks [22].

6.2.1. 5-fold cross validation

We now assess the ability of the model to generalise the predictions. Fig. 16 presents the 5-fold cross-validation performance of the CNN model. The error is calculated relative to the tolerance associated with each parameter and exhibits performance equivalent to the model trained on the full dataset (see Fig. 15(c)). No underfitting or overfitting is observed. Fig. 16 is the result of a further assessment on unseen observations, corroborating the estimates provided by the main model trained on the full dataset.

6.2.2. Finite element validation

The ToF map computed by the grain-scale FE model (Section 3.3) is now used as input to the trained CNN to predict the parameters of the FE weld model. The performance of the model is presented in Fig. 17, where Fig. 17(a) depicts the mean tolerance-relative error of each parameter, and Fig. 17(b) shows the absolute error between the ToF map computed by the forward model with the input parameters predicted by the CNN and the ToF computed from the FE simulation.

Fig. 17(a) exhibits a trend similar to those observed in Figs. 15(c) and 16, except for the parameter C_{33} , whose average tolerance-relative




Fig. 16. 5-fold cross validation of the CNN model.

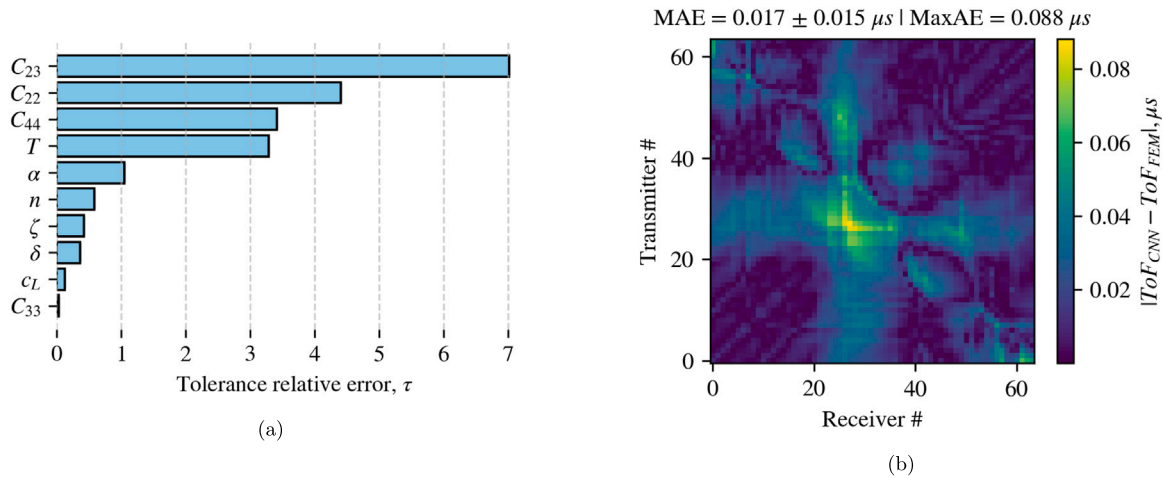


Fig. 17. CNN performance on FE data: (a) mean tolerance-relative error between predicted and true parameters, and (b) MAE between ToF map computed by the forward model with the input parameters predicted by the CNN and the ToF computed from the FE simulation.

error assumes the lowest value, and the parameter T , whose error is greater than 3τ . As for the parameters C_{44} , C_{22} and C_{23} , they carry higher tolerance-relative errors, reaching 7τ . Yet, the mean absolute error and standard deviation of $0.017 \pm 0.015 \mu s$, and the maximum absolute error of $0.088 \mu s$ in Fig. 17(b) are equivalent to the mean absolute error and standard deviation ($0.015 \pm 0.014 \mu s$), and maximum absolute error ($0.089 \mu s$) between the ToF obtained through the SRP model and the FE model, as shown in Fig. 10(a) - both generated with the exact same parameters.

The sources of discrepancy in Fig. 10(a) are known; they arise from the post-processing procedure applied to the FMC data and from the mismatch between the grain-scale FE model and the homogenised SRP model over which the network was trained. The error shown in Fig. 17(b) includes both of these sources and additionally incorporates the error propagated through the parameter prediction process. Considering the parameter prediction errors in Fig. 17(a), having achieved ToF results with the same levels of discrepancy displayed in Fig. 10(a) is a positive outcome. It also reinforces the conclusion that the material parameters c_L and C_{33} are more physics-relevant for the model, whereas the influence of C_{44} , C_{22} and C_{23} is minor.

Although the error ranking does not imply parameter importance on itself, the associated sensitivity analyses in Section 3.4 and the ToF error in Fig. 17(b) suggest that this metric can indicate which parameters are more influential. Yet, this is an isolated case and cannot be used to make definitive conclusions about the uncertainty around the predictions from FE data, especially because we do not fully understand how the parameter errors interact with the errors from post-processing the FMC and the mismatches in the grain-scale modelling. This is an analysis we will pursue in the future.

7. Experimental case study

In this section, we assess the DNN models presented in Section 4 on experimental data, i.e., an FMC dataset manually processed into a ToF map. We quantify their performance in correcting the delay laws for the austenitic weld sample introduced in Section 2. We evaluate the models from an application perspective, in their ability to improve the signal to noise ratio (SNR) and correct ultrasound images initially reconstructed under the isotropic assumption.

Following the procedure defined in [5], the SNR was calculated using the maximum of the image within ± 5 mm of the expected defect position, and the RMS value of the 10×10 mm² noise region (I_{noise}) centred around the maximum:

$$SNR = 20 \log_{10} \left(\frac{\max(I_{sig})}{\sqrt{I_{noise}^2}} \right). \tag{8}$$

Four TFM images are compared (Fig. 18), with delay laws computed based on: (i) isotropic medium; and anisotropic medium with grain orientation obtained from (ii) EBSD examination; (iii) FCNN-based grain orientation inversion from ToF; and (iv) CNN-based full weld inversion from ToF map. Under the isotropic assumption, we can identify the upper defect in Fig. 18(a). However, the lower defect position is unclear due to low SNR. The TFM image in Fig. 18(b) is computed with grain orientation information from an EBSD examination. Here, both artificial defects are evident, and we also note two naturally occurring defects from the welding process. Evidently, the use of EBSD data is incompatible with non-destructive testing. Thus, we utilise it as a benchmark for the inversion process. Note that the EBSD examination is also affected by errors related to the examination procedure and sample preparation, so deviations from the actual structure are still possible.

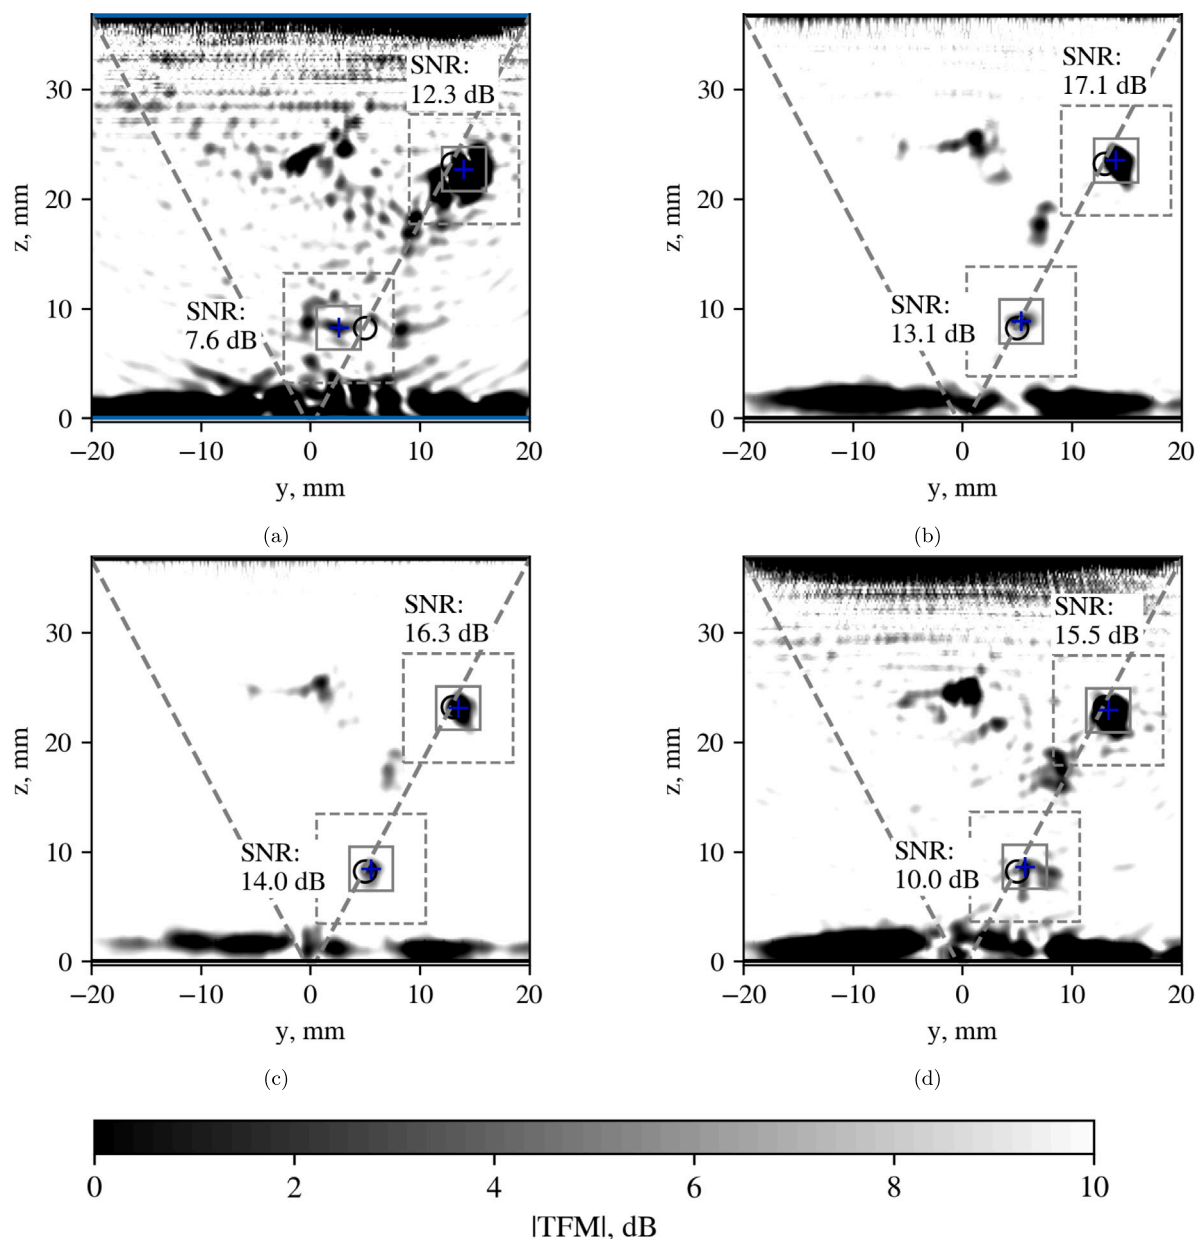


Fig. 18. Total focusing method imaging for defect localisation with (a) isotropic delay laws, and delay laws updated by grain orientations from (b) EBSD examination, (c) FCNN-ToF to grain orientation inversion, and (d) the CNN ToF map to weld description inversion.

Next, we employ the inversion models, and demonstrate their potential to correct ultrasound image from time-of-flight data. The TFM image in Fig. 18(c) is computed with grain orientation information from the FCNN model presented in Section 5.1. The model receives the ToF map obtained from post-processing the experimental FMC dataset and outputs the predicted grain orientation map. The orientation map is used to update the delays in the TFM computation. Compared to Fig. 18, the SNR has improved by 4 dB for the upper defect, and of 6.4 dB in the lower defect. Furthermore, using delay laws from inversion considerably reduced noise in the image and improved the location and focus of defect signatures.

The TFM image in Fig. 18(d) is computed with grain orientation information from the CNN inversion model (see Section 6.2). The CNN model receives the ToF map and outputs the weld parameters (See Table 1). The FCNN model required prior knowledge about the material constants, as only the grain orientations were predicted. In contrast, the CNN model provides the full generalised weld description, which includes the material parameters, grain orientation, and weld geometry

through the chamfer angle α . Despite the challenging task, the CNN model was capable of improving the SNR by 3.2 dB for the upper defect, and by 2.4 dB for the lower defect, as presented in Fig. 18(d).

Although the DNNs were trained on relatively simple homogenised models, they performed well when presented with measured data and allowed the identification of weld parameters, material properties, and geometry. The improvements in Fig. 18(d) are not as pronounced as in Fig. 18(c), but both effectively enable a more accurate defect signature location, noise reduction, and more focused signatures in comparison with the image obtained with the standard isotropic delay laws in Fig. 18(a).

8. Discussion and conclusion

Table 7 presents a summary of the four trained models, including the architecture, the shape of the input and output, the number of training, validation, and testing samples, as well as the training time,

Table 7
Training summary.

Section	Model	Input	Output	Number of samples			Training time	Inference latency	FLOPs	Parameter count
				Training	Validation	Testing				
5.1	FCNN	ToF (64 × 64)	$\theta(p)$ (19 × 20)	45 183	11 295	14 120	7.6 min	42.97 ms	1.0e6	1.1e6 (4.47 MB)
5.2	CNN						44.6 min	42.28 ms	5.6e6	6.4e6 (24.6 MB)
6.1	FCNN	ToF (64 × 64)	p (10.)	57 671	14 418	18 023	15.7 min	46.18 ms	2.3e6	3.4e6 (13.15 MB)
6.2	CNN+FCNN						75.5 min	43.49 ms	8.1e7	3.5e6 (13.57 MB)

Training was conducted using a 13th Gen Intel(R) Core(TM) i7-13850HX CPU, 2000 NVIDIA RTX Ada Generation Laptop GPU, and 32 GB RAM.

^a The dataset was generated using IRIDIS High Performance Computing at the University of Southampton.

and the prediction time to process one ToF map. The training times varied across the models and tasks, however they were relatively short in duration. Although the FCNNs trained faster than the CNNs, and the first task trained faster than the second, the difference in prediction times were negligible, at around 0.04 s.

The FCNN and CNN architectures exhibited similar performance on the first inversion task, both for the numerical and experimental observations. Since the FCNN model is simpler and trains quicker, it is the preferred option for this task. This conclusion diverges from the standard interpretation that CNN dominates image-to-image tasks. It also help us understand how the structure of the data influences the performance of DNN architectures: the pixels in the ToF and grain orientation maps are globally intertwined, i.e., all output and input pixels are related; and the pixels in both images depend on their position. Despite being an image-to-image regression, these observed long-range dependencies and positional perception favours the use of FCNNs [52] whereas CNN prioritises local patterns [53].

On the second inversion task, two fundamental changes were observed in the output/input relationship: (i) the global dependency was reduced, e.g., ToFs from neighbouring array elements positioned on the parent material depended only on c_L ; and (ii) the positional perception was removed due to the now varying (δ) phased-array location. Considering (i), (ii), the varying parameter sensitivity, and the image-to-vector regression format, a combined CNN-FCNN model was the logical choice to leverage on the relationship between output and input and the capabilities of both networks to extract local and global features. As indicated by the sensitivity analysis, only parameters with clear impact on the output such as δ , c_L , and C_{33} were determined to a high degree of accuracy.

Given the challenging task of a multi-dimensional regression, the performance of the models was promising, achieving error levels lower than those from the literature. Furthermore, the demonstration on the experimental dataset in Section 7 confirmed the effectiveness of the inversion in correcting TFM images — yielding improved image quality, and defect size and location. While the results were encouraging, further analysis is required in setups more representative of the industrial context, e.g., when the weld cap has not been removed, and under other transducer configurations. Especially for the second task (ToF to parameters), we foresee an expansion in application due to the flexibility of the model and reduced initial assumptions.

CRedit authorship contribution statement

Lucas Q. Machado: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Thomas Blumensath:** Writing – review & editing, Methodology. **Vykintas Samaitis:** Writing – review & editing, Investigation, Data curation. **Michael J.S. Lowe:** Writing – review & editing, Supervision, Methodology, Funding acquisition, Conceptualization. **Michał K. Kalkowski:** Writing – review & editing, Writing – original draft, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This research was supported by funding from the Euratom Research and Training Programme 2021–2027 (grant agreement no 101061359). University of Southampton and Imperial College London were funded by UK Research and Innovation under the UK Horizon Europe Guarantee Scheme (grant numbers 10044428 and 10040812). The authors acknowledge the use of the IRIDIS High Performance Computing Facility, and associated support services at the University of Southampton, in the completion of this work.

Data availability

Data will be made available on request.

References

- [1] Diaz J, Chassignole B, Fouquet T, Schumm A, Duwig V. Structural noise in modalisation. In: Proceedings of the 9th European conference on NDT. Berlin, Germany; 2006, URL <https://www.ndt.net/?id=3873>.
- [2] Connolly GD, Lowe MJS, Temple JAG, Rokhlin SI. The application of fermat's principle for imaging anisotropic and inhomogeneous media with application to austenitic steel weld inspection. Proc R Soc A: Math Phys Eng Sci 2009;465(2111):3401–23. <http://dx.doi.org/10.1098/rspa.2009.0272>.
- [3] Chassignole B, El Guerjouma R, Ploix M-A, Fouquet T. Ultrasonic and structural characterization of anisotropic austenitic stainless steel welds: Towards a higher reliability in ultrasonic non-destructive testing. NDT & E Int 2010;43(4):273–82. <http://dx.doi.org/10.1016/j.ndteint.2009.12.005>, URL <https://www.sciencedirect.com/science/article/pii/S0963869509001698>.
- [4] Li C, Wang R, Liu Z, Ji H, Qin S, Chen R, Zhong Y, Wu H. Study on phased array ultrasonic testing techniques for austenitic stainless steel butt welds. J Fusion Energy 2025;44(1):13. <http://dx.doi.org/10.1007/s10894-025-00475-6>.
- [5] Kalkowski MK, Lowe MJS, Samaitis V, Schreyer F, Robert S. Weld map tomography for determining local grain orientations from ultrasound. Proc R Soc A: Math Phys Eng Sci 2023;479(2277):20230236. <http://dx.doi.org/10.1098/rspa.2023.0236>.
- [6] Hamide M, Ecole Nationale Supérieure des Mines PF. Numerical modelling of steel arc welding. 2008.
- [7] Kalkowski MK, Bézi Z, Lowe MJS, Schumm A, Spisák B, Szavai S. Supporting imaging of austenitic welds with finite element welding simulation—Which parameters matter? Appl Sci 2023;13(13). <http://dx.doi.org/10.3390/app13137448>, URL <https://www.mdpi.com/2076-3417/13/13/7448>.
- [8] Ogilvy J. An iterative ray tracing model for ultrasonic nondestructive testing. NDT & E Int 1992;25(1):3–10. [http://dx.doi.org/10.1016/0963-8695\(92\)90002-X](http://dx.doi.org/10.1016/0963-8695(92)90002-X), URL <https://www.sciencedirect.com/science/article/pii/096386959290002X>.
- [9] Ogilvy J. Computerized ultrasonic ray tracing in austenitic steel. NDT Int 1985;18(2):67–77. [http://dx.doi.org/10.1016/0308-9126\(85\)90100-2](http://dx.doi.org/10.1016/0308-9126(85)90100-2), URL <https://www.sciencedirect.com/science/article/pii/0308912685901002>.
- [10] Moysan J, Apfel A, Corneloup G, Chassignole B. Modelling the grain orientation of austenitic stainless steel multipass welds to improve ultrasonic assessment of structural integrity. Int J Press Vessels Pip 2003;80(2):77–85. [http://dx.doi.org/10.1016/S0308-0161\(03\)00024-3](http://dx.doi.org/10.1016/S0308-0161(03)00024-3), URL <https://www.sciencedirect.com/science/article/pii/S0308016103000243>.

- [11] Holmes C, Drinkwater BW, Wilcox PD. Post-processing of the full matrix of ultrasonic transmit–receive array data for non-destructive evaluation. *NDT & E Int* 2005;38(8):701–11. <http://dx.doi.org/10.1016/j.ndteint.2005.04.002>, URL <https://www.sciencedirect.com/science/article/pii/S0963869505000721>.
- [12] Tant KMM, Galetti E, Mulholland AJ, Curtis A, Gachagan A. Effective grain orientation mapping of complex and locally anisotropic media for improved imaging in ultrasonic non-destructive testing. *Inverse Probl Sci Eng* 2020;28(12):1694–718. <http://dx.doi.org/10.1080/17415977.2020.1762596>, arXiv:<https://doi.org/10.1080/17415977.2020.1762596>.
- [13] Gueudre C, Le Marrec L, Moysan J, Chassignole B. Direct model optimisation for data inversion. Application to ultrasonic characterisation of heterogeneous welds. *NDT & E Int* 2009;42(1):47–55. <http://dx.doi.org/10.1016/j.ndteint.2008.07.003>, URL <https://www.sciencedirect.com/science/article/pii/S0963869508000753>.
- [14] Zhang J, Hunter A, Drinkwater BW, Wilcox PD. Monte carlo inversion of ultrasonic array data to map anisotropic weld properties. *IEEE Trans Ultrason Ferroelectr Freq Control* 2012;59(11):2487–97. <http://dx.doi.org/10.1109/TUFFC.2012.2481>.
- [15] Tant KMM, Galetti E, Mulholland AJ, Curtis A, Gachagan A. A transdimensional Bayesian approach to ultrasonic travel-time tomography for non-destructive testing. *Inverse Problems* 2018;34(9):095002. <http://dx.doi.org/10.1088/1361-6420/aaca8f>.
- [16] Fan Z, Mark AF, Lowe MJS, Withers PJ. Nonintrusive estimation of anisotropic stiffness maps of heterogeneous steel welds for the improvement of ultrasonic array inspection. *IEEE Trans Ultrason Ferroelectr Freq Control* 2015;62(8):1530–43. <http://dx.doi.org/10.1109/TUFFC.2015.007018>.
- [17] Holm EA, Cohn R, Gao N, Kitahara AR, Matson TP, Lei B, Yarasi SR. Overview: Computer vision and machine learning for microstructural characterization and analysis. *Met Mater Trans A* 2020;51(12):5985–99. <http://dx.doi.org/10.1007/s11661-020-06008-4>.
- [18] Singh J, Tant K, Curtis A, Mulholland A. Real-time super-resolution mapping of locally anisotropic grain orientations for ultrasonic non-destructive evaluation of crystalline material. *Neural Comput Appl* 2022;34(6):4993–5010. <http://dx.doi.org/10.1007/s00521-020-06670-8>.
- [19] Singh J, Tant K, Mulholland A, MacLeod C. Deep learning based inversion of locally anisotropic weld properties from ultrasonic array data. *Appl Sci* 2022;12(2). <http://dx.doi.org/10.3390/app12020532>, URL <https://www.mdpi.com/2076-3417/12/2/532>.
- [20] Patel R, Li W, Smith RJ, Clark M. Developing neural networks to rapidly map crystallographic orientation using laser ultrasound measurements. *Scr Mater* 2025;256:116415. <http://dx.doi.org/10.1016/j.scriptamat.2024.116415>, URL <https://www.sciencedirect.com/science/article/pii/S1359646224004500>.
- [21] Cantero-Chinchilla S, Wilcox PD, Croxford AJ. Deep learning in automated ultrasonic NDE – Developments, axioms and opportunities. *NDT & E Int* 2022;131:102703. <http://dx.doi.org/10.1016/j.ndteint.2022.102703>, URL <https://www.sciencedirect.com/science/article/pii/S0963869522001025>.
- [22] Nowers O, Duxbury DJ, Zhang J, Drinkwater BW. Novel ray-tracing algorithms in NDE: Application ofijkstra and A* algorithms to the inspection of an anisotropic weld. *NDT & E Int* 2014;61:58–66. <http://dx.doi.org/10.1016/j.ndteint.2013.08.002>, URL <https://www.sciencedirect.com/science/article/pii/S0963869513001163>.
- [23] Kaipio JP, Somersalo E. Inverse problems and interpretation of measurements. In: *Statistical and computational inverse problems*. New York, NY: Springer New York; 2005, p. 1–5. http://dx.doi.org/10.1007/0-387-27132-5_1.
- [24] Mark A, Fan Z, Azough F, Lowe M, Withers P. Investigation of the elastic/crystallographic anisotropy of welds for improved ultrasonic inspections. *Mater Charact* 2014;98:47–53. <http://dx.doi.org/10.1016/j.matchar.2014.09.012>, URL <https://www.sciencedirect.com/science/article/pii/S1044580314002873>.
- [25] Huthwaite PE. Quantitative imaging with mechanical waves. 2012. <http://dx.doi.org/10.25560/9765>, URL <http://hdl.handle.net/10044/1/9765>.
- [26] Aeronautics N, Administration S. Process specification for ultrasonic inspection of welds. Tech. Rep. PRC-6510, NASA; 2020, URL <https://www.nasa.gov/wp-content/uploads/2023/03/prc-6510-current.pdf>. [Accessed 4 May 2026].
- [27] Quey R, Dawson P, Barbe F. Large-scale 3D random polycrystals for the finite element method: Generation, meshing and remeshing. *Comput Methods Appl Mech Engrg* 2011;200(17):1729–45. <http://dx.doi.org/10.1016/j.cma.2011.01.002>, URL <https://www.sciencedirect.com/science/article/pii/S004578251100003X>.
- [28] Smith WD. The application of finite element analysis to body wave propagation problems. *Geophys J Int* 1975;42(2):747–68. <http://dx.doi.org/10.1111/j.1365-246X.1975.tb05890.x>, arXiv:<https://academic.oup.com/gji/article-pdf/42/2/747/1840090/42-2-747.pdf>.
- [29] Drozd MB. Efficient finite element modelling of ultrasound waves in elastic media. 2008, URL <http://hdl.handle.net/10044/1/7974>.
- [30] Van Pamel A, Sha G, Rokhlin SI, Lowe MJS. Finite-element modelling of elastic wave propagation and scattering within heterogeneous media. *Proc R Soc A: Math Phys Eng Sci* 2017;473(2197):20160738. <http://dx.doi.org/10.1098/rspa.2016.0738>, arXiv:<https://royalsocietypublishing.org/doi/pdf/10.1098/rspa.2016.0738>. URL <https://royalsocietypublishing.org/doi/abs/10.1098/rspa.2016.0738>.
- [31] Van Pamel A, Brett CR, Huthwaite P, Lowe MJS. Finite element modelling of elastic wave scattering within a polycrystalline material in two and three dimensions. *J Acoust Soc Am* 2015;138(4):2326–36. <http://dx.doi.org/10.1121/1.4931445>, arXiv:https://pubs.aip.org/asa/jasa/article-pdf/138/4/2326/13448901/2326_1_online.pdf.
- [32] Huthwaite P. Accelerated finite element elastodynamic simulations using the GPU. *J Comput Phys* 2014;257(PA):687–707.
- [33] Sobol I. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Math Comput Simulation* 2001;55(1):271–80. [http://dx.doi.org/10.1016/S0378-4754\(00\)00270-6](http://dx.doi.org/10.1016/S0378-4754(00)00270-6), URL <https://www.sciencedirect.com/science/article/pii/S0378475400002706>. The Second IMACS Seminar on Monte Carlo Methods.
- [34] Herman JD, Reed PM, Wagener T. Time-varying sensitivity analysis clarifies the effects of watershed model formulation on model behavior. *Water Resour Res* 2013;49(3):1400–14. <http://dx.doi.org/10.1002/wrcr.20124>, arXiv:<https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1002/wrcr.20124>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/wrcr.20124>.
- [35] Herman J, Usher W. SALib: An open-source python library for sensitivity analysis. *J Open Sour Softw* 2017;2(9). <http://dx.doi.org/10.21105/joss.00097>.
- [36] Iwanaga T, Usher W, Herman J. Toward SALib 2.0: Advancing the accessibility and interpretability of global sensitivity analyses. *Socio-Environmental Syst Model* 2022;4:18155. <http://dx.doi.org/10.18174/sesmo.18155>, URL <https://sesmo.org/article/view/18155>.
- [37] Santner TJ, Williams BJ, Notz WL. Sensitivity analysis, validation, and other issues. In: *The design and analysis of computer experiments*. New York, NY: Springer New York; 2003, p. 189–203. http://dx.doi.org/10.1007/978-1-4757-3799-8_7.
- [38] Tosin M, Côrtes AMA, Cunha A. A tutorial on sobol global sensitivity analysis applied to biological models. In: da Silva FAB, Carels N, Trindade dos Santos M, Lopes FJP, editors. *Networks in systems biology: Applications for disease modeling*. Cham: Springer International Publishing; 2020, p. 93–118. http://dx.doi.org/10.1007/978-3-030-51862-2_6.
- [39] de Amorim LB, Cavalcanti GD, Cruz RM. The choice of scaling technique matters for classification performance. *Appl Soft Comput* 2023;133:109924. <http://dx.doi.org/10.1016/j.asoc.2022.109924>, URL <http://dx.doi.org/10.1016/j.asoc.2022.109924>.
- [40] O'Malley T, Bursztein E, Long J, Chollet F, Jin H, Invernizzi L, et al. KerasTuner. 2019, GitHub repository. <https://github.com/keras-team/keras-tuner>.
- [41] Yu T, Zhu H. Hyper-parameter optimization: A review of algorithms and applications. 2020, CoRR abs/2003.05689. arXiv:2003.05689. URL <https://arxiv.org/abs/2003.05689>.
- [42] Maier HR, Galelli S, Razavi S, Castelletti A, Rizzoli A, Athanasiadis IN, Sánchez-Marrè M, Acutis M, Wu W, Humphrey GB. Exploding the myths: An introduction to artificial neural networks for prediction and forecasting. *Environ Model Softw* 2023;167:105776. <http://dx.doi.org/10.1016/j.envsoft.2023.105776>, URL <https://www.sciencedirect.com/science/article/pii/S1364815223001627>.
- [43] Alzubaidi L, Duan Y, Al-Dujaili A, Ibraheem IK, Alkenani AH, Santamaría J, Fadhel MA, Al-Shamma O, Zhang J. Deepening into the suitability of using pre-trained models of ImageNet against a lightweight convolutional neural network in medical imaging: An experimental study. *PeerJ Comput Sci* 2021;7:e715. <http://dx.doi.org/10.7717/peerj-cs.715>.
- [44] Cira C-I, Díaz-Álvarez A, Serradilla F, Manso-Callejo M. Convolutional neural networks adapted for regression tasks: Predicting the orientation of straight arrows on marked road pavement using deep learning and rectified orthophotography. *Electronics* 2023;12(18). <http://dx.doi.org/10.3390/electronics12183980>, URL <https://www.mdpi.com/2079-9292/12/18/3980>.
- [45] Jernelv IL, Hjelme DR, Matsuura Y, Aksnes A. Convolutional neural networks for classification and regression analysis of one-dimensional spectral data. 2020, arXiv:2005.07530. URL <https://arxiv.org/abs/2005.07530>.
- [46] Wei Y, Xiao H, Shi H, Jie Z, Feng J, Huang TS. Revisiting dilated convolution: A simple approach for weakly- and semi-supervised semantic segmentation. 2018, arXiv:1805.04574. URL <https://arxiv.org/abs/1805.04574>.
- [47] Connolly GD. Modelling of the propagation of ultrasound through austenitic steel welds Phd thesis, London, UK: Imperial College London, University of London; 2009, URL <https://www.imperial.ac.uk/media/imperial-college/research-centres-and-groups/non-destructive-evaluation/GDConnolly-thesis.pdf>. Submitted for the degree of Doctor of Philosophy and for the Diploma of Imperial College.
- [48] Connolly GD, Lowe MJS, Temple JAG, Rokhlin SI. Correction of ultrasonic array images to improve reflector sizing and location in inhomogeneous materials using a ray-tracing model. *J Acoust Soc Am* 2010;127(5):2802–12. <http://dx.doi.org/10.1121/1.3372724>, arXiv:https://pubs.aip.org/asa/jasa/article-pdf/127/5/2802/14777205/2802_1_online.pdf.
- [49] Verma VK, Saxena K, Banodha U. Analysis effect of k values used in k fold cross validation for enhancing performance of machine learning model with decision tree. In: Garg D, Rodrigues JJPC, Gupta SK, Cheng X, Sarao P, Patel GS, editors. *Advanced computing*. Cham: Springer Nature Switzerland; 2024, p. 374–96.

- [50] Akhtar N, Ragavendran U. Interpretation of intelligence in CNN-pooling processes: A methodological survey. *Neural Comput Appl* 2020;32(3):879–98. <http://dx.doi.org/10.1007/s00521-019-04296-5>.
- [51] Sunkara R, Luo T. No more strided convolutions or pooling: A new CNN building block for low-resolution images and small objects. In: *European conference on machine learning and principles and practice of knowledge discovery in databases*. ECML pKDD, 2022, p. 443–59.
- [52] Ding X, Xia C, Zhang X, Chu X, Han J, Ding G. RepMLP: Re-parameterizing convolutions into fully-connected layers for image recognition. 2022, [arXiv:2105.01883](https://arxiv.org/abs/2105.01883). URL <https://arxiv.org/abs/2105.01883>.
- [53] Zhang J, Li J. Chapter 13 - machine learning algorithm for cognitive engine. In: Zhang J, Li J, editors. *Spatial cognitive engine technology*. Academic Press; 2023, p. 169–85. <http://dx.doi.org/10.1016/B978-0-323-95107-4.00011-1>, URL <https://www.sciencedirect.com/science/article/pii/B9780323951074000111>.