



Interpolation techniques for ultrasonic data

Georgios Sarris^{*,1}, Michael J.S. Lowe², Peter Huthwaite³

¹Department of Mechanical Engineering, Imperial College London, Exhibition Road, South Kensington, London, SW7 2AZ, United Kingdom

ARTICLE INFO

Keywords:

Interpolation
In-silico data generation
NDE qualification
FE modelling

ABSTRACT

Many applications where ultrasound is used for diagnostics exist where limited data is preventing a particular approach from being fully exploited; for example, sufficient data availability would allow the qualification of non-destructive evaluation (NDE) methods in-silico, and would potentially also enable the training of machine learning algorithms related to ultrasound and its applications. Real, experimental ultrasonic data is often scarce, and while it is already known that finite element (FE) modelling produces data which is sufficiently realistic to augment real data, the computational cost associated with its generation at the scales required for the aforementioned purposes is often prohibitive. In this work, we propose the use of interpolation techniques in combination with results from FE modelling to rapidly generate more data without the need to solve additional FE models. We present the relevant methods to achieve this, and validate them through four exemplary cases of increasing complexity. Validation is achieved through the comparison of interpolation-generated results to those generated by full FE modelling, demonstrating that our method is capable of producing results for different physical setups and signals of various degrees of complexity. The results were typically within less than 1% away from the expected, but generated at a fraction of the typical computational cost, and, while the validation cases examined are of interest to the NDE community, the method extends to other fields where ultrasonic data is of interest.

1. Introduction

Safety cases are built on non-destructive evaluation (NDE) methods, and it is therefore critical that these are reliable: they must perform as they should in all scenarios to avoid unexpected failures. More specifically, this means that NDE methods need to be capable of detecting all critical defects in all anticipated cases [1], which includes various factors such as defects being in different positions and orientations as well as variations including environmental conditions. Inspection qualification is a formalised process, often matching an industrial standard, in order to establish that a particular technique meets its needs. Traditionally this involves testing physical test-pieces containing artificial defects, but because of the relative expense of their production these are limited and care must be dedicated to select the most challenging cases to conservatively demonstrate the technique. There are also issues with artificial defects not fully capturing the physics of their naturally occurring counterparts [2–4], and the quality and quantity of results being functions of the operators, whose performance has been shown to affect the results' reliability [3,5,6]. These factors make the qualification process conservative, and slow the qualification of new NDE methods.

Digital qualification aims to bypass these issues by utilising simulated data to qualify methods instead. This can reduce both the

time and cost to apply a new method in industry as well as limit conservatism, by allowing a much wider range of cases to be assessed. Although this represents a significant improvement, given the range of possible parameters in a problem, there is still a challenge to quickly generate sufficient data as these parameters vary.

Besides NDE-related applications, many research and industrial applications exist where simulated data is useful; this includes medical diagnostics [7–10], seismology [11–13], structural acoustics [14–16] and metamaterials [17,18], amongst others. Additionally, data availability for machine learning has become an issue, especially for ultrasonic applications, where real-world data is often either limited compared with the amount sufficing for training, or even if attainable, difficult to obtain in the required quantities [19].

While there do exist methods able to produce accurate results equivalent to those from experiments, such as statistical modelling [20–22], spectral collocation [23,24], surrogate models [25–27], and more recently, data augmentation [28], machine-learning related techniques [29–31] and combinations of the above [32], they are usually tailored to specific applications. A popular, general method for generating data in-silico for a plethora of applications is finite element (FE) modelling. FE modelling does not make any assumptions relating to the physics of

* Corresponding author.

E-mail address: g.sarris18@imperia.ac.uk (G. Sarris).

<https://doi.org/10.1016/j.ultras.2025.107823>

Received 15 April 2025; Received in revised form 10 September 2025; Accepted 11 September 2025

Available online 23 September 2025

0041-624X/© 2025 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

the model to be solved, besides the discretisation of the displacement field; any errors associated with this are well understood and can be addressed via suitable spatial and temporal mesh refinement [33]. This makes it suitable for generating results that fully capture the wave physics present in real, experimental measurements and that match existing theoretical models [34–38]. The main challenge with generating FE data, especially for applications in realistic qualification scenarios where many parameters need to be considered [20,39], is the associated computational cost.

In this work, we propose the use of mathematical techniques in combination with FE-generated data corresponding to strategically selected points within a parameter space, to interpolate the results between those and populate the parameter space fully. This provides a uniform approach by removing any prerequisites from the available data, as it is tailored to work with any FE generated ultrasonic signal as an input. The work presented here focuses on NDE-related applications, with a particular emphasis on phased arrays, which are inherently particularly computationally expensive to simulate, due to their multiple transmission and reception pairs. The underlying theory, however, is universal to any area where ultrasound is used, and hence we invite readers to recognise the application of our method to other fields where numerical modelling of ultrasound is commonly used.

This paper is organised as follows: Section 2 presents the relevant theory, both in terms of the interpolation itself and the methods used to apply this method to more complicated signals. Next, Section 3 discusses the FE models generated to evaluate the performance of our method, with the results and relevant comparisons being presented in Section 4. Finally, Section 5 summarises and concludes this work.

2. Theory

This section presents the theory and relevant equations used to implement our interpolation routine. First, the theory directly related to the interpolation itself is discussed, initially for interpolating across measurement locations and then for interpolating across any other variable in a parameter space. The challenges associated when working with signals containing multiple pulses are also discussed, and a suitable method to mitigate those are presented.

2.1. Interpolation

Parameter spaces in NDE-related problems comprise many dimensions; examples from the perspective of transmission/reception include transducer location and orientation or excitation frequency, while from a physical standpoint examples include defect location, orientation and size, number of scatterers or defects, or even environmental inputs like temperature. Analogous combinations of parameters are found in other ultrasonic applications outside NDE too. Here, we reduce the number of full FE models required to explore any of the aforementioned parameters; in this subsection we will focus on presenting the theory relating to interpolating across transducer location and will subsequently generalise this theory in the next subsection.

Let us suppose that having solved an FE model, the displacement at two nodes, at locations A and C, has been obtained and we wish to predict the displacement of a node located at B, at an arbitrary location between A and C, as shown in Fig. 1. For most of the results presented in this work, the measurement of interest will be the displacement of the nodes. However, the method should be valid for any physical measure from waves, such as velocity or pressure.

We denote the distance between A and B, and between B and C by x_{AB} and x_{BC} respectively, and the total distance between A and C by x_T , where $x_T = x_{AB} + x_{BC}$, as shown in the same figure. Denoting the displacements of the nodes at A and C by u_A and u_C respectively, we can predict the displacement at B, $\hat{u}_B$, using the standard definition of linear interpolation as follows:

$$\hat{u}_B = \frac{u_C - u_A}{x_T} x_{AB} + u_A. \quad (1)$$




Fig. 1. Schematic of two nodes, at locations A and C, whose displacements over time are known, and a node located at B where the displacement is to be determined via interpolation of the displacements at A and C. The distances between the nodes are also shown.

Eq. (1) can be rearranged, having $\hat{u}_B$ expressed as a weighted sum of the displacements at locations A and C:

$$\hat{u}_B = w_A u_A + w_C u_C. \quad (2)$$

In Eq. (2), the weights are a function of the distance of locations A and C with respect to B, and are given by

$$w_A = 1 - \frac{x_{AB}}{x_T}, \quad (3)$$

and

$$w_C = 1 - \frac{x_{BC}}{x_T}. \quad (4)$$

While Eqs. (1) and (2) are equivalent, it is easier to understand the physical implications of Eq. (2); say that location B was located at the same location as location A. In that case w_C would be equal to 0, and the displacement at location B would be the same as the known displacement at location A, as expected, since locations B and A are the same for this example, and vice versa if location B was selected at the same spot as location C. The weights for predicting the displacement at any location between A and C are defined such that the influence of the known displacements is a function of how close B is to them.

Interpolation of a physical field requires some form of inherent assumption about the underlying form of that field, whether explicitly stated or not. As a first approximation, we will consider a one dimensional variation of the form De^{ikx} where D is a complex amplitude, k is the wavenumber, x the propagation distance and i is the complex unit. This form can be considered to be fairly general for many parametric variations. For example, measuring across a wave incident at an oblique angle could be captured with that form, with a suitable wavenumber. Or, considering a change in temperature through the component would change the product kx , which could be approximated as a linear variation with temperature in the first instance. To account for this properly, we must manage this phase component appropriately in the interpolation. To better appreciate this, let us examine a simple, analogous case of interpolating two complex numbers, z_A and z_C , which inherently also possess both an amplitude and a phase, similarly to the time traces we are attempting to interpolate. Applying Eq. (2) to z_A and z_C in an attempt to find a complex number between them is equivalent to placing the complex number (z_B) onto the straight line connecting z_A and z_C , as shown on the Argand diagram in Fig. 2(a). However this Cartesian approach does not preserve neither the amplitude nor the phase. Instead, to obtain a meaningful answer z_B should be placed on the arc connecting z_A and z_C , as shown in Fig. 2(b). As shown schematically in the same subfigure, the arc is defined by the simultaneous interpolation of the amplitude and phase angle of z_A and z_C .

The same issue would arise when interpolating ultrasonic signals as the phase shift becomes significant, which is the case for almost all cases of interest. An example is shown in Fig. 3. In this figure, the known signals at locations A and C are plotted, as well as the signal resulting from applying Eq. (2). A comparison with the correct signal, assuming linear variation in amplitude and phase, and equal weights, is also shown. Since the phase difference between the signal at location A and the signal at location C is significant Eq. (2) gives results which are not physical.




Fig. 2. Argand diagrams of complex numbers z_A and z_C , and predicted complex number z_B between them. To estimate z_B , a Cartesian interpolation approach has been implemented in (a), while a polar interpolation approach has been deployed in (b). In (b), the amplitudes ($|z_A|$ and $|z_C|$) and phases (ϕ_A and ϕ_C) of z_A and z_C were interpolated to define the arc connecting z_A and z_C . The circle defined by z_A , z_B and z_C is also shown in grey.




Fig. 3. Interpolation result, $\hat{u}_B$, from Eq. (2), when applied to signals u_A and u_C for a case with a large phase difference between u_A and u_C . A comparison with the true signal, u_B is also shown.

As discussed above, the physically meaningful results in Fig. 3 would be achievable if the amplitude and phase were interpolated independently, before reassembling them to obtain the full solution. This becomes possible by working in the frequency rather than in the time domain; let $\mathcal{F}\{\cdot\}$ and $\mathcal{F}^{-1}\{\cdot\}$ denote the Fourier and inverse Fourier transforms respectively, and $|\cdot|$ and $\angle\cdot$ denote the amplitude and unwrapped phase respectively. Then, $\hat{u}_B$ is obtained from

$$\hat{u}_B = \mathcal{F}^{-1} \left\{ (w_A |U_A| + w_C |U_C|) e^{i(w_A \angle U_A + w_C \angle U_C)} \right\}, \quad (5)$$

where $U_A = \mathcal{F}\{u_A\}$, $U_C = \mathcal{F}\{u_C\}$ and $\angle\cdot \in (-\pi, \pi]$. It is worth noting that the term “unwrap” means that at every instance where two consecutive angles in the phase of either U_A or U_C are found, as the frequency increases, multiples of $\pm 2\pi$ are added until the difference is less than π . In our work, this was achieved by using the `unwrap()` function which is built into MATLAB (Natick, MA).

By utilising the fact that $w_A + w_C = 1$, Eq. (5) is written as:

$$\hat{u}_B = \mathcal{F}^{-1} \left\{ (w_A |U_A| + w_C |U_C|) e^{i\angle U_A} e^{iw_C(\angle U_C - \angle U_A)} \right\}. \quad (6)$$

By having a phase difference in the exponent, there will exist fewer instances of having phase differences larger than π , assuming that there have been no rapid changes between the signals recorded at A and at C. A challenge with phase unwrapping is when the signals change sufficiently rapidly that large phase jumps occur between samples, meaning that identifying where phase wrapping has occurred is difficult. One approach is to try to reduce the phase variation using the technique

in Eq. (6). Another approach is to zero-pad the signal such that the frequency domain signals are finely sampled. It should be noted that since we are dealing with simulated data, which will be relatively clean, we avoid the challenges associated with noise being present in the data.

Another advantage of using Eq. (6) is that we are no longer restricted by traditional, wavelength related, maximum allowable step sizes. If the underlying variation of kx with respect to the interpolation variable is indeed linear, then x_T can be arbitrarily large, since the phase term in Eq. (6) will perfectly capture the variation across any distance. In cases where this variation is not linear across the entire span of the dimension explored in the parameter space, good results are achievable by selecting a maximum step where the variation can be sufficiently approximated as linear. What is considered acceptable is application dependent. Additionally, since we are presenting a general signal interpolation method, it is likely that the maximum allowable step size is highly dependent on the application. Some cases may only require a rough estimate of the amplitude of the signal, while others may need nearly perfect estimation of the value of the signal at each time step. Therefore, it is infeasible to derive a maximum interpolation step which would be applicable to any general case. In practice, a user can establish an effective step size by performing a simple convergence test on a representative portion of the parameter space. By comparing the estimated signal against its true counterpart at when a large step is used, the user can confirm whether the assumption of local linearity holds for their required level of accuracy, and if not either the step can be reduced or interpolation avoided.

Implementing Eq. (6) instead of Eq. (2) results in the amplitudes and phases being treated independently and thus correctly matching the De^{ikx} form during the interpolation. It is worth noting that since the Fourier transform is a linear process, linear interpolation in the time and frequency domain are equivalent. This means that Cartesian interpolation is still possible in the frequency domain, by interpolating the Fourier transforms directly, i.e. replacing u_A and u_C with U_A and U_C respectively in Eq. (2), and will yield the same results as the time domain approach. However, it is exactly this ability to separate the amplitude and phase of the signals, and treat them separately, which makes the Fourier approach more desirable.

2.2. Interpolating across other variables

Eqs. (5) and (6), and specifically the weight terms, w_A and w_C , assume that the interpolation is implemented across a spatial dimension. However, the same equations are applicable regardless of the dimension of the parameter space over which the interpolation takes

place; for example, we can vary the size of a crack and interpolate the signals obtained at different crack sizes, inferring the corresponding signals at intermediate crack sizes, or, we can vary the tilt of a crack with respect to the propagation direction and by using the signals from a few tilt angles, infer the results at any intermediate angle.

Modifying Eqs. (3) and (4) is sufficient to achieve interpolation across any variable. The generalised form of these equations is

$$w_A = 1 - \frac{X_A}{\Delta X}, \quad (7)$$

and

$$w_C = 1 - \frac{X_C}{\Delta X}, \quad (8)$$

where X_A and X_C are the smallest and largest values of the parameter over which the interpolation is to take place, and $\Delta X = X_C - X_A$. For example, if we would like to obtain signals from a crack while varying its length, X_A and X_C would be the smallest and largest cracks for which we have true FE results respectively, while if obtaining signals from cracks at different tilt angles, X_A and X_C would be the smallest and largest tilt angles from which a full FE model has been solved respectively.

2.3. Signals containing multiple pulses

As discussed above, it was assumed that the signals to be interpolated follow the standard, e^{ikx} behaviour. There will be cases where the variation in the total phase, kx , with respect to the interpolation variable may not be common between the pulses comprising a signal. For example, since k is a function of the propagation speed, signals containing pulses which travelling at different speeds, such as multiple longitudinal and shear waves, and therefore different k values with respect to each other, will not be interpolated correctly if Eq. (6) is applied to the entire signal. In this example, the signals must be decomposed into their constituent pulses prior to interpolation.

To achieve this, each signal to be used in the interpolation is multiplied by a custom-made window, $w(t)$

$$w(t) = \begin{cases} 0 & 0 \leq t < t_p - \frac{t_w}{2} \\ T(t) & t_p - \frac{t_w}{2} \leq t \leq t_p + \frac{t_w}{2} \\ 0 & t > t_p + \frac{t_w}{2} \end{cases} \quad (9)$$

The location of the peak of the pulse to be extracted, t_p , is first identified. A suitable window size in the time domain, t_w , is selected such that the entirety of the pulse fits within $T(t)$, where $T(t)$ is the standard Tukey window, with a Tukey parameter equal to 0.5. The width of the window should be sufficient to contain the entire pulse, and depending on the signal being processed, different t_w values can be used on different pulses, if the pulses vary in length with respect to each other. In terms of the Tukey parameter, a value of 0 corresponds to a rectangular window (sharpest tapering), while a value of 1 corresponds to a Hann window (smoothest tapering). If the pulses being extracted are well separated, and t_w is chosen generously, then the tapering value makes little difference, as the entire signal will fall beneath the flat region of the window. For cases where pulses are closer together and hence smaller t_w values are required, then it is preferable to choose Tukey parameters closer to 0, to ensure that the tapered edges of the window do not alter the shape of the extracted pulse during the multiplication of the window with the signal. It is worth noting that while this process was completed manually in our work, future methods could automate this.

Multiplying a signal with $w(t)$ extracts the pulse of that signal whose peak is at $t = t_p$. The process is repeated, each time centring $w(t)$ around the pulse to be extracted, until the signal is decomposed into its constituent pulses. An example of this process is shown in Fig. 4; if the window shown in orange in Fig. 4(a) was to be multiplied by the signal, the first pulse would be extracted, while the same process with the window shown in Fig. 4(b) would extract the second pulse.

A limitation of the method presented in this work is interpolating signals containing multiple, overlapping pulses, which may not be separable using the method in this section. There are two cases to consider here, based on whether the derivative of the phase with respect to the interpolation variable is the same across the pulses comprising the signal. If the signals being interpolated contain multiple overlapping pulses where this derivative is common between the pulses, and this is a known fact, then interpolation can take place as usual. An example of such a case would be a signal containing multiple pulses travelling at the same speed, and interpolation is taking place in the spatial dimension. If the signals contain multiple overlapping pulses whose phase derivative with respect to the interpolation variable is not the same, or this is an unknown fact, then interpolating those would not yield meaningful results, for the reasons discussed earlier in this section.

3. FE modelling & test cases

Having developed the required theory, a number of test cases to assess the performance of the interpolation routine was generated. This section discusses the relevant FE theory and the model design & details for each test case.

3.1. FE modelling theory

The FE models here were solved using the Graphics Processing Unit (GPU) based, FE package Pogo [40]. Pogo uses standard explicit time domain elastodynamic theory and methods to solve FE models [41].

In all models, we defined some nodes to both act as source and monitor nodes; source nodes possess pre-determined forcing which subsequently excite the ultrasonic waves in the models while the displacement of the monitor nodes is stored during the solution and used in post-processing. The mesh refinement was set to be at least 25 elements per operating wavelength, calculated at the central frequency of the simulation. This level of refinement is necessary to ensure sufficient convergence of the solution [33], for NDE applications. The timestep was set to be at most 30% of the critical time step, calculated based on the aforementioned level of refinement and the speed of the wave that was to be excited in the model, to satisfy the Courant–Friedrichs–Lewy condition [42].

3.2. Test case details

To assess the performance of the interpolation routine, we generated and solved FE models for each test case, and subsequently attempted to use a small subset of measurements to generate the full set of results. Having access to the true, FE-generated set of results allows for the direct comparison of the interpolation results against the true, FE modelling generated results. Fig. 5 shows schematics of the four test cases used to assess the performance of the interpolation technique. The details of each test case are discussed below.

3.2.1. Test case I – Raster scan of a component with linearly reducing thickness

The interpolation routine presented here was validated using examples of increasing complexity. Therefore, as an initial example, we chose to simulate a component with a linearly reducing thickness, as shown in Fig. 5(a). This initial case was selected due to the phase being expected to vary linearly between measurements, in line with the assumptions made earlier.

The thinnest part of the domain was set to 75 mm, and the angle with which the thickness was set to reduce was defined to be 5°. As discussed above, the model was set up in a pulse-echo configuration, where the source nodes acted also as monitoring nodes. Each of these nodes (shown in Fig. 5(a) in red), was activated successively, from left to right, in a separate simulation, while all other source/monitor




Fig. 4. Example of a signal and the $w(t)$ required to extract the first (a) and second (b) pulse. To execute the extraction, the signal is multiplied by the respective window and the result is stored for further use.




Fig. 5. Schematics of the FE models used to evaluate the performance of our interpolation approach. Subfigure (a) shows a component with a gradually reducing thickness, (b) shows a component with an EDM notch of height h , (c) shows a component with a notch rotated by θ with respect to the vertical, and (d) shows a setup for performing a raster scan across a rectangular crack. In all schematics, the nodes acting both as source and monitor nodes are shown in red. In (a) and (c), embedded sources were used, while in (b) and (d) the sources were placed on the surface of the model; both excitation configurations were considered as they are common in FE modelling.

nodes remained inactive. We defined 31 such nodes, with a spacing of one operating wavelength, calculated at the central frequency of the simulation which was set to 2.25 MHz. A longitudinal wave was excited in each of the 31 simulations, by applying a 3-cycle Hann-windowed toneburst to the active source node, in the vertical direction. While sources comprising a single point, and hence single force, are inherently not able to excite a pure longitudinal or shear wave, the point source was displaced vertically to preferentially encourage the excitation of a longitudinal wave, while any shear waves will propagate orthogonally to it.

The material parameters were set to those of aluminium ($E = 70$ GPa, $\nu = 0.33$, $\rho = 2700$ kg/m³, where E is the Young's modulus, ν is the Poisson's ratio and ρ is the density). Finally, stiffness-reducing absorbing layers (ALs) with a thickness of at least three operating wavelengths were defined at the top, left and right regions of the domain [43,44], to absorb any unwanted reflections.

3.2.2. Test case II – Reflection from Electrical Discharge Machining (EDM) notch

So far, an example where the interpolation variable is the location of the transducer have been presented; however, as discussed in Section 2, we are free to interpolate across other dimensions of a parameter space,

adjusting the weights according to the method discussed in Section 2.2. For this example, a domain containing an EDM notch was simulated, as shown in Fig. 5(b). The height of the notch, h , was varied from 10 mm to 20 mm, in 0.5 mm intervals while the width of the notch was also kept constant at 0.25 mm. To more closely resemble the geometry of EDM notches, a semicircular tip with a radius of 0.125 mm was defined, instead of a flat one. A separate mesh was generated for every notch height to be investigated.

The position of the transducer was kept constant, directly above the notch. In the previous example single nodes were used for the excitation — here, a more realistic excitation and reception configuration was defined. To simulate a realistic transducer, a source line with a length of 25.4 mm (equivalent to a 1-inch transducer) was defined on the top edge of the domain. All nodes in the source line were excited simultaneously, using a 3-cycle Hann-windowed toneburst with a central frequency of 5 MHz. The same nodes were also set to act as monitor nodes, similarly to a transducer in a pulse-echo configuration. Gaussian amplitude apodisation was applied across the source, to suppress the generation of Rayleigh and tip-diffracted waves. The far-field radiation pattern from a source is related to the Fourier transform of its apodisation function [45]; the Fourier transform of a uniform window is the sinc function, which possesses strong side lobes, while the Fourier transform of a Gaussian window is also a Gaussian. Therefore, when

the former is used, energy radiates in unwanted direction, while the latter ensures better energy concentration in the desired direction. Also, when uniform windows are used, the discontinuity in displacement between the two edges of the source line and the rest of the domain acts like a sharp tip, generating tip diffracted waves. This effect diminishes with Gaussian apodisation, as the amplitude of the source gradually diminishes to zero at its edges. The gradual reduction in amplitude towards the edges also suppresses the generation of Rayleigh waves, which can be excited via abrupt localised disturbances on the surface of a material. The material parameters were set to those of steel ($E = 216.9$ GPa, $\nu = 0.287$, $\rho = 7930$ kg/m³), and ALs with a thickness of at least three operating wavelengths were defined at the left and right regions of the domain.

3.2.3. Test case III – Reflections from a notch at an angle

Varying the tilt of a defect with respect to the propagation direction of the wave and recording the response for post-processing, is common practice during the qualification of an NDE method, especially in the nuclear and aerospace industries [20,39,46]. To recreate this scenario, an FE domain containing a flat notch, tilted by θ with respect to the horizontal, was generated, as shown in Fig. 5(c). The angle of the notch was varied from 0° to 10°, in 1° increments, and a model was solved per notch angle.

A sole node, acting both as a source and a monitor node was defined near the top right edge of the model. It was excited using a 3-cycle Hann windowed toneburst with a central frequency of 2.25 MHz. The material properties of the domain were set to those of aluminium ($E = 70$ GPa, $\nu = 0.33$, $\rho = 2700$ kg/m³) and ALs were defined at the left, right and top boundaries of the domain.

3.2.4. Test case IV – Raster scan of a horizontal crack

We considered a final common case explored during the qualification of NDE methods — that of a raster scan over a crack. In this case, the amplitude of the reflection is recorded and stored, to generate plots showing the variation of the amplitude of the reflection from the defect as a function of the lateral position of the transducer [47–49]; those plots are then used to both assess the performance of a transducer in an inspection scenario, and also, the limits of the crack detection from the perspective of the ability to observe the reflection.

To investigate this case, a rectangular domain with a horizontal crack at the centre of the domain was generated. The length and width of the crack were set to 3 mm and 0.5 mm respectively. A schematic of this setup is shown in Fig. 5(d). A source line with a length of 30 mm was defined on the top boundary of the domain. However, to simulate a raster scan, not all nodes were activated simultaneously. We chose to simulate a 3 mm transducer, and take measurements at 0.5 mm intervals, scanning from left to right. For the first measurement, only the first m nodes of the 30 mm source line were activated, where m is the number of nodes corresponding to a 3 mm source. Then, for the second scan, nodes $m/2$ to $3m/2$ were active, and so forth, up until the end of the source line was reached. This process of grouping the nodes is shown schematically by the black ellipses in Fig. 5(d) – in each simulation, only the nodes in one ellipse were active, and Gaussian apodisation was applied to them to more closely replicate the physics of a real transducer. The material parameters and AL setup were kept the same as the previous example, with the sole difference being the absence of an AL on the top region of the material; to simulate a more realistic scan, as the excitation was taking place at that region.

4. Results & discussion

The solution of the four test cases discussed above yielded a full set of FE-generated results for each test case. We then attempted to use a subset of those to obtain the full set of results via interpolation — details of the exact choice of nodes with known and unknown measurements will be given when discussing each test case. In terms of

assessing the quality of the interpolation-assisted results with respect to the true FE generated solution, the error in the peak amplitude (AE_i) of the signal, and the error in the location (LE_i) of this peak, between the i th measurement and the i th prediction, were calculated as follows:

$$AE_i := \frac{\max(|u_i^{\text{int}}(t)|) - \max(|u_i^{\text{FE}}(t)|)}{\max(|u_i^{\text{FE}}(t)|)}, \quad (10)$$

and

$$LE_i := \frac{t_i^{\text{int}} - t_i^{\text{FE}}}{t_i^{\text{FE}}}, \quad (11)$$

where $\max(|u_i^{\text{int}}(t)|)$ and $\max(|u_i^{\text{FE}}(t)|)$ denote the absolute signal maxima of the i th measurement predicted by interpolation and obtained by the FE model respectively, and t_i^{int} and t_i^{FE} are the times of arrival of those maxima.

We also recognise that it is not sufficient to assess whether the interpolated pulses arrive at the correct time and with the correct amplitude; for most applications, an overall match in the shape of the signal between the FE and interpolation generated signals will be required. Additionally, it may not be sufficient to compare the estimated signal with its corresponding, FE generated counterpart; erroneous interpolation may result into the estimated signal resembling another signal in the parameter space more closely. To assess both concerns, a cross-correlation based metric, r_{ij} was defined, taking into account the full time traces of both the estimated signal at the i th measurement in the parameter space $u_i^{\text{int}}(t)$, and comparing it with the true, FE generated signal of the j th measurement $u_j^{\text{FE}}(t)$:

$$r_{ij} := 1 - \frac{\int_{-\infty}^{\infty} u_i^{\text{int}}(t) \cdot u_j^{\text{FE}}(t) dt}{\sqrt{\int_{-\infty}^{\infty} |u_i^{\text{int}}(t)|^2 dt \int_{-\infty}^{\infty} |u_j^{\text{FE}}(t)|^2 dt}}. \quad (12)$$

This metric is based on the cross correlation of the FE and interpolated signals, at zero lag. This means that r_{ij} would have a value of 0 for a perfect match (i.e. $u_i^{\text{int}}(t) = u_j^{\text{FE}}(t)$), and a value of 1 for no correlation between $u_i^{\text{int}}(t)$ and $u_j^{\text{FE}}(t)$. Values between 1 and 2 are also possible, with a value of 2 indicating that $u_i^{\text{int}}(t)$ and $u_j^{\text{FE}}(t)$ have the same shape but are out of phase by π . For the correct interpolation results, r_{ij} should be minimum (ideally 0) when $i = j$.

In Eqs. (10)–(12) the meaning of the term “ i th” measurement changes slightly based on the example case, but always denotes the signal obtained when the variable of the dimension of the parameter space being investigated is set to its “ i th” value, whether that be measurement location, characteristic defect size, crack angle etc.

When presenting the results in this section, we will quote AE , LE and r values respectively, where $AE = \langle AE_i \rangle$, $LE = \langle LE_i \rangle$ and $r = \langle r_{ii} \rangle$ with $\langle \cdot \rangle$ denoting the ensemble average across i .

Sections 4.1–4.4 present and discuss the results for each of the test cases presented in the previous Section, and Section 4.5 presents the efficiency gains attainable when interpolation is deployed.

4.1. Raster scan of a component with linearly reducing thickness

As mentioned in Section 3.2.1, we had defined 31 nodes in the model acting both as source and monitor nodes. Having solved a separate model per such node, we were able to plot a B-scan comprising all the results obtained from the 31 models. A B-scan is an image generated by collating the signals produced at many measurement locations and plotting them against their respective measurement location. This is shown in Fig. 6(a). As expected the time of arrival of the reflection from the bottom surface decreases linearly from the first (leftmost — node 1) to the last (rightmost — node 31) node.

Given the simplicity of this model, for the interpolation it was sufficient to only use the signals from the leftmost and rightmost nodes to predict the displacement at all other nodes. The B-scan resulting from this process is shown in Fig. 6(b). Visually, there is a very good match




Fig. 6. True FE generated (a) and interpolation-assisted (b) B-scans from solving the FE model shown schematically in Fig. 5(a). A comparison of a true, FE generated and interpolation generated signal in the time (c) and frequency (d) domain, at an arbitrary node is also shown. The signals used in the interpolation process are plotted in (e), while the r_{ij} plot for all i, j combinations is shown in (f).

between the FE generated and interpolation-assisted B-scans. The mean absolute AE and LE were found to be 0.123% and 0.002% respectively, and the r value was 1.5658×10^{-5} .

There may exist cases where there is interest in quantities besides the amplitude of a signal and its time of arrival. Therefore, it is important to establish that the shape of the interpolated signals is preserved at all points in time. A comparison between a time trace at an arbitrary node and the corresponding interpolation result, as well as their frequency domain depictions are shown in Fig. 6(c) and (d) respectively. In both cases there is a very good match between the true FE generated and interpolation results. This is also evident by the small r value quoted above.

Fig. 6(e) shows the true FE signals which when interpolated, generated all the signals plotted in Fig. 6(b) in the form of a B-scan. The signals are clearly separated by more than half a wavelength, which is a traditional limit for simple signal interpolation, showing the ability of our method to interpolate across larger gaps in a parameter space. For clarity and conciseness, we have not included the signals used as inputs for interpolation on the other plots in this paper.

Finally, Fig. 6(f) shows how well each interpolated signals matches with any of the signals in the dataset. It is clear that r_{ij} is minimised when $i = j$, demonstrating that every signal generated via interpolation matches the corresponding FE generated signal most closely. This also means that quoting the r value, defined previously as $\langle r_{ii} \rangle$, is meaningful.

4.2. Reflection from an EDM notch

Unlike the previous example, here the variable was set to be the height of the EDM notch, and not the transducer position. Ten models were solved, each containing an EDM notch whose height varied from 10 mm to 20 mm. For each case, the time of arrival of the reflection, as well as the peak amplitude of the reflection were recorded and stored. The results from this example also relate well with part of the motivation of this study, which was the rapid generation of data for NDE qualification, as the reflection amplitude from defects is a typical quantity of interest during the qualification process [21,47].

Then, using the signals solely from the models containing the smallest and largest notches, we interpolated to obtain the signals corresponding to all intermediate heights. Those signals were processed to extract their arrival time and peak amplitude — the results are shown in Fig. 7. In the same figure, an example of an interpolation-generated A-scan versus the true FE signal is also plotted.

Fig. 7(b) and (c) clearly show a good agreement between the true FE and interpolation-assisted results. The mean absolute AE and LE were found to be 0.241% and 0.000% (NB given to standard precision used throughout the paper) respectively, while the r value was 3.5617×10^{-5} . The r_{ij} map in Fig. 7(d) clearly demonstrates that there are no mismatches between the signals generated via interpolation, and their true FE counterparts.




Fig. 7. Comparison between an FE and an interpolation-generated signal (a), obtained by solving the model shown schematically in Fig. 5(b). The time of arrival of the signals generated via interpolation, as well as their amplitude were extracted for each simulation solved, and Figures (b) and (c) show the reflection amplitude and time of arrival as a function on the height of the notch respectively. The r_{ij} plot for all i, j combinations is shown in (d).

At this point it is important to highlight that we are not directly interpolating the datapoints in Fig. 7(b) and (c) – the interpolation-assisted results are extracted directly from A-scans, similar to the one shown in Fig. 7(a). Therefore, the plots here are solely an example of potentially useful diagrams; the signals themselves remain available for any other post-processing deemed suitable.

4.3. Reflections from a notch at an angle

Similarly to the example above, here we have again not interpolated across transducer locations — in this example, the response from a surface breaking notch at an angle was recorded, as this angle varies. Eleven models were solved, varying the angle of the notch from 0° to 10° , in 1° increments.

The signals obtained from these models possess a greater degree of complexity. To demonstrate this, Fig. 8(a) shows the signals obtained from the solution of models with a notch angle between 0° and 5° . Three pulses exist in all signals; the leftmost pulse is a reflection from the bottom surface of the domain, the middle pulse arises from corner trapping, and the rightmost pulse is a reflection directly from the right face of the notch. The pulses also exhibit different properties amongst themselves, with the first pulse varying neither in amplitude nor in time of arrival, the second one varying subtly in amplitude but significantly in time of arrival, and the third one having small variation in the time of arrival but noticeable variation in amplitude. Therefore, the windowing routine described in Section 2.3 was deployed, to decompose the signals to be interpolated into their constituent pulses. Each pair of pulses was interpolated separately, and the three resulting signals were summed to give the final result for each interpolation point. The decomposed pulses were interpolated using the same method deployed throughout this work, i.e. by assuming linear variation in amplitude and phase.

We are not restricted to selecting the results at two extremes of the dimension of the parameter which is to be populated via interpolation; we are free to select any suitable subset. To demonstrate this, we chose to use the simulation results at 0° and 5° to generate the corresponding results at 1° , 2° , 3° and 4° , and then to use the results from 6° and 10° , to generate the results at 7° , 8° and 9° . An example of the interpolation

result, and the corresponding true FE signal, from one case from each of these subdomains is shown in Fig. 8.

A good match between the interpolation-assisted and true FE results was found in all cases examined. However, using the interpolation generated results has reduced the required number of simulations by 60%. This is assuming that solely those 11 results were needed; given that the results for those are satisfactory, results for very finely spaced intermediate angles can be generated, with minimal additional computational time and resources required in comparison with obtaining those results from full FE modelling.

While the overall shape of the signal is predicted well, the most noticeable discrepancies between the signals generated via interpolation and the true FE signals are observed in the third pulse, as shown in Fig. 8(c) and (d). The underlying assumption of linear interpolation throughout this work is the linear variation of the amplitude and the phase of the signals in the particular section of the parameter space being investigated. Discrepancies between the interpolation-assisted and true FE solutions arise when these assumptions are not satisfied, which may be what we are observing in the bottom panels of Fig. 8(b) and (c). In terms of quantifying how well the signals match, the r value for this example was 3.6911×10^{-4} . Again, the r_{ij} map in Fig. 8(f) clearly shows that there are no mismatches between the signals generated via interpolation, and their true FE counterparts.

4.4. Raster scan of a horizontal crack

In this last case, a raster scan of a transducer over a vertical crack was simulated. For the post processing we chose to extract the amplitude of the reflection from the crack and plot this as a function of the location of the measurement; as mentioned before, this type of plot is commonly used in qualification of NDE methods [47–49]. A total of 60 models were solved, which is equivalent to a 3 mm transducer scanning over a distance of 30 mm, with a 0.5 mm overlap between measurements.

The measurement locations are numbered from 1 to 60, from left to right. The results from the odd numbered measurement locations were used to predict the analogous measurement at the even numbered locations. The resulting reflection amplitude versus transducer lateral




Fig. 8. Example of signals obtained from solving the domain shown schematically in Fig. 5(c) for different θ values (a), and comparison between the signal obtained via interpolation at 3° (b), and 7° (c). To aid visualisation, subfigures (d) and (e) show the portion of the signals in subfigures (b) and (c) corresponding to the smaller pulses respectively. The r_{ij} plot for all i, j combinations is shown in (f).

position plot is shown in Fig. 9(a), alongside with two examples of the signals used to construct this plot in Fig. 9(b) and (c). A comparison with the corresponding true FE results is plotted in the same figures.

Both the reflection amplitude curve and the signal from the interpolation-assisted results match well with the results from the true FE model, but with 48% fewer simulation results. The mean absolute AE in Fig. 9(a) is 0.472%, while the r value was 8.5107×10^{-4} .

There is an interesting feature in Fig. 9(d). The locus of the points where r_{ij} is minimum for every i is no longer a straight line occurring when $i = j$ – there is mirror symmetry around the centroid of the plot. This however, is expected, given the symmetry of the model, in which the measurement locations are symmetric around the crack, resulting in the signal at each location matching that of its corresponding location equidistant from the opposite end (e.g., the signal from measurement 1 matches the signal from measurement 60, the signal from measurement 2 matches the signal from measurement 59, etc.).

The method discussed here also copes well with asymmetrical signals/signals with more complicated shapes, as shown in Fig. 9(b), where the signal does not possess the classical shape of a Hann-windowed toneburst. Since this signal has been received away from the crack, it is likely that this disturbance in shape is due to the interaction between the incident wave and the tip of the crack closest to it, while the signal in Fig. 9(c) is a direct reflection from the face of the crack.

We can further reduce the number of results required to obtain the reflection amplitude versus transducer location curve. Suppose

the node numbering is maintained, and the results from locations 1 and 4 are used to predict the results at locations 2 and 3, then the results from locations 4 and 7 are used to predict the results at location 5 and 6 and so forth. This reduces the required number of simulations by approximately 63%. The resulting reflection amplitude versus transducer lateral position plot is shown in Fig. 10(a), alongside two examples of the signals used to construct this plot in Fig. 10(b) and (c). A comparison with the corresponding true FE results is plotted in the same figures.

The results in this case are still satisfactory – the signals match well, while the mean absolute AE for the reflection amplitude versus transducer position curve is 0.981% while the r value was 7.970×10^{-2} . As expected, the reconstruction quality and associated errors have increased, however fewer simulations were required to obtain the results. This phenomenon is expected to hold true in any interpolation example and the acceptable level depends on the specific problem being solved. There are no signal mismatches between the signals generated by interpolation and the true FE signals, as shown in Fig. 10(d).

4.5. Efficiency gains

We would like to highlight the significant increase in efficiency which is achieved when interpolation is used to populate a parameter space. In the second column of Table 1 we are reporting the FE modelling solution time to obtain the full set of results for each of




Fig. 9. Comparison between a reflection amplitude vs. transducer position curve obtained via processing of full FE and interpolation-assisted results (a). The FE results were obtained by solving the model shown schematically in Fig. 5(d). Examples of the signals used for the construction of the curves in (a) from measurement locations away (b) and just above (c) from the crack are also shown. The r_{ij} plot for all i, j combinations is shown in (d).




Fig. 10. Comparison between a reflection amplitude vs. transducer position curve obtained via processing of full FE and interpolation-assisted results (a). The FE results were obtained by solving the model shown schematically in Fig. 5(d). Examples of the signals used for the construction of these curved from measurement locations away (b) and just above (c) from the crack are also shown. Compared with Fig. 9, fewer nodes with known measurements were used to generate the results here. The r_{ij} plot for all i, j combinations is shown in (d).

the cases examined in this work. The third column shows the time that linear interpolation required to generate the interpolation-assisted results.

As shown in Table 1, interpolation results in a significant increase in efficiency. The full FE solution time here is not prohibitive for any of the four cases, however, FE solution time can increase significantly, for example, when taking finer steps in the parameter space or working with larger or three-dimensional spaces. Increasing the number of generated signals has little effect in the interpolation time, as seen when comparing the times reported in the second column for the two variations of Case IV. For the latter, since the interpolation time is

largely a function of the length of the signal, we anticipate no additional interpolation time if the signals originate from three-dimensional FE modelling.

It is worth noting that hardware influences the times reported in Table 1. For the FE modelling, the times reported are a result of solving the models using one Nvidia RTX 2080Ti graphics card with 11 GB of memory. A standard desktop computer with an Intel Core i7-7700K CPU processor and 32 GB of RAM was used to generate the interpolation data. Hardware differences are expected to affect these times, however it is clear the interpolation reduces the data generation times by many orders of magnitude.

Table 1

Comparison between the time taken to populate the parameter space in each of Cases I–IV using FE modelling and interpolation. The interpolation time is the time required to interpolate the signals, and populate the equivalent parameter space.

Case number	Full FE solution time (s)	Interpolation time (s)
Case I	93	0.089
Case II	10	0.021
Case III	275	0.088
Case IV (Fig. 9)	30	0.034
Case IV (Fig. 10)	30	0.027

5. Conclusion

In-silico methods for the qualification of NDE techniques are certainly desirable, however, data availability is an issue, as the time required to generate sufficient data to qualify a method in-silico with the required level of confidence might be prohibitive. Here we have presented a method which uses a few FE generated results in combination with an interpolation approach to rapidly generate more data without the need to solve the full FE model for every parameter step required. This method takes into account the inherent property of a signal to possess both an amplitude and a phase, both of which must be carefully preserved during interpolation.

We have derived the necessary equations to interpolate across different dimensions in a parameter space. The focus of this paper has been interpolating across measurement locations; this is applicable to NDE scenarios where raster scans, or multiple measurements on different locations are required. We have also demonstrated the applicability of our method to other scenarios, where the variable of the interpolation was the height or the tilt of the defect.

In all cases interpolation results match full FE results very well, however, they have been generated in a fraction of the time it would have taken to solve full FE models. It is worth re-iterating that this method does not interpolate the post-processed results but the signals directly, meaning that the interpolation results can be treated and used in the same fashion as true FE modelling generated results.

We have also demonstrated the ability of the method discussed in this paper to yield good results from complicated signals, containing multiple pulses not necessarily travelling at the same speed, and for signals containing pulses with elaborate shapes other than the typical, Hann-windowed pulses. The limitations of the method were also demonstrated, by examining a case where discrepancies between the signals generated via interpolation and FE modelling arose, potentially due to the assumption of linear variation in the amplitude and the phase of the signals not being satisfied.

It should be recognised that no interpolation method will be able to interpolate arbitrary patterns to infinite accuracy; here we aim to capture typical patterns that appear in wavefields. It is unlikely that abrupt changes will exist as the underlying parameters vary; the presence of diffraction will inevitably smooth the data.

This accurate and rapid generation of data allows for the exploration of broader or higher order parameter spaces in a more feasible fashion with applications in the NDE qualification field, as well as machine learning or any other area where ultrasonic data availability is an issue.

CRediT authorship contribution statement

Georgios Sarris: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation. **Michael J.S. Lowe:** Writing – review & editing, Supervision, Funding acquisition, Conceptualization. **Peter Huthwaite:** Writing – review & editing, Supervision, Software, Resources, Project administration, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Georgios Sarris, Peter Huthwaite, Michael Lowe report financial support was provided by Engineering and Physical Sciences Research Council. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors are funded by UK Research and Innovation – Project reference EP/X02427X/1.

Data availability

Data will be made available on request.

References

- [1] N. McLean, J.P. McKenna, A. Gachagan, J. McArthur, G. Hayward, The design of a multi-agent NDE inspection qualification system, AIP Conf. Proc. (2007).
- [2] K. Markulin, M. Vavrou, M. Kemppainen, I.V. R., Paussu, J. Pirinen, Effect of the different artificial flaw type in qualification purposes from the point of view of the inspection company, in: 10th International Conference on Non Destructive Evaluation in Relation To Structural Integrity for Nuclear and Pressurized Components, 2013, pp. 409–416.
- [3] H. Wüstenberg, A. Erhard, Problems with artificial test reflectors at the performance demonstration of ultrasonic inspections, in: Proceedings of 6th European Conference on Non Destructive Testing, 1994, pp. 741–746.
- [4] M. Kemppainen, I. Virkkunen, J. Pitkänen, R. Paussu, H. Hänninen, Advanced flaw production method for in-service inspection qualification mock-ups, Nucl. Eng. Des. 224 (1) (2003) 105–117.
- [5] R. Singh, Three decades of NDT reliability assessment and demonstration, Mater. Eval. 59 (7) (2001) 856–860.
- [6] H.M. Stephens, Performance based NDE personnel qualifications, in: Proceedings of the ASME 2014 Pressure Vessels & Piping Conference PVP2014, 2014, pp. 1–6.
- [7] J.L. Warren, J.E. Yoo, C.A. Meyer, D.S. Molony, H. Samady, H.N. Hayenga, Automated finite element approach to generate anatomical patient-specific biomechanical models of atherosclerotic arteries from virtual histology-intravascular ultrasound, Front. Med. Technol. 4 (2022).
- [8] D. Mills, S. Smith, Finite element comparison of single crystal vs. multi-layer composite arrays for medical ultrasound, IEEE Trans. Ultrason. Ferroelectr. Freq. Control 49 (7) (2002) 1015–1020.
- [9] N. Fougeron, P.-Y. Rohan, D. Haering, J.-L. Rose, X. Bonnet, H. Pillet, Combining freehand ultrasound-based indentation and inverse finite element modeling for the identification of hyperelastic material properties of thigh soft tissues, J. Biomech. Eng. 142 (9) (2020).
- [10] A.R. Westervelt, M. Fernández, M.G. House, J. Vink, C. Ling Nhan-Chang, R.J. Wapner, K.M. Myers, A parameterized ultrasound-based finite element analysis of the mechanical environment of pregnancy, J. Biomech. Eng. 139 (5) (2017).
- [11] J. Peikert, A. Hampel, M. Bagge, Relative importance of poroelastic effects and viscoelastic relaxation for postseismic velocity fields after normal and thrust earthquakes: Insights from 2D finite-element modelling, Tectonophysics 838 (2022) 229477.
- [12] E. Sotelo, M. Favino, R.L. Gibson, Application of the generalized finite-element method to the acoustic wave simulation in exploration seismology, Geophysics 86 (1) (2021) T61–T74.
- [13] M. Galis, P. Moczo, J. Kristek, A 3-D hybrid finite-difference-finite-element viscoelastic modelling of seismic wave motion, Geophys. J. Int. 175 (1) (2008) 153–184.
- [14] H. Zhou, X. Song, W. Xiong, H. Wu, Q. Li, C. Cai, Vibration and underwater noise simulation of a river-crossing tunnel based on a 2.5D structural-acoustic finite element method, Structures 64 (2024) 106502.
- [15] Z. He, G. Liu, Z. Zhong, X. Cui, G. Zhang, A. Cheng, A coupled edge-/face-based smoothed finite element method for structural acoustic problems, Appl. Acoust. 71 (10) (2010) 955–964.
- [16] E. Li, Z.C. He, Y. Jiang, B. Li, 3D mass-redistributed finite element method in structural-acoustic interaction problems, Acta Mech. 227 (3) (2015) 857–879.
- [17] M. Miniaci, A. Krushynska, F. Bosia, N.M. Pugno, Large scale mechanical metamaterials as seismic shields, New J. Phys. 18 (8) (2016) 083041.

- [18] J.M. De Ponti, A. Colombi, R. Ardito, F. Braghin, A. Corigliano, R.V. Craster, Graded elastic metasurface for enhanced energy harvesting, *New J. Phys.* 22 (1) (2020) 013013.
- [19] S. Uhlig, I. Alkhasli, F. Schubert, C. Tschöpe, M. Wolff, A review of synthetic and augmented training data for machine learning in ultrasonic non-destructive evaluation, *Ultrasonics* 134 (2023) 107041.
- [20] S.G. Haslinger, M.J.S. Lowe, R.V. Craster, P. Huthwaite, F. Shi, Prediction of reflection amplitudes for ultrasonic inspection of rough planar defects, *Insight - Non-Destr. Test. Cond. Monit.* 63 (1) (2021) 28–36.
- [21] S.G. Haslinger, M.J. Lowe, P. Huthwaite, R.V. Craster, F. Shi, Elastic shear wave scattering by randomly rough surfaces, *J. Mech. Phys. Solids* 137 (2020) 103852.
- [22] Z. Wei, F. Shi, Z. Wang, Modeling stochastic elastic wave diffraction by the tips of randomly rough defects, *J. Mech. Phys. Solids* 190 (2024) 105744.
- [23] E. Georgiades, M.J. Lowe, R.V. Craster, Leaky wave characterisation using spectral methods, *J. Acoust. Soc. Am.* 152 (3) (2022) 1487–1497.
- [24] D.A. Kiefer, M. Ponschab, S.J. Rupitsch, M. Mayle, Calculating the full leaky lamb wave spectrum with exact fluid interaction, *J. Acoust. Soc. Am.* 145 (6) (2019) 3341–3350.
- [25] M. Ivanyina, R. Qu, R. Hu, C. Zhong, J. Chan-Lau, Surrogate data models: Interpreting large-scale machine learning crisis prediction models, *IMF Work. Pap.* 2023 (041) (2023) 1.
- [26] J.T. Hwang, J.R. Martins, A fast-prediction surrogate model for large datasets, *Aerosp. Sci. Technol.* 75 (2018) 74–87.
- [27] J.A. Cain, S. Visagan, M.M. Monti, S.M.A.R.T. F.U.S: Surrogate model of attenuation and refraction in transcranial focused ultrasound, *PLoS One* 17 (10) (2022) e0264101.
- [28] C. Shorten, T.M. Khoshgoftaar, A survey on image data augmentation for deep learning, *J. Big Data* 6 (1) (2019).
- [29] I. Virkkunen, T. Koskinen, O. Jessen-Juhler, J. Rinta-aho, Augmented ultrasonic data for machine learning, *J. Nondestruct. Eval.* 40 (1) (2021).
- [30] A.K. Supreet, J.B. Harley, Transfer learning of ultrasonic guided waves using autoencoders: A preliminary study, in: 45th Annual Review of Progress in Quantitative Nondestructive Evaluation, 2019, 050013–1–050013–8.
- [31] T. Di Ianni, L. De Morelli, A. Perelli, A. Marzani, Compressive sensing of full wave field data for structural health monitoring applications, *IEEE Trans. Ultrason. Ferroelectr. Freq. Control.* 62 (7) (2015) 1373–1383.
- [32] P. Paialunga, F. Shi, S.G. Haslinger, J. Corcoran, A machine learning surrogate model for time of flight diffraction measurements of rough defects, *NDT E Int.* 144 (2024) 103089.
- [33] M.B. Drozd, Efficient Finite Element Modelling of Ultrasound Waves in Elastic Media (Ph.D. thesis), Imperial College London, 2008.
- [34] G. Sarris, S.G. Haslinger, P. Huthwaite, P.B. Nagy, M.J.S. Lowe, Attenuation of Rayleigh waves due to three-dimensional surface roughness: A comprehensive numerical evaluation, *J. Acoust. Soc. Am.* 154 (2) (2023) 808–818.
- [35] P. Xu, R. Jones, G. Sarris, P. Huthwaite, Efficient generation of realistic guided wave signals for reliability estimation, *Struct. Heal. Monit.* (2025) 1–23.
- [36] J. Zhang, B.W. Drinkwater, P.D. Wilcox, Defect characterization using an ultrasonic array to measure the scattering coefficient matrix, *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* 55 (10) (2008) 2254–2265.
- [37] Z. Andreas, P. Huthwaite, B. Pavlakovic, High-resolution thickness maps of corrosion using SH1 guided wave tomography, *Proc. R. Soc. A: Math. Phys. Eng. Sci.* 477 (2245) (2021).
- [38] A.V. Pamel, G. Sha, S.I. Rokhlin, Michael, Finite-element modelling of elastic wave propagation and scattering within heterogeneous media, *Proc. R. Soc. A: Math. Phys. Eng. Sci.* 473 (2197) (2017) 20160738.
- [39] C. Curtis, A. Walker, Inspection qualification and NDE reliability, in: *Proceedings of 58th Annual Conference of the British Institute of Non-Destructive Testing-NDT*, 2019, p. 1.
- [40] P. Huthwaite, Accelerated finite element elastodynamic simulations using the GPU, *J. Comput. Phys.* 257 (2014) 687–707.
- [41] K.J. Bathe, *Finite Element Procedures*, Prentice-Hall, 1996.
- [42] R. Courant, K. Friedrichs, H. Lewy, On the partial difference equations of mathematical physics, *IBM J. Res. Dev.* 11 (2) (1967) 215–234.
- [43] J. Pettit, A. Walker, P. Cawley, M. Lowe, A stiffness reduction method for efficient absorption of waves at boundaries for use in commercial finite element codes, *Ultrasonics* 54 (7) (2014) 1868–1879.
- [44] P. Rajagopal, M. Drozd, E.A. Skelton, M.J.S. Lowe, R.V. Craster, On the use of absorbing layers to simulate the propagation of elastic waves in unbounded isotropic media using commercially available finite element packages, *NDT E Int.* 51 (2012) 30–40.
- [45] K.J. Parker, Apodization and windowing functions, *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* 60 (6) (2013) 1263–1271.
- [46] A. Ballisat, P. Wilcox, R. Smith, D. Hallam, Model assisted qualification of NDE techniques, *AIP Conf. Proc.* (2017).
- [47] T.A. Gray, R.B. Thompson, B.P. Newberry, Model-based ultrasonic NDE system qualification methodology, *Rev. Prog. Quant. Nondestruct. Eval.* (1987) 957–965.
- [48] T.P. Lerch, L.W. Schmerr, A. Sedov, Characterization of spherically focused transducers using an ultrasonic measurement model approach, *Res. Nondestruct. Eval.* 8 (1) (1996) 1–21.
- [49] L.W. Schmerr, *Fundamentals of Ultrasonic Nondestructive Evaluation : A Modeling Approach*, Springer International Publishing, Cham, 2016.