

Development of an Automated Ultrasonic Scanning Gantry for Investigating Metal Microstructure

A. Alyubi, O. Bruce-Gardyne, O. Djuric, H. Sabri, T. Tabet, R. Zabedee

April 2026

Supervised by:

Dr. Alexander Velichko

University of Bristol

School of Electrical, Electronic and Mechanical Engineering

2025-2026

Summary

Modern manufacturing has a growing need to characterise metal microstructure, and thanks to its success in non-destructive evaluation, ultrasonic testing stands out as a faster and cheaper alternative to conventional techniques. Investigating sub-millimetre features using ultrasound requires accurate and repeatable data collection methods, along with computationally efficient algorithms to reconstruct the internal structures. Large-area inspection requires precise probe positioning and the ability to combine successive scans into continuous volumetric images.

A 3-axis gantry was designed to position 1D or 2D ultrasonic arrays over a sample surface, guided by automated control to capture structured scan sequences. The raw Full Matrix Capture data collected was filtered, cropped and compressed before using the Total Focusing Method to reconstruct the volumetric images. This revealed the internal backscattering diffraction patterns in the metal samples.

Synthetic and experimental studies showed that microstructural backscatter could be stitched into larger images, with statistical validation indicating that stitching reliability depends on overlap and scan quality. A final end-to-end experiment demonstrated successful volumetric imaging of a pure copper sample, although mechanical limitations and incomplete rotational performance mean the current system remains a prototype. Overall, the project establishes a strong foundation for automated ultrasonic imaging and identifies the refinements needed for a more robust and fully validated platform.

Declaration

This project report is submitted towards an application for a degree in Mechanical Engineering at the University of Bristol. The report is based upon independent work by the candidates. All contributions from others have been acknowledged and the supervisor is identified on the front page. The views expressed within the report are those of the authors and not of the University of Bristol.

We hereby assert our right to be identified as the authors of this report. We give permission to the University of Bristol Library to add this report to its stock and to make it available for consultation in the library, and for inter-library lending for use in another library. It may be copied in full or in part for any bona fide library or research worker on the understanding that the users are made aware of their obligations under copyright legislation.

We hereby declare that the above statements are true.

Name	Signature
A. Alyubi	
O. Bruce-Gardyne	
O. Djuric	
H. Sabri	
T. Tabet	
R. Zabedee	
Date:	27/04/2026

©Copyright, {A. Alyubi, O. Bruce-Gardyne, O. Djuric, H. Sabri, T. Tabet, R. Zabedee, 2026}

Certification of ownership of the copyright in a dissertation presented as part of and in accordance with the requirements for a degree in Mechanical Engineering at the University of Bristol. This report is the property of the University of Bristol Library and may only be used with due regard to the authors. Bibliographical references may be noted but no part may be copied for use of quotation in any published work without prior permission of the authors. In addition, due acknowledgement for any use must be made.

Work Allocation

No.	Section / Task	AA	OB	OD	HS	TT	RZ
1	Introduction	-	-	100%	-	-	-
2	System Overview and Specification	-	-	100%	-	-	-
3	Concept Generation and Design Specification	90%	-	-	5%	-	5%
4	Design, Iterations, and Manufacturing	45%	-	-	55%	-	-
5	Automated Data Collection	-	-	-	-	-	100%
6	Gantry Performance Evaluation	33%	-	-	34%	-	33%
7	Preprocessing and Imaging	-	5%	83%	-	12%	-
8	Volume Stitching	-	100%	-	-	-	-
9	Synthetic Data Generation	-	-	-	-	100%	-
10	Stitching Validation	-	100%	-	-	-	-
11	Final Experimental Procedures	-	-	100%	-	-	-
12	Conclusion	16%	16%	16%	16%	16%	16%

We confirm that the information on this page accurately describes our individual contributions to the project.

Name

Signature

A. Alyubi



O. Bruce-Gardyne



O. Djuric



H. Sabri



T. Tabet



R. Zabedee



Date:

27/04/2026

Contents

1	Introduction (OD)	5
2	System Overview and Specification (OD)	6
3	Concept Generation of Scanning Gantry (AA)	8
4	Design, Iterations, and Manufacturing (AA, HS)	10
4.1	Base Frame (HS)	10
4.2	X-axis Slider (HS)	12
4.3	Y-axis Slider (AA)	12
4.4	Z-axis and Rotational Mechanism (AA)	13
4.5	Array Clamp (HS)	14
4.6	Summary of Axes and Motion Parameters (HS)	14
4.7	Design for Manufacture (AA, HS)	15
5	Automated Data Collection(RZ)	17
5.1	Electronics (RZ)	17
5.2	Gantry Automation (RZ)	18
5.3	Calibration (RZ)	20
6	Gantry Performance Evaluation (AA, HS, RZ)	22
7	Preprocessing and Imaging (OB, OD, TT)	23
7.1	Filtering (OD)	23
7.2	Speed of Sound (OD)	25
7.3	Total Focusing Method (OD)	26
7.4	Hardware Acceleration (OD)	28
7.5	2D Array Processing (OB, OD)	30
7.6	Radon Transform (TT)	31
8	Volume Stitching (OB)	32
8.1	Stitching Implementation	33
8.2	Algorithm Pipeline	35
8.3	Stitching in Two Dimensions	35
9	Synthetic Data Generation (TT)	36
9.1	Motivation	36
9.2	Initial Method	37
9.3	Born Forward Model	37
9.4	Validation of Synthetic Data	41
10	Stitching Validation (OB)	43
10.1	Bootstrap Tests	43
10.2	Certainty Metrics	44
10.3	Angular Misalignment Tests	47
11	Final Experimental Procedure and Results (OD)	48
12	Conclusion (All)	49
	References	51
A	Appendix A	54
B	Appendix B	55

1 Introduction (OD)

Non-destructive testing (NDT) has become increasingly crucial in modern manufacturing and quality assurance processes, particularly in industries where material integrity is paramount and it may be expensive or disadvantageous to use destructive testing [1]. The objective of NDT techniques is to ensure the quality of an assembly or part by detecting and classifying indications based on size, position, and nature. This includes critical applications such as checking for defects in nuclear power plant pipes [2, 3], weld joints [4], and the quality of parts produced by additive manufacturing (AM) [5]. While radiography, penetrant testing, and Eddy current testing are widely standardised, ultrasonic testing (UT) stands out by providing deep, volumetric penetration for detecting internal flaws unlike surface-only methods.

Early industrial applications of UT included detecting macroscopic flaws such as cracks and voids [6] relying on single-element transducers acting as both transmitter and receiver [7]. However, the fixed focus and single angle of interrogation restricted the ability to characterise complex geometries or distinguish subtle material features [8]. The development of multi-element array transducers allowed the timing of pulses from individual elements to be controlled to steer and focus the ultrasound beam, producing images without moving the probe [9]. This also made post-processing techniques like Full Matrix Capture (FMC) possible. By recording all transmit–receive element combinations, FMC provides a rich dataset that can be reconstructed using the Total Focusing Method (TFM), allowing the inspection region to be focused at every point [10]. This offers major improvements in spatial resolution and defect detectability compared to conventional ultrasonic methods [11].

Beyond conventional defect detection, modern manufacturing has a growing need to understand the material itself to characterise the microstructure, elastic properties, and potential degradation before a flaw forms. Traditionally, techniques such as electron backscatter diffraction are used to characterise grain structure, but these are comparatively expensive, time-consuming, and not readily suited for rapid volumetric inspection [12]. As microstructural features such as grain cluster boundaries interact with ultrasound, UT could also provide information about the internal state of a material. If such information could be reconstructed reliably, UT offers a much faster, cheaper and non-invasive alternative to traditional methods. This was successfully demonstrated by Wang et. al (2025) [13] where an iterative process analysed 100 FMC datasets and accurately calculated the size of grains in copper as $33\mu m$, whilst metallography returned $35.5\mu m$. Grain clusters are much smaller than the ultrasonic wavelength and produce weaker, more distributed signals than defects. This makes the imaging problem significantly more sensitive to positioning errors, signal processing choices and reconstruction quality.

To produce meaningful volumetric reconstructions of microstructural backscatter, two methods can be used. Firstly, the use of 2D arrays can directly image volumes. In 2D arrays, the elements are organised on a plane and the data collected can be processed to give a 3D image of the internal structure. Conversely, 1D arrays have their elements aligned along a single axis and scan 2D slices through the volume. These can be used to generate 3D volumes by rotating the array on the surface of a target material to then be reconstructed using a Radon transform to give a volumetric representation of the imaged area [14].

To image an area larger than the array, multiple volumes must be scanned and stitched together. When imaging defects, it is easy to visually verify if a stitch has been successful due to the strong signals

present [15]. When stitching backscatter, this becomes unreliable, particularly in 3D. Therefore, a synthetic data pipeline is crucial in producing 3D TFM data to give statistical confidence for the stitching algorithms. Synthetic data would verify the expected error when imaging experimental data when there is no ground truth. For the synthetic data to reliably replicate the results seen experimentally, ultrasonic interactions must be accurately modelled.

Manual scanning, while suitable for basic defect detection, introduces significant operator-dependent variability. A 15MHz wave in copper has a wavelength of around 0.3mm, much larger than the expected grain sizes from 8.5 - 60 μm [16]. Whether using 1D or 2D arrays for imaging, inconsistent placement can lead to inaccuracies and render the data unusable. Small errors over many scans can lead to un stitchable images. This necessitates an automated and reliable system for positioning and rotation of the array. The data collected should be repeatable, ensuring similar experiments to that shown in [13] can be done more easily. Such a system would support the possibility of an end-to-end system in manufacturing, where target samples could be automatically scanned and processed to give the operator the full volumetric representation of the internal structure, whether that is for defect detection or microstructural backscatter. It requires the integration of mechanical design, motion control, signal processing, imaging, stitching and validation despite limited availability of experimental data into a single workflow.

However, this process produces large volumes of computationally expensive data. Modern ultrasonic inspection systems are increasingly leveraging the parallel processing power of Graphics Processing Units (GPUs) [17]. Originally designed for rendering graphics, GPUs excel at performing the same operation on many data points simultaneously. By offloading the image reconstruction to a GPU, inspection times can be reduced from minutes to seconds, enabling near real-time imaging and feasibility of large-area scans [18]. Along with efficient raw data processing and filtering, this is crucial for producing the synthetic data required, as well as reconstructing the experimental data in feasible time frames.

Each part of this report addresses a necessary contribution to that workflow. The system overview and specification define the overall requirements and establish the performance targets for the platform. The gantry design sections describe the mechanical structure that enables accurate and repeatable probe positioning. The automation and preprocessing sections then show how raw ultrasonic data can be collected consistently and converted into a form suitable for reconstruction. The imaging and stitching sections address the core signal processing challenges, including TFM reconstruction, hardware acceleration, and methods for combining multiple scans into larger images or volumes. Finally, the validation and experimental results sections assess how reliably these methods perform, both on synthetic datasets and on experimental scans, so that the limitations and strengths of the full system can be understood. All of these components are required if automated ultrasonic microstructural imaging is to become practical and repeatable.

2 System Overview and Specification (OD)

This system required two sub-groups to be formed. The "Build Team" provided the physical framework and spatial fidelity for accurate 3D positioning of the array on the target material, and automating the data collection process. The "Code Team" provided the signal fidelity, where the vast amounts of raw data were processed, imaged and stitched by validated algorithms.




Figure 1: System overview.

As shown in figure 1, the roles of each sub-group were clear. The data collection via the gantry and gantry controller was handled by the "Build Team" and the data coming from the PC to be processed was handled by the "Code Team". The MicroPulse FMC was a device used in the UNDT Laboratory at the University of Bristol which connects the ultrasonic arrays to the PC. The program used to interact with the MicroPulse FMC was Brain. The system required individual Product Design Specifications (PDSs) for each sub-group but also an overall PDS for the process. Some of the key requirements are listed below, with a supplementary PDS included in Appendix B in table 12.














Table 1: Key product design specifications.

Group	Category	Requirement	Priority	Assessment
All	Performance	System must produce spatially consistent volumetric reconstructions across full scan area.	MUST	Scan sample with defect present and verify no spatial distortion or discontinuities in stitched volume.
All	Performance	2D array scan and processing time must be within 40 minutes.	MUST	Measure time from scan start to final imaged output.
All	Validation	System must produce repeatable results under identical conditions.	MUST	Repeat full scanning procedure and compare similarity of outputs.
All	System	Mechanical motion, data acquisition and processing must operate as a unified pipeline.	MUST	Perform a full scan and confirm that minimal manual operation is required.
Build Team	Performance	The platform must maintain structural stability during operation without excessive vibration or deflection.	MUST	Visual inspection during a full scan to verify that the platform remain stable with minimal vibration and deflection.
Build Team	Performance	The automated scanning platform must provide translational positioning accuracy within $\pm 1mm$.	MUST	Positional error will be determined by comparing the actual position to the programmed position.
Build Team	Structural	The platform must be able to hold both 1D and 2D arrays for operation.	MUST	Operational test to verify that the design of the clamp and the arrays are compatible.
Code Team	Performance	Ultrasound scans must be imaged using TFM to produce 2D images of microstructural backscattering.	MUST	Visually inspect TFM images to ensure permanent features persist across images.
Code Team	Performance	Stitching algorithm must correctly recover expected translation between scans.	MUST	Evaluate performance using synthetic data and experimental data with known offsets.
Code Team	System	Code must accept configurable parameters (e.g. speed of sound).	MUST	Run pipeline with varied parameters and confirm correct behaviour.

3 Concept Generation of Scanning Gantry (AA)

A morphological analysis was used to develop ideas for a prototype structure and the mechanisms involved within the platform rig. This is shown in Table 2. Each pathway is shown, with each colour representing an option for a complete automated scanning gantry.

Table 2: Morphological analysis of the automated platform rig concept.

Feature	Concept 1	Concept 2.1	Concept 2.2	Concept 3
Structure 				
Array Holding Mechanism (z-axis)				
Translation Mechanism (x and y-axis)				
Rotation Mechanism (θ)				

- **Blue pathway** – The blue pathway involves a platform with an array holder and is attached to a gantry beam which utilises a slider mechanism to move in the Y direction. The array holder uses a linear actuator concept, which converts rotational motion from the motors into a linear motion to move the array in the Z direction. The rail beam is connected to the base by an X-axis slider at each end. The slider will be connected to the base using a lead screw and wheels. Another motor will be added to the array holder and connected to a clamp to add a rotational axis.
- **Purple pathway** – The purple pathway uses a turntable approach, where the test sample plate rotates manually on a ball bearing. An electronic angle sensor is integrated to measure the rotational position accurately. The sample plate is connected to the base and moves in the X and Y directions. The scanning array will be attached to a gantry beam using a clamp to move it up and down.
- **Green pathway** – The green pathway uses a manual gear system to rotate the test sample plate.

A small input gear is turned by hand and drives a larger gear that is attached to the underside of the sample turntable, which provides a reduction in gears that makes rotation smoother and easier to control. This gear system allows to control angular position, with the turntable angle calculated from the input gear rotation using a fixed ratio. For example, 1 full turn could equal 2° of the turntable rotation. The rest of the system follows the same idea as the purple pathway.

- **Orange pathway** – The orange pathway utilises a different workflow to the other ideas. It uses a concept of a robotic arm as the primary motion mechanism. This multi-axis arm provides the required movement in all direction, which allows the array to be placed and moved freely over the test surface.

The blue pathway is a simple, easy to operate, and flexible mechanism, which allows it to scan different sample sizes. However, the connector between the array holder and the rail beam must be designed carefully so that it can hold the weight of the array holder and the array itself while still providing a smooth movement along the rail beam. For the purple path, this concept has reduced the control system complexity because the automation requirements are lower than a fully motorised rotary axis. However, manual rotation can also introduce variability in speed changes and inconsistent positioning, which increase the risk of misalignment during scanning. The green path approach has the benefit of more controlled manual rotation when compared to the previous concept because the gear ratio provides finer and smoother angular adjustments. However, the angle accuracy depends on the mechanism as it may cause backlash and gear wear which can introduce angular error. The orange path has the ability to adjust probe angle, which improves signal quality on curved or angled surfaces. However, more complex calibration and programming are required because of the multiple arm joints which may reduce scanning accuracy.

The method of down-selecting concepts involved a combination of Pairwise Comparison and Multi-Criteria Decision Analysis (MCDA). An MCDA is a quantitative method of choosing between multiple ideas when there are several conflicting objectives [19]. The first step was to generate a criteria to evaluate the ideas against each other. Once these had been made, a Pairwise Ranking Matrix was used to indicate the strength of each metric relative to one another [20]. This matrix is shown in Table 3 and was used to calculate the weighting for each metric used in the MCDA. When combined, the MCDA and Pairwise Comparison methods provided a fair way to strategically choose the concepts from the morphological analysis.

From Table 3, it was clear to see that mechanical stability was the most important characteristic for the concept. This is reflected in the MCDA weightings, as shown in Table 4. The MCDA matrix indicated that the blue path gave the strongest score, based on the importance of each criterion. It is also important to note that each pathway had benefits and disadvantages that may not be adequately addressed by the MCDA. A final concept that combined the most advantageous systems of each pathway was generated and is shown in Figure 2.

Table 3: Pairwise ranking matrix.

	Ease of Operation	Mechanical Stability	Accuracy	Safety	Manufacturing & Cost	Total
Ease of Operation		0	0	0.5	0	0.5
Mechanical Stability	1		1	0.5	1	3.5
Accuracy	1	0		0.5	1	2.5
Safety	0.5	0.5	0.5		1	2.5
Manufacturing & Cost	1	0	0	0		1

Table 4: MCDA for the automated platform rig concept.

Criteria	Weighting	Blue		Green		Purple		Orange	
	↓ $\Sigma 1$ ↓	Rating (1-5)	Weighted Score	Rating (1-5)	Weighted Score	Rating (1-5)	Weighted Score	Rating (1-5)	Weighted Score
Ease of operation	0.05	3	0.15	2	0.1	2	0.1	1	0.05
Mecahnical Stability	0.35	4	1.4	3	1.05	3	1.05	2	0.7
Accuracy	0.25	5	1.25	3	0.75	3	0.75	2	0.5
Safety	0.25	4	1	5	1.25	5	1.25	3	0.75
Manufacturing and Cost	0.10	4	0.4	3	0.3	3	0.3	1	0.1
	Final Score	4.2		3.45		3.45		2.1	




Figure 2: Final concept.

4 Design, Iterations, and Manufacturing (AA, HS)

4.1 Base Frame (HS)

The design process of the automated gantry followed: material selection, detailed design, manufacturability, testing, and iteration prior to finalising the design. Material selection for the platform rig was primarily guided by considering the availability and cost of the materials. All the components were sourced from verified suppliers or from the university inventory itself. Each of the materials were carefully chosen to ensure the mechanical requirements were met whilst abiding by the budget constraints.

Firstly, Aluminium T-slot Extrusion was used as the base frame. All structural beams were cut from 40 mm $\times$ 40 mm aluminium T-slot extrusion into four lengths: 750 mm for the two longitudinal base beams, 580 mm for the two lateral outer side beams, 500 mm for the two lateral inner side beams, and 700 mm for the gantry beam. The length to be cut for each beam was decided based on the requirement stated in the PDS under the geometry category (see Table 12 in Appendix B). The base frame formed a rigid rectangular structure that established the overall dimensions of the rig and provided sufficient support for all other subsystems, which is illustrated in Figure 3.




Figure 3: Technical drawing of base frame.

It was crucial that the base was rigid as it acted as the backbone of the platform rig to maintain positional accuracy throughout the scanning process. Beam with 40 mm $\times$ 40 mm cross-section was selected since it was compatible with the gantry wheels that were incorporated into the assembly for translation. This allowed the wheels to roll smoothly along the two outer faces of the extrusion and guide the sliders' translation. The beams were sourced directly from the university inventory, making it a practical and cost-effective choice. Procuring alternative materials would introduce additional cost and lead time. All beams were joined at the corner using angle brackets, forming rigid 90° joints.

Leadscrews were used as an actuator to translate all three linear axes. 10 mm diameter leadscrews were used for the X-axis and Y-axis translation to accommodate the high combined mass of the sliders, gantry beam, stepper motors, and array holder mechanism. Using a larger diameter provided a stiffer and more robust shaft, resisting any bending and torsional deflection during operation. On the other hand, an 8 mm diameter leadscrew was used for the array holding mechanism, where the moving mass was lower. A smaller diameter was sufficient to withstand the loading conditions and reduce the overall weight of the array holder. The leadscrew converted the rotational motion of the motor into the linear motion of the sliders. This allowed accurate positioning, suitable for automated operation. Each leadscrew was coupled to their respective shaft coupler: 5 mm to 10 mm for the 10 mm diameter leadscrew, and 5 mm to 8 mm for the 8 mm diameter leadscrew. The shaft coupler accommodated the misalignment between the motor shaft and the leadscrew, therefore allowing the motor to rotate the leadscrew.

4.2 X-axis Slider (HS)

The X-axis sliders translated the gantry beam along the lateral side beams. The sliders were 3D printed using PLA, which is discussed in subsection 4.7. Figure 4a shows the initial design of the slider. The hole shown at the top of the slider in Figure 4b is where the leadscrew is placed for Y-axis translation. The slider on each side of the base frame translated the array holding mechanism simultaneously. Each slider was equipped with eight gantry wheels, with two gantry wheels on each of the two beam faces to aid for smooth and consistent translation along the rail. Installing high number of gantry wheels helped distribute the combined weight of gantry beams, array holding mechanism, and the ultrasonic array over a larger contact area.




Figure 4: (a) Initial and (b) final designs of the X-axis Slider.

After initial testing, some issues were noted. Firstly, the tolerance of the slider was too tight leading to excessive friction introduced between the slider and the beam. Moreover, the installation of the gantry axes were complicated due to the tight tolerances. The initial design of the X-axis slider was too large, resulting in long printing times. The design was therefore broken down into multiple smaller components. This allowed for easier maintainability, modification, and further iteration. The height was also increased from 200 mm to 260 mm which will allow more taller samples.

4.3 Y-axis Slider (AA)

The gantry beam positioned across the top of the rig would provide the support for the array holding mechanism to translate along the Y-axis. Initially, the Y-axis slider was designed using only a wheel-guidance system, as shown in Figure 5a. This offered a simple solution that could be manufactured easily. The testing phase highlighted a few issues, including the Y-axis slider tilting forward due to the weight of the array holding mechanism, and poor dimensional tolerances from 3D printing. As a result, this initial design did not provide sufficient stability for reliable scanning operations and affected the motion's smoothness.

The final design integrated the use of a rail-guidance system to the slider, as shown in Figure 5b. The use of both wheels and rails helped to provide a counter-moment, reducing unwanted tilting during scanning. The assembly consisted of steel rods and four gantry wheels that were connected to the slider and moved along the gantry beam. Extenders were designed and added to the top of the X-axis slider to support the rail-guidance system and provide attachment points for both ends of the steel

rods. The slider carriage acted as the central connecting element between the wheels and the rails, and the array holding mechanism. This allows the load to be transferred efficiently to the beam. The combination of the slider block, wheels, rails, and leadscrew allowed a stable and smooth translation along the beam.




Figure 5: (a) Initial and (b) final designs of Y-axis Slider.

4.4 Z-axis and Rotational Mechanism (AA)

The array holding mechanism was initially designed using only a rail-guidance system for motion along the Z-axis. Initially, the main body was 3D printed because it allowed for rapid iteration during early development, as shown in Figure 6a. However, the movement of the Z-axis slider carriage was not smooth due to friction between the slider and the main body. The design was then iterated to improve the structural rigidity and provide smoother movement.

The final array holding mechanism used a C-beam aluminium extrusion sourced from the university inventory as the main structural guide for vertical movement. This also modified the attachment method between the Y-axis slider and the array holding mechanism. The height of the Y-axis slider was increased, and T-slot nuts were inserted into the rear slot of the C-beam and fastened using bolts through the slider body. This made the connection between the two parts more secure and provided a more suitable mounting method for the extrusion. A slider carriage was designed to move the scanning array along the extrusion and was mounted within the C-beam. The linear movement was driven by an ACME leadscrew mounted along the C-beam. The ACME leadscrew was used because the threads are strong and it would provide smooth and controlled motion without slippage [21]. Four small wheel bearings and an ACME nut block were attached to the back of the slider carriage to improve guidance and stability.

The rotational mechanism was integrated with the array holder assembly to allow controlled rotation of the mounted array. This aligned with the PDS shown in Table 12, where the array holding mechanism must incorporate a rotational mechanism. A mounting block was attached to the Z-axis slider and served as the support for the rotational motor. The motor was mounted directly to the block and connected to a steel rod via a flexible shaft coupling. A clamp was connected to the lower end of the steel rod, which is discussed in subsection 4.5. This configuration provided a practical solution for the array to translate and rotate, while maintaining a simple arrangement that was easy to assemble and

operate. Figure 6b shows the assembled array holding mechanism and rotational mechanism.




Figure 6: (a) Initial and (b) final design of array holding mechanism and rotational mechanism.

4.5 Array Clamp (HS)

The array clamp was connected to the rotational mechanism, and held the array securely during operation. Figure 7 shows the design of the array clamp. When designing the array clamp, a space for the array's data cable to pass was included, whilst ensuring a strong grip on the array. There were interchangeable inserts at the side where fasteners can slide along it to adjust according to the size of the array. The array clamp used a simple fastener mechanism to tighten the clamp. The ultrasonic gel used on the sample led to the array suctioning to the material's surface, so it is important that the array clamp can hold the array without slipping. A rubber tape was taped on the surface of the bottom part of the clamp to increase the friction. The opening of the array clamp was modular to accommodate both 1D and 2D arrays, satisfying the PDS requirement shown in Table 1




Figure 7: CAD model of the array clamp.

4.6 Summary of Axes and Motion Parameters (HS)

Once assembled, the platform rig functioned as a 3-axis Cartesian motion system. All translations were driven by leadscrews connected to stepper motors via shaft couplers, in addition to gantry wheels to ensure smooth and repeatable motions. The CAD model of the platform rig and the full assembly

of the platform rig are shown in Figure 8a and Figure 8b respectively. During testing, the test sample was placed within the testing area, and it would remain fixed throughout the scanning.




Figure 8: Final platform rig.

4.7 Design for Manufacture (AA, HS)

The platform rig was manufactured using simple processes such as cutting, drilling, and 3D printing. As stated in the PDS, it was decided to limit the number of manufacturing processes and items to be made by the university technical services to eliminate any additional constraints.

4.7.1 Structural Frame (HS)

Aluminium T-slot extrusion was determined as the best material to be used as the structural frame due to its significantly simplified manufacturing process. All the beams could easily be cut to the desired length and connected perpendicularly to each other using an angle connector. The platform rig was assembled in a way that all the fasteners could still be accessed once fully assembled. It was designed to be modular so that any modification, replacement, or adjustment of any components could be done without having to disassemble the whole structural frame, except the inner side beams since it would require the removal of the outer side beams.

4.7.2 3D Printed Components (AA, HS)

For parts with complicated geometries, 3D printing was chosen since it allowed more room for testing and optimisation of the design. 3D printing is less time consuming and low cost compared to ordering parts from manufacturers, or fabricating parts from specialist workshops. Polylactic acid (PLA) was used for all 3D-printed components in the assembly, which included the X-axis and Y-axis sliders, X-axis extender, array clamp, motor housings, and bearing slots. CAD model for the motor housings and bearing slot are shown in Figure 49 in Appendix A. PLA was chosen primarily due to its availability and the cheap filament cost compared to other polymers. PLA is commonly used in fused deposition modelling and it is capable of producing dimensionally accurate parts at fine layer resolutions [22]. Standard PLA also possesses adequate compressive and flexural strength for the prototype.

The 3D printed parts were designed with manufacturability in mind. Print orientation was considered

during the design stage to minimise the amount of support needed. This reduced the post processing time and material waste since all the holes were self-supporting or would require minimal supports. In addition, large assemblies were divided into multiple smaller components. This allowed individual parts to be reprinted quickly during modification or repair without having to reprint the entire assembly. Dimensional tolerance was also an important consideration when designing the printed parts. Clearances were provided for bolts, rods, and sliding interfaces to account for the dimensional variation inherent in the fused deposition modelling printing process. The print settings in Bambu Studio were customised to balance dimensional accuracy, surface finish, strength, and printing time. The main settings used for the 3D printed components are summarised in Table 5.

Table 5: Customised print settings from Bambu Studio.

Parameter	Print Settings	Reasoning
Material	Generic PLA	Readily available.
Layer Height	0.16 mm	Good dimensional accuracy. [23]
Infill Pattern	Gyroid	Good overall strength. [24]
Infill Density	20%	Balance between strength and printing time. [25]
Support Type	Tree supports	Easy to remove.

4.7.3 Rotational Block (AA)

A block for the rotational mechanism was designed considering ease of manufacture, material availability, and structural stiffness. This block was attached to the Z-axis slider to integrate the rotational mechanism with the array holding mechanism and served as a support for the rotational motor. Aluminium was selected due to its good strength-to-weight ratio [26]. The design consisted of flat surfaces, simple hole patterns, and straight edges, which made it suitable for university workshop processes. This is shown in Figure 9. The main body was cut from aluminium stock using straight cuts, which simplified fabrication and improved dimensional consistency. Milling was then used to achieve the required dimensions. The holes were drilled according to the specified dimensions. The geometry of the motor support was kept simple to reduce manufacturing complexity. Simple hole patterns were preferred to reduce fabrication time and minimise dimensional errors.





Figure 9: Drawing of the rotational block.


5 Automated Data Collection(RZ)

5.1 Electronics (RZ)

The electronics system provided the interface between the automation and the gantry. It comprised of two independent Arduino microcontrollers, a CNC shield [27] carrying four DRV8825 stepper motor driver modules used for the translational motors, a DRV8833 dual H-bridge driver module for the rotational motor, and an external 24V power supply. Both Arduinos were powered by a PC as they were constantly connected via USB during use. This is shown in Figure 10.



(a) DRV8833 motor driver with Arduino UNO



(b) CNC shield with DRV8825 motor drivers

Figure 10: Motor control hardware configurations used in the system.

The CNC shield provided dedicated sockets for up to four DRV8825 stepper driver modules, along with the signal routing needed between the Arduino's digital output pins and each driver's step, direction, and enabled inputs. This avoided the need for manual wiring of the step and direction signals. Additionally, a jumper was connected between the X and A axis step and direction pins so that the A axis socket would clone the X axis (as there were two X-axis motors). The motor drivers were the main heat source thus, to prevent overheating of the electronics, a heat sink was placed on top of each driver and the potentiometers of the drivers were adjusted so that the current supplied would not exceed half the maximum current rating of the drivers. This adjusted value was set and measured to be 0.7A using a multimeter.

The decision to use two separate Arduinos rather than a single board was a design decision driven by hardware constraints and interface incompatibility. G-Code is mainly used for linear Cartesian systems, and while it is possible to map rotational motion onto a linear axis by treating degrees as millimetres, doing so would require modifying the G-Code Reference Block Library (GRBL) firmware and introducing a new coordinate system which can make physical commands unclear. A dedicated Arduino for the rotational motor avoided this issue, and allowed the rotational axis to be commanded directly in degrees. An Arduino Mega with an extended CNC shield would have provided the extra space to install another motor, but this would have incurred extra cost and procurement time. The two Arduino solution was assembled entirely from components found in the lab with no additional cost.

5.1.1 Motor Selection (RZ)

This section outlines the motor selection process for the X, Y, and Z axis drive system for the gantry assembly. The minimum requirements of the motor are based on driving the X-axis conditions as they are most demanding. These system parameters are shown in Table 6.

Table 6: X-axis motion system parameters.

Parameter	Value	Notes
Total moving mass	4,736 g	Measured
Motors	2	One per lead screw, load shared equally
Axis orientation	Horizontal	No axial gravity component
Linear guide type	V-slot extrusion with POM wheels	
Lead screw outer diameter (d)	10.1 mm	Measured
Lead screw mean diameter (d_m)	9.1 mm	
Lead screw pitch (p)	2 mm	Measured
Lead screw starts	1 (single start)	Assessed by inspection
Lead screw lead (l)	2 mm/rev	$l = \text{pitch} \times \text{starts}$
Required precision	A few millimetres	Open-loop stepper acceptable

There were two motors that shared the total load equally when driving in the X-axis. Hence, the mass being driven per motor was 2.37kg. The torque required to drive a lead screw under axial load was calculated using equation 1, derived from [28].

$$T_R = \frac{F d_m}{2} \left(\frac{l + \pi f d_m}{\pi d_m - f l} \right), \text{ where } d_m = d - \frac{p}{2} \quad (1)$$

T_R is the torque required to drive the load, F is the total axial force, d_m is the mean thread diameter, and f is the coefficient of friction at the thread interface and p is the pitch of the lead screw. The value of d_m was found to be 9.1mm. A target of 0.5 ms^{-2} is assumed to be a reasonable acceleration and the thread friction coefficient of the aluminium threaded block against clean and dry steel was assumed to be $f = 0.61$ [29], although it should be noted that the lead screws were lubricated prior to testing which reduced the friction value. Substituting all values into equation 1, a required torque of 34.4 mNm was calculated.

A stepper motor was selected as the most appropriate motor type for this application for several reasons. Unlike DC brushed or Brushless DC motors, the stepper motors chosen moved in a well-defined angular increment of 1.8 degrees under no load which makes them inherently suited to open-loop positioning. This removes the need for an encoder to determine its position. This is sufficient for the precision requirement of this gantry. NEMA 17 stepper motors, which are used in similar projects [30], with 420mNm torques were chosen for this application, allowing a safety factor of over 12 times the required amount.

5.2 Gantry Automation (RZ)

The automation system was responsible for positioning the ultrasonic array at each scan zone whilst also controlling the rotational orientation. It was made up of three primary layers: a 'Scan PC' where a python script is left to run, two Arduino microcontrollers, and the mechanical gantry driven by the stepper motors. The system overview is shown in Figure 11. GRBL was adopted for the motion control

firmware of the translational axes. The translational Arduino controlled the translational axes (X, Y, Z) of the gantry, which ran GRBL firmware. GRBL is an open-source firmware which is optimised for stepper motor machines and uses G-code as a standardised language [31]. Key motion parameters including steps per millimetre, maximum feedrates, and axis acceleration limits were configured directly within GRBL. The rotational Arduino drives the rotational axis and runs on a custom firmware.




Figure 11: Automation system overview

5.2.1 Scan path strategy (RZ)

The automation system was designed to follow a calculated predefined motion path on the scan object, as shown in Figure 12. A raster scan strategy and boundary offset principle was adopted, consistent with established practice in automated ultrasonic NDT [32]. It was assumed that every scanned object was a flat rectangle. The scan area was defined by specifying the four corner coordinates (relative to the origin which is manually set before scanning starts) of the object. The Scan PC then extrapolated the boundary of the object and identified the starting position of the scan by selecting the corner whose Euclidean distance from the origin is smallest as the start of the first scan area.




Figure 12: Diagram of machine pathway

The scan zone centres were found by distributing the zones at uniform intervals across the area of the object. The step between adjacent zone centres were found via Equation 2. The first and last zone centres were positioned so that the area of the array remained within the boundary of the object.

Equation 3 was used where Min is the smallest coordinate and Max is the largest coordinate for X and Y respectively.

$$\text{Step} = \text{Scan Size} - \text{Overlap} \quad (2)$$

$$\text{First Zone} = \text{Min} + \frac{\text{Scan Size}}{2}, \quad \text{Last Zone} = \text{Max} - \frac{\text{Scan Size}}{2} \quad (3)$$

Equation 3 ensures that the centre of the first and last zone was offset inward from the lower and upper boundary of the scan object, respectively by half the scan footprint. This is necessary as a zone centre positioned exactly by the boundary of the object would place half the array outside the defined scan area, which may result in faulty data acquisition.

5.2.2 Scan Sequence Execution (RZ)

Prior to scanning, the coordinates of the 4 corners of the object were inputted and the translational Arduino was initialised. This Arduino transmitted a set of configuration commands to the GRBL firmware that defined the operation parameters. These commands established the unit system (mm), positioning so that every target coordinate was relative to its home position, and feed rate interpretation, which were adjusted for each axes. The gantry tracked its position in X,Y, and Z, by counting its motor steps relative to the origin point. The Scan PC then executed the scan sequence shown in Figure 13. Firstly, the gantry would move to the scan centre. To ensure the physical position of the gantry and the software state remain synchronised throughout, the Scan PC actively queried the translational Arduino to confirm it had come to rest before proceeding. Once stationary, the Z-axis was lowered to the surface of the object at a reduced speed to ensure controlled contact. The rotational Arduino was then commanded to rotate at each specified angle, confirming each move before the user acquired data and the next angle was complete. Once a full rotation was done, the array returned to its 0 degree reference point.




Figure 13: Scanning sequence of the automated gantry system per scan area

5.3 Calibration (RZ)

Prior to scanning, the gantry system was calibrated to ensure it was accurately moving to desired positions. Firstly, the correct steps-per-millimetre conversion was established. This parameter defined how many motor steps were needed to produce the required distance of linear motion along the sliders. The setup in Figure 14 was used to calibrate the axes. The Z axis did not need to be calibrated as the scan height would be manually set before each scan. The contact quality was visually verified by looking at the time-trace plot on the UT PC.




Figure 14: Calibration Setup

The gantry was first set to a known reference point and then programmed to travel a known distance. The actual distance moved was measured from the reference point and substituted into Equation 4. This process was repeated for each axis until the correction produced no further meaningful reduction in error, at which point any source of error stemmed from the gantry and not the motors themselves.

$$\text{Calibrated Step [mm]}^{-1} = (\text{Initial Steps [mm]}^{-1}) \left(\frac{\text{Programmed Distance}}{\text{Actual Distance}} \right) \quad (4)$$

The default value of 500 steps/mm for each axes produced large positional error. Table 7 shows that the X axis converged to a stable step/mm value within two iterations, whereas Table 8 shows the Y axis still showed a large error even after iteration. This disparity is linked to several mechanical factors that disproportionately affect the Y axis' measurements.

Table 7: X axis steps/mm calibration.

Iteration	Steps/mm	Commanded (mm)	Actual (mm)	Error (%)
0 (default)	500.0	—	—	—
1	131.6	10	38.0	+280.0
2	132.5	142	141.0	−0.7
3	132.5	100	100.3	+0.3

Table 8: Y axis steps/mm calibration.

Iteration	Steps/mm	Commanded (mm)	Actual (mm)	Error (%)
0 (default)	500.0	—	—	—
1	149.6	10	9.5	−5.0
2	157.5	10	9.5	−5.0
3	143.1	50	55.0	+10.0

A main source of error was the poor interface between the array clamp to the Z-axis slider. This led to the array being able to rotate freely during measurement. The unintended rotation altered the effective position of the reference point on the array relative to the ruler, which introduced a systematic offset into the measured displacement. The X-axis sliders also had its structural limitations. Due to the weight of the Z axis slider and its attached components, the gantry had a constant tilt, causing one side of the array to have more pressure than the other. These mechanical limitations meant that a

full calibration procedure for the rotational mechanism was unsuccessful. Visual inspection confirmed that the array could not rotate to desired angles. Had the mechanical assembly been sufficiently rigid, calibration would have been carried out by programming the array to a series of known angles and physically measuring the actual angle it reached using a protractor. This process would be repeated iteratively and use the same approach in Equation 4.

6 Gantry Performance Evaluation (AA, HS, RZ)

The automated gantry was evaluated against the PDS after testing. Table 9 shows the PDS points which were failed or partly achieved, with all remaining points having passed. There were several issues which introduced limitations to the gantry’s capabilities.

Table 9: Gantry PDS compliance summary.

PDS Requirement	Priority	Target	Achieved	Status
Translational positioning accuracy	MUST	± 1 mm	X-axis: ± 1 mm; Y-axis: ± 5 mm	Partial
Structural stability during operation	MUST	No excessive vibration or deflection	Noticeable forward tilt; X-axis swinging during translation	Fail
Rotational mechanism incorporated	MUST	Array holder can rotate	Rotation integrated but imprecise; calibration unsuccessful	Partial
Automatic scanning operation	MUST	Full scan cycle without human intervention	Manual data acquisition trigger required at every scan position	Partial

First, there was a slight tilt in the gantry during scanning. The X-axis slider and its attached components caused the gantry beam to tilt forward due to loose tolerancing. Although this forward tilt was reduced after several design iterations, there was still a significant effect on the accuracy of testing. For example, when pressing the array onto the scan object, the gantry would tilt the other way to accommodate for the force being applied to the scan object. Additionally, the X-axis sliders did not have enough structural rigidity during translation and the slider would swing during testing due to this lack of support. To resolve this, guide rails could be installed along the X-axis, similar to the guide rails used in the Y-axis as shown in Figure 5b. A longer contact length between the slider and the leadscrew might be able to compensate the tall slider and improve the stability. Other future improvements would include to replace the C-beam with 3D printing of different slider designs compared to the failed initial design shown in Figure 6a which would significantly reduce the mass.

Secondly, the metal plates that were incorporated in the final assembly, shown in Figure 8, could be better integrated in the slider design. The metal plates were introduced as a mitigative action against the inaccuracy of in translation from 3D printed parts. The 3D prints failed to withstand repetitive translation and the threading of the leadscrew destroyed the slider hole. This can be prevented by using metal as the contact point for translation with the leadscrew. Another source of translational inaccuracy was the Y-axis slider which did not have smooth motion, specifically near the ends of each side of the leadscrew. Metal lubricant was initially sprayed on all the leadscrews to reduce friction but the accuracy of translation remained the same. It was assumed that the mechanical stalling were caused by damage to the leadscrew as it was procured from the university inventory. Swapping this out for an ACME leadscrew may resolve the issue. It could also be that the metal-to-metal contact

between the leadscrew and metal plates were not smooth, causing friction and stalling.

The rotational mechanism was successfully integrated into the array holding mechanism, but it did not achieve the required precision needed for rotation. This is likely due to small clearances in the rotational block and at the connection between the array clamp and the rod. The clearance of the hole in the rotational block was not intentional, as the drawing submitted to the technical services had an exact dimension of 10 mm, as shown in Figure 9. However, the actual hole dimension of 10.2 mm was measured after the block was manufactured. In addition, uneven pressure on the scan object caused by the forward tilt of the supporting sliders also increased friction on one side of the array and reduced rotational accuracy. This meant that using the Radon transform to image 3D volumes using 1D arrays was not possible as it required very high rotational accuracy. This would reduce the reliability of volumetric reconstruction and make stitching more difficult. To address this issue, the rotational block should be manufactured with tighter tolerances. The rotational shaft could also be supported with a bushing or bearing to reduce wobble and help the array to remain normal to the sample surface during scanning. Alternatively, a stepper motor with a higher holding torque could be utilised to hold the array in place and resist the forces which affect its rotational position.

The gantry achieved semi-autonomous operation. Translational and rotational motion were coordinated by the Scan PC without manual input at each scan position. However, two manual dependencies remained. First, data acquisition had to be manually triggered on the University PC at every scan position. Establishing a link between the Scan PC and the University PC would be the primary step required to achieve full autonomy. Secondly, the Z-axis scan height had to be set manually by the user prior to each scanning session. Although this was functional, this introduced variability in the contact force between the array and the scan object, potentially causing inconsistencies between scans. A more robust solution would be to integrate a compression load cell [33] between the array and the clamp. This would provide a direct measurement of the contact force at each scan position allowing a repeatable contact condition to be set automatically. However, this would require a redesign of the existing clamp. As a first prototype, the platform establishes a good foundation for future work and validates the core concept of automated ultrasound testing for microstructure characterisation.

7 Preprocessing and Imaging (OB, OD, TT)

7.1 Filtering (OD)




Figure 15: Data process shown on data with a circular defect present.

The raw data collected by the automated gantry had to be processed prior to imaging. This included filtering, cropping, compressing, and generating the metadata required by the imaging algorithm such as the speed of sound and element positions. An overview of the raw data pipeline is shown in Figure 15

and all information can be found in the project’s GitHub [34]. The raw data is stored in .mat files by the MicroPulse FMC and PC. Each scan was averaged 64 times to reduce the effect of random noise and any signal sample which crossed the threshold (1000) was replaced by the subsequent sample. The raw .mat files were converted into multiple .csv files (metadata, element positions, firing orders) and a .h5 file for the time-trace data for quicker read/write speeds. The time-trace data was cropped to reduce memory bandwidth for future operations by removing unnecessary data, and compressed after filtering by limiting the floating point precision to 32 bits. Reducing this from 64 bits provided sufficient accuracy while significantly reducing memory requirements. This is shown in Figure 16 where the compression led to no noticeable difference in image quality. The mean percentage difference was $9.83\text{e-}07\%$ and the maximum percentage difference was $2.57\text{e-}04\%$.




Figure 16: Comparison of TFM images for an aluminium sample using a 128 element, 10MHz array: (a) TFM image obtained from 64-bit precision, (b) TFM image obtained from 32-bit precision, and (c) the percentage difference between the images.

The filter used for these operations was a bandpass filter with a cosine-tapered transition [35]. Each signal was first centred by subtracting the mean to eliminate DC offsets. The signals were then transformed into the frequency domain using a Fourier Transform. A bandpass filter was applied over the frequency domain from f_{min} to f_{max} (the minimum and maximum frequencies that the filter would affect). These were defined as a percentage of the centre frequency, f_{centre} . Rather than a simple rectangular window, a tapered cosine window was applied to the selected frequency indices. This is characterised by the α parameter which controlled the drop-off from the peak of the filter. This allows the filter to be more rectangular or Gaussian depending on the data being processed. An example of the filter being used on a pure sample signal is shown in Figure 17.

As seen in Figure 17, the bandpass filter successfully removed unwanted noise from the time domain signal. For pure samples, a smooth filter ($\alpha = 1$) with a wide frequency band ($f_{centre} \pm 80\%$) produced high-contrast, low-noise reconstructions suitable for imaging microstructural backscatter. On the other hand, the best results for imaging defects used a near-rectangular filter ($\alpha = 0.2$) and a much shorter frequency band ($f_{centre} \pm 20\%$). An example of the filter used for samples with defects is shown in Figure 46 in appendix A. The importance of correct filtering is highlighted in subsection 7.3.

For defect detection, the dominant signals arise from strong, spatially localised reflections, producing coherent responses concentrated around the transducer’s centre frequency. A narrow, rectangular bandpass filter was sufficient to isolate the high amplitude components while suppressing out-of-band noise. In contrast, microstructural backscattering arises from the cumulative interaction of




Figure 17: Time domain A-scan (left) of a pure aluminium sample filtered using a bandpass filter in the frequency domain (right).

the ultrasonic wave with many grain boundaries of varying orientation and size. This produces a more distributed and stochastic signal, where the energy may spread over a broader frequency range. Consequently, a wider and smoother bandpass filter was required to preserve the full spectral content of the backscattered signal.

7.2 Speed of Sound (OD)

Accurate knowledge of the material's speed of sound is essential for TFM imaging, as it directly governs pixel scaling and the spatial localisation of reconstructed features. It was crucial for the stitching operations to have accurate pixel sizing. To obtain a reliable estimate, the target material was scanned at a minimum of three distinct locations using a 5MHz array, and the mean of the results was taken to reduce the effects of local variability. A 5MHz array was selected to ensure sufficient penetration depth, enabling clear observation of back wall reflections. Higher frequency waves experience greater attenuation within the material, which can weaken the reflections and reduce measurement reliability.




Figure 18: Time trace with triggers for speed of sound calculation.

As shown in Figure 18, two trigger points were placed on the first and second back wall reflections. The speed of sound was calculated from the time difference between these reflections. Measuring the time-of-flight from the initial excitation to the first back-wall reflection would introduce latency associated with the transducer and acquisition system. By instead measuring between successive back-wall reflections, this constant offset was eliminated, resulting in a more accurate estimate of the propagation time within the material.




Figure 19: A closer look at the speed of sound triggers.

Consistent trigger placement was critical to ensure repeatability of the measurement. For each material, an amplitude threshold was defined such that a trigger was placed at the point where the signal first exceeded the value. For the pure aluminium sample, a threshold of 0.2 was used, as shown in Figure 19. To further improve accuracy, the trigger position was linearly interpolated between the two nearest samples, reducing discretisation error. Although this approach produced consistent alignment of trigger points across successive back wall reflections, it required some level of trial and error depending on the material used. By visually inspecting the time traces, an appropriate threshold was chosen for the triggers to correctly intersect the back wall reflections as shown in Figure 19.

The time difference between the two trigger points was measured and combined with the known sample thickness, obtained using callipers, to calculate the speed of sound. It should be noted that this measured time corresponds to a round-trip propagation between reflections, and therefore represents twice the physical thickness of the sample. The speed calculated in aluminium and copper using this method was 6369.9 m/s and 4636.1 m/s. This is in accordance with the accepted values of 6420 m/s and 4760 m/s [29]. This corresponds to errors of 0.78% and 2.60% respectively.

7.3 Total Focusing Method (OD)

The Full Matrix Capture (FMC) acquisition process provides a comprehensive dataset containing all transmit–receive element combinations. However, this data alone does not directly yield an interpretable image of the material. An imaging algorithm is therefore required to reconstruct the spatial distribution of reflectors within the inspection region.

The Total Focusing Method (TFM) is widely used for this purpose, as it enables high-resolution image reconstruction by synthetically focusing at every point within the domain. Unlike conventional phased array imaging methods, which rely on predefined focal laws and limited angular coverage during acquisition [10], TFM performs dynamic focusing retrospectively using the complete FMC dataset.

For each pixel location, the travel time from every transmit–receive element pair is calculated, and the corresponding signals are coherently summed. This ensures that all available aperture information contributes to every image point, resulting in improved spatial resolution and increased signal-to-noise ratio. However, this comes at the cost of significant computational complexity. The image intensity at a given point was calculated by summing the time-aligned signals from all transmitter–receiver pairs, as defined in Equation 5:

$$I(x, y, z) = \left| \sum_{t=1}^{N_{tx}} \sum_{r=1}^{N_{rx}} \left[A_{tr} \left(\frac{d_t(x, y, z) + d_r(x, y, z)}{c} \right) \right] \right| \quad (5)$$

where $I(x, y, z)$ represents the reconstructed image intensity at the point (x, y, z) within the inspection region. The summations were performed over all transmitting and receiving array elements, where N_{tx} and N_{rx} denote the total number of transmit and receive elements respectively, and the indices t and r correspond to a specific transmitting and receiving element pair. The term A_{tr} represents the recorded A-scan signal corresponding to transmission from element t and reception at element r . The quantities $d_t(x, y, z)$ and $d_r(x, y, z)$ represent the geometric distances between the image point (x, y, z) and the transmitting and receiving array elements respectively. These distances are defined as

$$d_t(x, y, z) = \sqrt{(x - x_t)^2 + (y - y_t)^2 + (z - z_t)^2}, \quad d_r(x, y, z) = \sqrt{(x - x_r)^2 + (y - y_r)^2 + (z - z_r)^2} \quad (6)$$

where (x_t, y_t, z_t) and (x_r, y_r, z_r) denote the spatial coordinates of the transmitting and receiving elements. The total propagation time from the transmitting element to the image point and back to the receiving element is therefore $(d_t(x, y, z) + d_r(x, y, z))/c$, where c is the speed of sound in the target material. Linear interpolation was used to extract the respective intensity from the A-scan, A_{tr} .

To improve the visual representation of the reconstructed images, a Hilbert transform was applied as a post-processing step to extract the envelope of the TFM image. While the Hilbert transform is conventionally defined for a 1D time-signal, the reconstructed TFM image retains oscillatory behaviour along the depth of the image due to the underlying ultrasonic wave field [36]. This treated each column of the image as a 1D signal. The analytic signal is defined as

$$I_a(x, y, z) = I(x, y, z) + iH\{I(x, y, z)\} \quad (7)$$

where $H_z\{\cdot\}$ is the Hilbert transform applied along the depth axis (z). The envelope image was then obtained via

$$A(x, y, z) = |I_a(x, y, z)| \quad (8)$$

This operation suppresses the high-frequency oscillations associated with the ultrasonic carrier wave, producing a smoother intensity field that better represents the spatial distribution of backscattered energy. This is shown in Figure 20 where the use of the Hilbert envelope significantly enhances the visibility of microstructural backscatter in the pure sample. Furthermore, the reduced oscillatory content improves the robustness of cross-correlation during stitching in section 8, as the envelope provides a more stable intensity pattern.

The Hilbert transformed TFM image in Figure 20 was further improved by normalising the image amplitude with respect to its maximum value and converted to a decibel scale using a logarithmic




Figure 20: Comparison of TFM images for an aluminium sample using a 128 element, 10MHz array: (a) TFM image obtained from the filtered data, (b) Hilbert transform applied to the image, and (c) normalised image in decibel scale from 0 to -20dB.

transformation. The dynamic range compression enhanced the visibility of low amplitude features while preserving relative contrast. The importance of filtering is also highlighted in Appendix A in Figure 47, where the same images were processed on unfiltered data, and any semblance of microstructural backscattering is not visible.

Prior to stitching, it was found that cropping the warped parts of the images improved the stitching results. The reason for this is presented in Figure 21.




Figure 21: Warping occurring on a persistent feature during a scanning sweep in aluminium.

The scans shown in Figure 21 were taken during an experimental procedure with pure aluminium where a 3D printed guide was used to move the scanning array by 5mm increments over the surface of the material. This experimental procedure is shown in Figure 22a. Figure 21 highlights the effect of warping on the extremities of the array's length, where the red contoured feature maintains structure directly under the array (middle two images) and stretches on the extremities (first and last image). This occurs as features on the sides are being imaged by elements on the opposite end of the array, at steep angles. The use of an angular filter was explored where only elements within a certain angle of a pixel would contribute to the imaging, but this led to less spatial resolution compared to using all elements available at each point in the image. The use of a binary threshold was also found to improve the reliability of the stitching and this is explained further in section 8. An example of the cropping and binary threshold is shown in Figure 22b.

7.4 Hardware Acceleration (OD)

The computational cost of TFM is significant due to the requirement to evaluate contributions from all transmit-receive element pairs at every pixel in the image domain. For typical imaging parameters,




Figure 22: Experimental scanning setup and pre-stitch processing.

e.g. 400×400 pixels and a 128 element array, this results in 160,000 transmitter-receiver distance calculations with 16,384 time traces to look across. A naive CPU implementation therefore becomes prohibitively slow when reconstructing many images, generating large volumes of simulated data, or performing 3D imaging on 2D arrays.

The code was written and compiled in C++ with the use of the OpenMP library [37] for multi-threading on the CPU. Moreover, an initial GPU-based implementation was developed to accelerate the TFM computation by exploiting the inherent parallelism of the problem. Since each pixel calculation is independent, these could be assigned to individual GPU threads [38]. In this baseline approach, each image was processed sequentially by loading raw FMC data and transferring it to an AMD Radeon RX 6900XT with 16GB of VRAM. A GPU kernel was launched for the image, and after the calculations were completed, the results were copied back to host memory. This process repeated for subsequent images. While this reduced image computation time compared to the CPU implementation, it introduced substantial overhead due to repeated kernel launches, frequent CPU-GPU memory transfers, and redundant preparation of geometric data. Additionally, the GPU remained idle during CPU-side file I/O, leading to poor overall utilisation. This is illustrated in Figure 23.




Figure 23: Sequential GPU processing. Diagram not to scale.

A batched processing pipeline was implemented to improve GPU utilisation and overall throughput. Instead of processing images sequentially, multiple FMC datasets were grouped into batches and processed simultaneously on the GPU. A single kernel was now used to process an entire batch of images rather than launching a kernel per image. This improved kernel launch overhead efficiency and

GPU occupancy. Moreover, the geometric quantities required for TFM were invariant across images within a scan. These were computed once and reused for all subsequent batches, eliminating redundant computation and memory transfers. Another improvement was the use of pipeline parallelism between the CPU and GPU. While the GPU processed batch n , the CPU simultaneously loaded and prepared batch $n + 1$. This overlap hid some of the I/O latency and increased GPU utilisation. The GPU batching is illustrated below in Figure 24.




Figure 24: Batched GPU processing. Diagram not to scale.

The optimisation shown in Figure 24 compared to that in Figure 23 focused on maximising throughput rather than minimising latency for a single image. During intensive scanning operations, this was preferred as the GPU preparation overhead was negligible thanks to the efficiency gains observed in the overall processing time. On the other hand, when imaging a few scans with varying parameters, the sequential GPU processing was more suitable. The results of the hardware accelerations is summarised in Table 10 which shows that using a batched GPU method was 726 times faster than the Python implementation, and 120 times faster than multi-threaded C++. Future implementations could extend the parallelism to the synthetic data’s FMC simulations and increasing the speed of stitching. More work could be done to optimise the initial batch load, which currently heavily affects the overall imaging time.

Table 10: Time taken for a 400x400 TFM image.

Language	Single/Multi Threaded	Hardware	Time Taken [s]
Python	Single	CPU	131.4
C++	Single	CPU	108.8
C++	Multi	CPU	21.81
C++	Multi	GPU	0.223
C++	Multi (Batch)	GPU	0.181

7.5 2D Array Processing (OB, OD)

Rotating a 2D array about its centre was found to be beneficial in improving stitching performance in experimental scans. During an experimental procedure performed on a copper sample shown in Figure 25a in Appendix A, four scans were taken 90° to each other at each position, and the array was moved 2mm between each position. By rotating the array as shown in Figure 25a, the effective coverage of the internal volume was improved. Each rotation gave unique angles of the same volume, strengthening persistent features.

Three methods were compared to accumulate the 4 rotations into 1 image, maximum, mean and median. The accumulation of each successive translational position was stitched by the methods described in section 8 and the results are shown in Figure 26. All accumulation methods identified the




Figure 25: 2D array rotations and the corresponding isosurface produced via max accumulation.

correct experimental shift. However, max accumulation provided the lowest dip either side of the peak. This gave it more certainty than the other two methods, and so was the choice of accumulation when combining 2D array data. The normalised curves of Figure 26a are shown in Figure 26b. Normalisation was used out so the graphs' shapes could be compared.




Figure 26: Original and normalised correlation graphs of different accumulation methods.

7.6 Radon Transform (TT)

The Radon reconstruction method can be employed to reconstruct 3D volumes from 2D slices imaged by a 1D array. As described in Section 4.4, the array could rotate around a fixed axis normal to the sample surface through N angular positions, θ_k uniformly spaced over $[0^\circ, 180^\circ)$. Each θ_k produced a 2D TFM slice of the volume, and the N slices were stacked into a 3D array indexed by (z, x, θ) , representing the same physical region imaged from progressively rotated orientations. The angular step between successive scans was set to 5° , corresponding to $N = 36$ slices over the half-rotation. This value was identified as the point at which image quality converges, with finer sampling producing no observable improvement [39]. At each depth z_0 , the row of pixels across every angular slice was assembled into a sinogram in (θ, x_θ) space, and an inverse Radon transform recovered the reflectivity distribution at that depth. Stacking the per-depth reconstructions produced the full 3D volume, as

shown in Figure 27.

This method relied on each row of the 2D TFM slice at z_0 behaving as a line integral of the reflectivity through the full 3D volume, rather than through a narrow slab at $y = 0$. This was because TFM pixel values were computed from true 3D element-to-pixel distances: scatterers off the $y = 0$ plane contributed to the reconstructed row with the correct time of flight and were therefore included in the projection. Element elevation directivity attenuates contributions from scatterers far from the scan plane, placing a practical limit on the effective reconstruction volume. The minimum required angle for accurate, low noise reconstruction using this method was 5° [39]. The implementation used was verified against synthetic volumes, as shown in Figure 27.




Figure 27: Radon reconstruction pipeline. **(a)** The 1D array (red) rotates through angle θ about a fixed vertical axis at the centre of the sample. Three scatterers s_0 , s_1 , s_2 are placed at known positions inside the volume. **(b)** Three TFM slices at representative scan angles. **(c)** Sinogram at $z_0 = 20$ mm, formed by stacking the z_0 row of every angular scan. Overlaid curves show the theoretical projection trajectories for s_0 at $(x, y) = (0, 0)$ mm and s_1 at $(x, y) = (5, 3)$ mm. **(d)** Inverse-Radon reconstruction. All three scatterers appear at their known positions.

8 Volume Stitching (OB)

Once the images were processed, stitching was performed to reconstruct the full scanned volume. Multiple stitching techniques were evaluated to determine the most suitable approach for this application. Table 11 summarises the methods considered, along with their underlying principles, capabilities, and limitations.

Template matching methods (TMMs) were selected as the most appropriate approach for this application, as the expected variation between scans was limited to translation. This assumption was

Table 11: Comparison of image stitching methods.

Method	Principle	Strengths	Limitations
Phase Correlation	Estimates translation using phase differences in the Fourier domain.	Highly robust to noise, precise translation estimation.	Sensitive to boundary effects and shape changes. [40]
Feature-Based Methods	Detects key points between images.	Handles translation, rotation, and scaling, widely applicable.	Struggles with repetitive or textureless images [41].
Template Matching	Searches for the best match of a template using similarity metrics.	Simple and intuitive, effective for translation detection.	Limited to translation, computationally expensive, sensitive to variations in scale, illumination, and warping [41].

based on the design of the gantry, where the array would either return to its original orientation or stay at a 180° angle for the next scan, meaning there should be no change in angle from the scans. Accounting for rotation or scale variations would introduce unnecessary computational complexity and increase the risk of error without providing additional benefit. The effects of rotational misalignment are discussed in subsection 10.3.

Different TMMs rely on different similarity measures, which include the sum of absolute differences (SAD) and sum of squared differences (SSD) [41]. However, these approaches operate directly on raw pixel intensities and are therefore sensitive to changes in brightness and intensity [42]. Normalised cross-correlation (NCC) was selected as the similarity measure due to its robustness to variations in signal amplitude and brightness. Unlike SAD and SSD, NCC standardises the signals to have zero mean and unit variance before comparison. As a result, NCC compares the similarity of the intensity patterns between signals. This makes it invariant to linear changes in brightness and contrast.

However, NCC still depends on pixel intensities. It evaluates how well the relative variations in intensity between two signals align, making it well suited for detecting translations in signals with consistent underlying features but varying amplitude [41]. NCC provides a convenient interpretation, as it produces a bounded similarity score and is less sensitive to large local deviations compared to SSD, which squares intensity differences.

8.1 Stitching Implementation

Direct voxel-wise comparison of two 3D volumes is computationally expensive and sensitive to noise. Conversely, collapsing the entire volume into a single 1D signal (as is often done in 2D stitching) results in a loss of spatial information, making it difficult to distinguish between multiple features and can lead to the wrong shift being chosen, especially when the images in question have random grain structures.

To mitigate noise sensitivity while avoiding substantial loss of spatial resolution, a grid-based approach was adopted. The volume was split into multiple tiles, each of which was reduced to a 1D signal along the stitching axis by averaging over its perpendicular plane.

As illustrated in Figure 28, the grid is defined in the plane perpendicular to the direction of translation. It should be noted that Figure 28 presents only a subset of the full grid and is intended solely as a

schematic demonstration of how individual tiles operate on an example volume.





Figure 28: Illustration of the grid-based stitching approach.

Each tile produces its own signal and cross-correlation is computed across the entire overlap in each tile. For example, if there was a 4x4 grid, and the overlap being tested is 10 voxels, there would be 160 cross-correlation calculations for this one step. For a 200x200x400 image, the algorithm must check 400 potential shifts (± 200) and 800 pixels in each plane. This showcases the need for a reasonable grid size. Stitching images of these dimensions took between 5.4 and 34.5 seconds for 6x6 and 100x100 grids respectively. A 15x15 grid was used throughout this project. With a stitching time of 6.4 seconds, it was one of the faster grids but it also preserved spatial resolution. The NCC score is computed with Equation 9 below.

$$C_{\text{norm},r,c}(k) = \frac{\sum_x p_{1,r,c}(x) p_{2,r,c}(x+k)}{\sqrt{\sum_x p_{1,r,c}(x)^2 \sum_x p_{2,r,c}(x+k)^2}} \quad (9)$$

where $p_{1,r,c}(x)$ and $p_{2,r,c}(x)$ represent the one-dimensional signals extracted from the corresponding tile at position (r, c) in volumes 1 and 2, respectively, and x indexes the coordinate along the stitching axis. By implementing Equation 9, the signals from corresponding tiles in each volume should match each other when the correct shift is chosen. Figure 29 shows the signals from corresponding tiles that voted for the chosen shift, the resulting NCC curve and a 2D slice of the reconstructed image where the white spots are where the two images' hotspots overlap.



(a) Example signals extracted from corresponding tiles in two overlapping volumes. (b) Example normalised cross-correlation curve. (c) Example stitch overlay.

Figure 29: Intensity graph overlap, normalised cross-correlation result and resulting stitch.

Figure 29a shows that while the signals are not perfect matches, they share the same general shape.

8.2 Algorithm Pipeline

The volumes first undergo preprocessing to reduce systematic correlation. As described in subsection 7.3, unreliable regions, such as the top of the volume, which exhibit a strong and consistent signal from the array's initial signals being reflected are cropped out. These strong, consistent signals are illustrated in Figure 30 which shows adjacent 3D images at $z=0$ and the percentage difference between them. Some percentage differences are large because it compares every pixel. This means if a pixel is white when it was black in the other image, a large percentage difference will be found. The mean percentage difference between these images is 11% but the near identical shape shows how the signals are very similar.




Figure 30: Comparing the top slice between 3D TFM images.

The sides of the volumes are also omitted from the cross-correlation calculation as the score can often increase towards the sides of the volumes which can be seen in Figure 29b. This is likely due to the reduced number of voxel planes in these regions, which limits the averaging of random intensity variations and can lead to artificially elevated correlations. The volumes are then divided into the grid in the plane perpendicular to the axis of translation e.g. y - z plane when translation is in x . The 1D signals are then extracted by averaging the pixel intensities in each tile in the grid. This 1D signal is then normalised to zero mean and unit variance and NCC is computed for each tile across the candidate shifts.

The NCC scores are then summed for each candidate shift and the shift with the highest score is chosen as the global shift to reconstruct the volume. Finally, the images are blended together across the overlap so there aren't any seams. Blending is carried out with intensity ramps which transfer intensity values linearly from a 1:0 ratio at one side of the overlap to a 0:1 ratio at the other side of the overlap.

8.3 Stitching in Two Dimensions

As the project required scanning a complete material sample, the stitching algorithm was extended to handle multiple images across both axes, introducing the challenge of overlap in two dimensions. As shown in Figure 31, scans 1&2 and 3&4 overlap purely along the x -axis, while scans 1&4 and 2&3 overlap along the y axis. However, all four scans also overlap in the centre, so each stitch could be influenced by non-adjacent images. To avoid this, stitching was restricted to immediately adjacent scans. When two images were stitched, any subsequent image was aligned only with its direct

neighbour, rather than with an already-stitched composite.

The full volume was then assembled after all pairwise shifts were calculated. Shifts along the x-axis were computed first to form rows. These rows were then aligned by estimating a y-axis shift, which was assumed to be globally consistent for each pair of rows. This shift was calculated using the final image of the previous row and the first image of the next row, as this was the only pair whose relative displacement was purely in the y-axis. Other inter-row pairs couldn't be assumed to align purely in y, due to inaccuracies in the gantry positioning and therefore weren't used for a y-axis stitch estimation.

To ensure no seams appear after stitching, blending was carried out on each row of images independently and then the rows were blended together using the same technique as outlined in subsection 8.2






Figure 31: 2x2 scan of a block.

The algorithm can stitch any number of images together. For example, the stainless steel block shown in Figure 32a was reconstructed into Figure 32b from 9 scans.



(a) Stainless steel block.



(b) Reconstructed block.

Figure 32: 9 scans reconstructing a stainless steel block.

9 Synthetic Data Generation (TT)

9.1 Motivation

The purpose of the synthetic pipeline is to generate images that are statistically representative of experimental TFM scans with similar intensity distributions, speckle structures, and correlation behaviours as the experimental scans. The validation target is therefore that synthetic images produce an NCC landscape of comparable shape and peak sharpness to the experimental one, meaning the stitching algorithm faces a comparable difficulty in both cases.

To validate the accuracy of the method, a known ground truth is required to compare the stitched volume and measure the predicted shift to the real shift. Experimental data alone cannot provide this. Although the gantry programmed a specific displacement between scans, mechanical backlash and positioning error meant the true physical displacement was not known to sufficient precision to serve as a reliable reference. Furthermore, the relationship between physical displacement and voxel shift depended on the pixel size calibration, introducing an additional source of uncertainty. As a result, experimental data could only be used to assess the consistency of the stitching algorithm through bootstrap testing, as described in Section 10. Consistency, however, does not imply correctness as a stitching algorithm could converge reliably to a systematically incorrect shift if, for example, a dominant feature in the overlap region biases the NCC landscape. Synthetic data with a known ground truth displacement is therefore the only method of directly measuring stitching accuracy.

Characterising algorithm performance across a range of overlaps, noise conditions, and feature densities would be impractical to achieve experimentally, as each test case would require a full physical scan. A synthetic pipeline allowed a large number of test cases to be generated automatically, enabling more thorough and systematic validation.

9.2 Initial Method

Initially, the volume was generated by filling a 3D array with uniformly distributed random values scaled to a low intensity range (0.05-0.15). A Gaussian filter with a fixed standard deviation was then applied to spatially correlate neighbouring voxels and simulate the grain structure of the sample. Defects could then be inserted as geometric shapes (cylinders, spheres...) at higher intensities (0.85-0.95). This treated each defect as a fixed object with the same appearance from every viewing angle, whereas in real TFM data the reconstructed shape of a defect depends on the angle from which the array images it. As shown in Figure 33, this produced a uniform correlation length throughout the volume, unlike the spatially varying speckle of experimental TFM backscatter [43], and yielded NCC peaks that were sharper and more well-defined than those observed experimentally. The stitching algorithm therefore performed artificially well on this data, motivating a physics-based forward model capable of reproducing the underlying scattering behaviour rather than mimicking its appearance.

9.3 Born Forward Model

When an ultrasonic wave propagates through a polycrystalline metal, it encounters a medium whose acoustic impedance is not uniform. Each grain has a slightly different crystallographic orientation, and this translates into a small spatial variation in density and elastic stiffness, and therefore in acoustic impedance $Z = \rho c$. The backscattered signal recorded by the array is the sum of weak echoes produced wherever the impedance changes, and in the case of pure metals with no defects, this occurs at grain boundaries. Reconstructing this signal from first principles requires solving the acoustic wave equation in a medium with spatially varying wave speed. After transformation into the frequency domain, this reduces to the Helmholtz equation [44],

$$(\nabla^2 + k_0^2) p(\mathbf{r}) = -k_0^2 \gamma(\mathbf{r}) p(\mathbf{r}) \quad (10)$$

where $p(\mathbf{r})$ is the total acoustic pressure, k_0 is the wavenumber in the background medium, and $\gamma(\mathbf{r})$ is a dimensionless contrast function describing local deviations from that background. Equation 10 is an




Figure 33: (a) Initial Gaussian-smoothed synthetic volume, (b) experimental TFM data, and (c) normalised NCC curves for the synthetic and experimental data.

implicit equation, and so solving this is analytically intractable for anything beyond isolated scatterers. Numerically, this requires finite element methods which is computationally expensive, especially for the large volume of 3D datasets the validation pipeline requires to statistically validate the stitching algorithm.

9.3.1 Born Approximation

The Born approximation circumvents this difficulty by exploiting the fact that in a polycrystalline metal the impedance variation between the background medium and any individual scatterer is small. If $\gamma(\mathbf{r})$ is small, the total field inside each scatterer differs only slightly from the incident field $p_i(\mathbf{r})$ that would exist in the unperturbed medium. Replacing $p(\mathbf{r})$ with $p_i(\mathbf{r})$ on the right-hand side of Equation 10 decouples incident and scattered contributions and linearises the problem: each scatterer now radiates independently in response to the known incident field, with no re-scattering. The scattered pressure at an observation point $\mathbf{r}$ can then be written as a volume integral over the contrast distribution,

$$p_s(\mathbf{r}) = k_0^2 \int G(\mathbf{r}, \mathbf{r}') \gamma(\mathbf{r}') p_i(\mathbf{r}') d^3\mathbf{r}' \quad (11)$$

where $G(\mathbf{r}, \mathbf{r}')$ is the free-space Green's function describing propagation from a point source at $\mathbf{r}'$ to the receiver at $\mathbf{r}$. For the weak-impedance-fluctuation regime relevant to polycrystalline grain scattering, the contrast function reduces to the impedance perturbation normalised by the background impedance,

$$\gamma(\mathbf{r}) \approx \frac{\delta Z(\mathbf{r})}{Z_0} \quad (12)$$

where Z_0 is the background acoustic impedance and $\delta Z(\mathbf{r})$ is its local perturbation. Specialising Equation 11 to a compact scatterer of volume ΔV over which γ is approximately constant reduces the volume integral to a point-scatterer amplitude a proportional to impedance contrast [45]. Adopting the weak-contrast reflection coefficient $R \approx \delta Z/(2Z_0)$ [44] as the amplitude normalisation gives

$$a \propto \frac{\delta Z}{2Z_0} \cdot \Delta V. \quad (13)$$

Every scatterer in the simulated volume contributes an echo whose amplitude is proportional to the local impedance jump and whose arrival time is determined by the geometric time of flight. The validity of this formulation rests on three assumptions [44]:

- The impedance contrast must be weak, $|\delta Z/Z_0| \ll 1$, so that the replacement of $p(\mathbf{r})$ by $p_i(\mathbf{r})$ introduces only a small error. In polycrystalline aluminium and copper, grain-to-grain impedance fluctuations are of order 1–3 %.
- Single scattering must dominate, requiring that the acoustic mean free path exceed the characteristic scatterer spacing, so that the neglected multiple-scattering terms remain negligible.
- The phase accumulated across a single scatterer must be small, $k_0 \cdot |\delta c/c| \cdot L \ll 1$, so that the incident field is approximately uniform over its volume and the integrand in Equation 11 can be treated as constant [46].

These conditions are satisfied for the microstructural backscatter regime that the stitching pipeline is designed to validate.

9.3.2 Voronoi Grain Structure

Polycrystalline metals consist of discrete grains separated by sharp boundaries, across which acoustic impedance changes abruptly due to differences in crystallographic orientation between neighbouring grains [47]. In Equation 13, these boundaries are the sites where $\delta Z \neq 0$, and they therefore constitute the population of weak scatterers responsible for the diffuse backscatter observed in experimental TFM images. To reproduce this in simulation, the volume is constructed using a Voronoi tessellation, which partitions 3D space into discrete cells by assigning each voxel to its nearest randomly seeded nucleus [48]. Each cell represents a single grain, and the sharp interfaces between cells represent grain boundaries. This approach is well established in the computational ultrasonics literature as a statistically representative model of polycrystalline microstructure and has been shown to reproduce ultrasonic scattering characteristics comparable to those measured experimentally [49]. Each grain is assigned an acoustic impedance drawn from a narrow Gaussian distribution about the background value as shown in Equation 14,

$$Z_i = Z_0(1 + \varepsilon_i), \quad \varepsilon_i \sim \mathcal{N}(0, \sigma^2) \quad (14)$$

with $\sigma \approx 0.03$, which represents the weak grain impedance contrast typical of metals such as copper [47], and sits within the weak-contrast regime required by the Born approximation. The mean grain size was set to 0.5 mm, within the $50\mu\text{m}$ to 2mm range reported for copper components [50]. At 7.5 MHz in copper, the wavelength is $\lambda \approx 0.6\text{mm}$, so the grain size is comparable to the wavelength. This produces the diffuse speckle pattern characteristic of coarse-grained polycrystalline backscatter

[51], and matches the experimental TFM images more closely than the Gaussian-smoothed approach (Figure 34).




Figure 34: Voronoi-based grain model (right), and Gaussian-smoothed volume (left)

9.3.3 Scatterer Extraction

The Voronoi tessellation defines the impedance field, but Equation 13 requires a discrete list of scatterers with associated amplitudes before the forward model can evaluate the scattered pressure. Since scattering occurs only where the impedance changes, the signed impedance gradient is computed along each of the three voxel-grid axes; wherever this gradient is non-zero, a scatterer is emitted carrying the magnitude and sign of that gradient. The corresponding Born amplitude follows directly from Equation 13,

$$\delta = \frac{\Delta Z}{2Z_0} \quad (15)$$

where δ is the normalised Born amplitude of a single voxel scatterer and ΔZ is the signed impedance difference across the voxel along the relevant axis. To keep the scatterer count manageable and reduce the computational cost without loss of physical content, voxels whose Born amplitude falls below a small threshold are removed from the scatterer list, since their contribution to the reconstructed image is negligible.

9.3.4 Waveform Synthesis and Elevation Directivity

The forward model evaluated the signal recorded at each transmit–receive element pair by summing the contribution of every scatterer in the volume. For a given transmitter–scatterer–receiver group, the arrival time of the scattered echo was computed geometrically from the sum of the transmit-leg and receive-leg propagation distances divided by the wave speed. The amplitude at the receiver is the scatterer’s Born amplitude multiplied by four factors: a spherical spreading term $1/r$ on each leg, an exponential material attenuation $e^{-\alpha r_{\text{tot}}}$ over the total round-trip path $r_{\text{tot}} = r_{ts} + r_{rs}$, and element directivity factors applied at both the transmitting and receiving elements [52]. Real array elements are not mathematical points but rectangular apertures of finite width in both the array-axis and elevation directions, so each element radiates and receives with an angle-dependent sensitivity

rather than isotropically [44]. This finite extent is captured by modelling the element directivity as a separable product of two sinc functions, one over each aperture dimension, so that scatterers lying off the element’s central axis are appropriately attenuated.

The collection of amplitude–arrival-time pairs for a given element pair forms an impulse train, which is convolved with a Gabor wavelet matched to the transducer’s centre frequency and fractional bandwidth to produce the physical waveform. The pulse-echo wavelet is used because it approximates the pulse shapes produced by real piezoelectric transducers [53]. To keep the computational cost low, the scatter-addition step is implemented as a vectorised accumulation into a two-dimensional element-by-time array, and the convolution is batched across all element pairs. The final volumes were then processed through the same TFM algorithm defined in Section 7.3.

9.3.5 Model Limitations

The Born forward model described above captures only single-scattered longitudinal echoes from grain-boundary impedance contrasts. Several other contributions present in experimental FMC data are therefore absent from the synthetic volumes. Near-field artefacts from the array surface and coupling layer as well as front-wall and back-wall reflections were omitted from the model. In the experimental pipeline, both were removed by the depth-cropping stage described in Section 7.3, so their absence was consistent with the processed experimental input.

Mode conversion between longitudinal and shear waves, and Kirchhoff scattering from hard defect surfaces were omitted from the model. Mode-converted contributions are an order of magnitude weaker than direct longitudinal grain scattering for the back-scattering geometry used here, so they were excluded [46]. Defect-surface Kirchhoff scattering was also excluded because the stitching validation targets the microstructural backscatter, where defect echoes would produce strong, well-localised NCC peaks that make the stitching task easier rather than harder, and so their absence does not weaken the validation. These assumptions were made to reduce computational cost without reducing validation fidelity, which is assessed in Section 9.4.

9.4 Validation of Synthetic Data

9.4.1 Pixel Intensity Distribution

To validate the quality and accuracy of the synthetic data, it was compared with experimental data. Visual inspection was insufficient to quantitatively assess the validity of the volumes, therefore the pixel intensity distributions of the images were analysed.

The distribution of pixel intensities provided a check that the bulk statistics of the images agreed. Figure 35 shows the normalised pixel intensity distributions for a representative copper volume pair. The two distributions exhibit a heavy-tailed distribution with agreement quantified by a Kolmogorov–Smirnov (KS) statistic of 0.018 (averaged from 59 experimental and synthetic distributions) [54]. Agreement at this level indicates that the synthetic and experimental volumes have comparable overall brightness statistics and dynamic range. This statistic is not on its own sufficient as two images with identical intensity distributions can have entirely different spatial structure.

The 2D spatial autocorrelation function (ACF) of a TFM slice characterises the length scale over which the speckle pattern decorrelates. This correlation length is dependent on the underlying microstruc-




Figure 35: Intensity distribution of the experimental and synthetic data (KS $D = 0.018$) (left), and CDF comparison (right).

tural scattering physics as it reflected the combined effect of grain size, ultrasonic wavelength, and the imaging point-spread function [55]. The synthetic pipeline must reproduce it if it is to represent experimental data faithfully. Therefore, the ACF combined with the pixel intensity distributions gave a strong indication of the validity of the synthetic data.

For each volume, the 1D ACF was computed independently along each of the three axes by averaging across the two transverse axes, and then normalised so that $R(0) = 1$. The mean ACF and $\pm 1\sigma$ band were then taken across $n = 56$ synthetic volumes and $n = 56$ experimental volumes drawn from separate copper scans. The $1/e$ crossing of each mean ACF was used as the summary statistic for correlation length.

The results are shown in Figure 36. The two datasets exhibit the same qualitative ACF shape on all three axes, and the $\pm 1\sigma$ uncertainty bands overlap throughout the lag range. The $1/e$ correlation lengths (synthetic / experimental) are $0.7 / 0.5$ mm on z , $0.8 / 0.6$ mm on y , and $0.8 / 0.8$ mm on x . These values are comparable to the ultrasonic wavelength in copper at 7.5 MHz ($\lambda \approx 0.6$ mm). The synthetic correlation lengths are therefore equal to or longer than the experimental on every axis, with the closest agreement on the x axis where the array provides the densest sampling and the largest disagreement on the depth (z) axis where the correlation length is set by the pulse envelope. This is consistent with the simulated Gabor pulse having a slightly longer envelope than the experimental transducer response. While the absolute correlation lengths do not match exactly, the shapes and decay rates of the ACFs are close enough that the speckle in the synthetic volumes decorrelates over comparable spatial scales to the experimental data, which is the property the stitching algorithm depends on.




Figure 36: Mean 1D autocorrelation functions of synthetic and experimental copper TFM volumes, with $\pm 1\sigma$ bands: (left) z -axis, (middle) y -axis, and (right) x -axis. The dashed line marks the $1/e$ threshold used to define correlation length.

9.4.2 NCC Curves

The intensity distribution and autocorrelation comparisons together establish that the synthetic volumes reproduce both the brightness statistics and the spatial structure of experimental TFM speckle. The final validation target is the NCC landscape itself, since this is the quantity the stitching algorithm acts on directly. Figure 37 compares representative NCC curves computed from a synthetic stitch and from an experimental stitch on the copper sample, across the same lateral shift range.

The two curves are qualitatively similar. Both feature a single dominant peak at the same shift with comparable peak width, and both show similar rising shoulders at the extremes of the search range arising from edge effects in the cross-correlation. The off-peak structure is also consistent between the two cases, with some fine-scale variation in the negative-shift region but a closely matched overall shape. Because the curves are individually normalised to a peak amplitude of one, this is a comparison of shape rather than absolute correlation magnitude. This does not affect the stitching algorithm, which identifies shifts from peak position relative to baseline rather than from absolute NCC.

This is a substantial improvement over the Gaussian-smoothed initial method, whose NCC peak was significantly narrower than experimental and sat against a near-zero off-peak baseline (Figure 33). The Born forward model therefore reproduces the NCC landscape that the stitching algorithm encounters experimentally, which was the validation target stated in section 9.

The principal limitation of this comparison is the small number of experimental stitches available as only four valid pairs were collected during the final experiment, which is sufficient to support a qualitative shape comparison but not to make distributional statements about peak sharpness, baseline noise, or peak-prominence variability. A larger experimental dataset would be needed to test whether the synthetic NCC curve matches experimental data in distribution as well as in shape.




Figure 37: Representative NCC curves from a synthetic and an experimental stitch on copper.

10 Stitching Validation (OB)

10.1 Bootstrap Tests

Bootstrap testing was used to evaluate the robustness of the stitching algorithm. Here, robustness refers to how sensitive the estimated shift is to variations in the input data, particularly given that the volumetric images contain random grain structures and localised hotspots that vary spatially.

First, the stitching algorithm was applied to the full volumes to obtain a baseline shift estimate,

which serves as a reference for assessing variability. Each volume was then divided along the z-axis into a set of 3D bands. New volume pairs were constructed by randomly sampling these bands with replacement, effectively generating alternative versions of the original data. This resampling process produces datasets in which some regions may be over-represented, under-represented, or omitted entirely.

The stitching algorithm was then applied to each resampled volume pair. This process was repeated 1000 times, producing a distribution of estimated shifts. A narrow distribution indicates that the baseline shift estimate is stable and well-defined, whereas a wide distribution suggests greater uncertainty, indicating that the algorithm does not converge to a single consistent solution for the given data. Two examples of this bootstrap analysis are shown in Figure 38.




Figure 38: Comparison of bootstrap tests using experimental and synthetic data.

In both cases, the distributions were centred on the reference shift, indicating that the stitching algorithm was generally robust. However, a clear difference in variability is seen between the two datasets. The synthetic dataset in Figure 38b produces an almost perfectly centred distribution, with 99% of bootstrap samples returning the same shift. This indicates extremely high stability and suggests that the signal features driving the correlation are consistent and distributed throughout the volume.

In contrast, the experimental dataset in Figure 38a exhibits a broader distribution, with the majority of bootstrap estimates still clustered tightly around the original shift, but with a noticeable spread of 2 voxels, or 0.08mm. The increased variability reflects the increased difficulty of stitching experimental data, where the shift-deciding features may be less uniformly spread over the z-axis and warping. As a result, resampling different bands can alter the dominant correlation signal, leading to changes in the estimated shift. Despite this variation, >95% of estimates were within 1 pixel of the original value which shows that the algorithm remained reliable. The limited spread suggests that while local structural differences influenced the correlation, they did not seriously compromise the overall alignment result.

10.2 Certainty Metrics

To validate the stitching algorithm, 580 synthetic volumes or 290 volume pairs were created over 42 overlaps. The cross-correlation graphs show that, while the same overlaps often produce similar peak structures and values, there is variation in height and shape within the same overlaps as seen in Figure 39. Another source of variation is in the baseline noise levels. Some curves have well-defined peaks while others display broader peaks or noise levels that trump the true peak. This behaviour is

illustrated in Figure 39. For an 80% overlap, the cross-correlation curve exhibits a dominant, narrow peak, indicating a high-certainty shift estimate. In contrast, a 40% overlap produces a broader and less distinct which increases the chances of an incorrect shift.




Figure 39: Comparison of cross-correlation curves for large and small overlaps.

This variability is reflected directly in the stitching accuracy. The predicted shifts deviate increasingly from the expected values as the overlap decreases as shown in Figure 40. Correct predictions (orange) are associated with smaller shifts (larger overlap) while errors (green to purple) occur increasingly as the shift increases.




Figure 40: Expected vs predicted shift with error markings.

Overall, the trend of error increasing with shift size is expected, but the outliers around the 100 pixel shift mark (seen in opaque pink at the top) show that there is a problem with the synthetic data. No physical explanation was found. Therefore, these data points were removed from the analysis. As discussed in subsection 9.4, the synthetic data exhibits similar cross-correlation behaviour and intensity distributions to the experimental measurements. However, it is important to note that the two datasets are not identical. In particular, experimental data is subject to additional sources of uncertainty, including warping towards the edges of the images and positioning uncertainty from the gantry. Consequently, even if the stitching algorithm demonstrates perfect stitching accuracy for synthetic data with overlaps down to 85%, this level of performance can't be directly assumed for experimental data. However, as a guideline for experimental data collection, targeted overlaps should not be smaller than 70% as it was found in the synthetic data that this is a threshold below which error starts to increase. However, this recommendation should be treated with caution, as further validation

using a larger set of experimental data would be required to confirm its general applicability.

As a result, it was insufficient to treat the output of the stitching algorithm as a deterministic result. Instead, it was necessary to quantify the certainty associated with each predicted shift. To quantify this, a set of certainty metrics was developed to characterise each correlation curve. These metrics describe key properties of the dominant peak in the cross-correlation curve and were designed to capture how reliably a shift can be identified.

The peak-to-side-lobe ratio (PSR) measures how strongly the dominant peak stands out relative to the surrounding baseline. A high PSR indicates that the peak is well separated from noise and competing structures, suggesting a reliable match. Conversely, a low PSR indicates that the peak is not clearly distinguishable, increasing the likelihood of incorrect shift selection. The peak width, measured at both 50% and 95% of the peak height relative to the baseline, quantifies the sharpness of the correlation maximum. Narrow peaks correspond to well-constrained and precise shift estimates, whereas broader peaks indicate ambiguity, where a range of shifts produce similar correlation values. Furthermore, the separation between the highest and second-highest peaks (peak ratio) captured the level of ambiguity in the solution. If multiple peaks have similar amplitudes, the algorithm may select an incorrect shift. A large separation between the dominant and secondary peaks therefore indicates a more reliable result.

Individual metrics only capture part of the underlying behaviour. While the PSR exhibited a general trend in which higher values corresponded to improved accuracy, other metrics such as peak width and peak ratio showed variability when considered independently. Combining these features allows more information to be incorporated into the prediction, resulting in a two-dimensional probability model that provides a more consistent prediction.

A more robust characterisation was obtained by combining features. Two-dimensional certainty maps revealed distinct regions in feature space where the probability of correct stitching was high, and others where it was significantly lower. In particular, combinations of high PSR and moderate peak width consistently corresponded to high-confidence stitching results as seen in Figure 41.




Figure 41: Shift certainty: PSR vs width at 95% of peak height

Using these empirical relationships, stitching could be treated as a probabilistic process. For a given

set of features $\mathbf{f}$ extracted from a correlation curve, a confidence value was assigned:

$$\text{certainty} = P(\text{correct} \mid \mathbf{f}) \quad (16)$$

This represented the probability that the predicted shift lay within a specified tolerance (e.g. ± 1 pixel) of the true value.

Although these certainty metrics were derived from synthetic data, their primary purpose was to enable interpretation of experimental stitching results where ground truth is unavailable. For each experimental stitch, the same features could be computed and mapped onto the empirical probability distributions obtained from synthetic data. This allows a confidence value to be assigned to each predicted shift, which allows high confidence stitches to be accepted and low confidence stitches to be further constrained. It should be noted that data generation was time constrained so the certainty metrics are not based off large quantities of data. In an ideal world, these certainty metrics would be based on more data from more robust synthetic generation that accurately captures the physics and nuance of the experimental system.

10.3 Angular Misalignment Tests

The gantry introduces positioning uncertainty in rotation as well as in translation. As the gantry is expected to be used with radon reconstruction or rotational imaging with 2D arrays to improve signal fidelity, it is likely that an angular misalignment could be introduced between scans. As the stitching algorithm was designed purely for translation, a robustness test was carried out to ensure that it would be able to handle angular misalignments. Tests were carried out with synthetic data to see at what angular misalignment the stitching algorithm would start to find incorrect shifts. Angles from 0.5° to 16° were tested on overlaps of 70% to 90%. The peak cross correlation scores and the certainty of results against the injected misalignments are shown below in Figure 42




Figure 42: Peak cross correlation and certainty plots vs injected misalignment.

Misalignment had no effect on the stitch error which is why it hasn't been shown here. This was not expected as angles are known to introduce error in NCC methods [56]. With the absence of error, other metrics were required to quantify the effect of misalignment on stitching. Peak cross-correlation scores should drop as the misalignment reduces the likeness between corresponding tiles and certainty

metrics such as PSR vs peak width at 95% peak height were expected to drop as well. However, from Figure 42, it can be seen that neither of the expected outcomes occurred, with peak cross-correlation scores and stitching certainty increasing marginally in two of the three tests in both plots. Overall, the distributions are flat which is unexpected but less surprising than increasing certainty. Increasing cross-correlation scores could have been caused by random correlation from adjacent hotspots but this is unlikely as there tends to be low-intensity space between hotspots which would reduce the cross-correlation score. Increasing or constant certainty in the results is surprising as well, as introducing misalignment should reduce correctly overlapping areas, therefore reducing peak height compared to noise levels and PSR in turn. The most likely reason for this strong performance is that there is a limitation of the synthetic data and further work is required to produce more realistic data where misalignment affects stitching results. However, as there is no experimental data to compare these results with, no conclusion can be drawn from these results currently.

11 Final Experimental Procedure and Results (OD)

To evaluate the capabilities of the current prototype, an experimental procedure was conducted using the end-to-end system. Two experiments were conducted to image a pure copper sample using a 2D ultrasonic array with 128 elements at 7.5MHz. The experimental setup is shown in Figure 43.




Figure 43: Final experimental setup.

Due to the issues explained in section 6, the experiment could not benefit from the rotations discussed in subsection 7.5 and so remained a purely translational test. Both experiments translated the array both in the X and Y directions. The two experimental paths taken by the array are shown in Figure 44a.

In experiment 1, the gantry was not calibrated prior to the scanning sequence. The gantry was then calibrated between experiments following the procedure shown in subsection 5.3 and this took 23 minutes. The scanning sequences took 15 minutes for experiment 1, and 10 minutes for experiment 2. The scans were processed and imaged as outlined in section 7. Processing the raw data took 1 minute and 13 seconds, and imaging the 12 volumes took 3 minutes and 51 seconds. This gave a total end-to-end process of 2 minutes and 55 seconds per volume in experiment 1, and 5 minutes and 54 seconds per volume in experiment 2. This remained within the time limits set out by the PDS.

As the array diameter was 12mm, shifts of 2mm and 3mm corresponded to overlaps of 83% and 75%




Figure 44: Comparison between experimental and calculated array shifts.

respectively. The imaged volumes were then stitched and the calculated shifts are shown in Figure 44b. It can be seen that in both experiments, the translation along the y-axis does not align well with the expected distances shown in Figure 44a. This could be from the issues discussed in section 6, as well as the lack of rotations leading to low quality scans, which couldn't be stitched. However, both experiments showed satisfactory translational accuracy along the x-axis. The full stitched volume for experiment 2 is shown in Figure 45a.




Figure 45: Experiment 2 stitched volume.

As shown in Figure 45b, the cross-section shows that the volumetric imaging was successful for experiment 2 as the backscattering signals were successfully stitched, especially at the centre of the image. The seams present could be removed via weighted average blending as discussed in section 8, but some areas have large differences across the seams. This suggests that the experimental procedure could still be greatly improved, such as ensuring that rotation is used to gather more information at each scan location.

12 Conclusion (All)

In conclusion, the project delivered a functioning integrated platform that combined mechanical positioning, automated scan control, preprocessing, imaging, stitching, and validation into a single work-

flow. The automated gantry, provided by Build Team, addressed most of its Product Design Specifications. The platform was designed around a 580mm x 750mm footprint using aluminium extrusion and the array holding mechanism was able to support both 1D and 2D arrays. During operation, the gantry maintained its structural stability without excessive vibration which would otherwise hinder its ability to collect FMC data. The X-axis slider, which is the main components when translating the gantry beam in X axes underwent iteration processes before it were successfully incorporated into the full assembly. The design could be improved by introducing guide rails, similar to the Y-axis slider to eliminate the wobbling during translation, and by making a longer slider to improve the stability. The Y-axis slider and the Z-axis array holding mechanism were successfully combined into a single moving subsystem. However, the added weight of this combined assembly caused a forward tilt during testing, which affected the array contact with the sample. Therefore, increasing the slider contact length and adding further guide support would solve this stability issue.

The automation strategy demonstrated that the mechanical motion and data acquisition could operate with limited manual intervention compared to traditional data collection. However, despite the calibration process producing acceptable translational accuracy of $\pm 1\text{mm}$ for the X-axis and $\pm 5\text{mm}$ for the Y-axis, the mechanical rigidity of the gantry remained insufficient for the rotational axis. The rotational axis was successfully integrated, but it could not achieve the required angular precision. This prevented any rotational scans from being performed autonomously. Future improvement should therefore focus on manufacturing the rotational block with tighter tolerances to increase stability. The rotational shaft should be supported with a bushing or bearing to improve rotational accuracy. Thus, addressing these structural limitations is a necessary step before the full automated scanning capabilities can be fully realised.

The signal processing, provided by Code Team, also succeeded in fulfilling many of its Product Design Specifications. The pipeline successfully accepts configurable parameters such as speed of sound, and the imaging workflow was shown to produce TFM reconstructions of both defects and microstructural backscatter. The speed of sound was found within 2.6% of accepted values, and the compression greatly reduced memory bandwidth while only deviating $2.57\text{e-}04\%$ in image accuracy.

A physics-based synthetic data pipeline was also developed to provide ground-truth data for stitching validation. The pipeline used a Born forward model with a Voronoi grain structure, replacing an initial Gaussian-smoothed model. Validation against experimental copper scans showed agreement on pixel intensity distributions to a Kolmogorov-Smirnov statistic of 0.018, $1/e$ autocorrelation lengths within 0.2 mm on every axis, and NCC curves matched in peak position, width, and off-peak structure. The synthetic data was crucial for the statistical analysis of stitching, identifying a 70% lower bound on overlap for reliable stitching and underpinning the certainty metrics that allow experimental stitches to be assigned a confidence value. The GPU-accelerated TFM framework was used for the simulation pipeline, achieving a 120-times speed-up over multi-threaded C++ on the CPU.

Future improvements could extend the forward model to include mode conversion, Kirchhoff scattering from defect surfaces, and wall reflections, which would allow the pipeline to validate defect detection in addition to microstructural stitching. A more realistic misalignment response is also needed, since the current synthetic data does not reproduce the expected sensitivity of NCC scores to angular misalignment, limiting the conclusions that could be drawn from the robustness tests. Implementing end-to-end hardware acceleration from synthetic data generation to stitching could also accelerate the optimisation process, and provide larger volumes of data.

Overall, the system was successful in delivering the core architecture of an automated ultrasonic microstructural imaging system. It integrated an automated scanning gantry with a configurable imaging pipeline, and demonstrated that microstructural backscatter could be reconstructed and stitched. The main improvements required are mechanical refinement of the gantry, more complete experimental repeatability studies, and further acceleration of the data pipelines. The platform has a clear potential to become a reliable tool for large-area ultrasonic inspection and future 3D microstructural characterisation.

References

- [1] S. Iliopoulos et al., “Detection and evaluation of cracks in the concrete buffer of the Belgian Nuclear Waste container using combined NDT techniques,” *Construction and Building Materials*, vol. 78, pp. 369–378, Mar. 2015. Accessed: Feb. 15, 2026. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0950061814013385>.
- [2] Y. Qin et al., “An Improved Phased Array Ultrasonic Testing Technique for Thick-Wall Polyethylene Pipe Used in Nuclear Power Plant,” *Journal of Pressure Vessel Technology*, vol. 141, no. 041403, May 2019. Accessed: Feb. 15, 2026. [Online]. Available: <https://doi.org/10.1115/1.4043384>.
- [3] S.-J. Song et al., “Development of an ultra sonic phased array system for nondestructive tests of nuclear power plant components,” *Nuclear Engineering and Design*, vol. 214, no. 1, pp. 151–161, May 2002. Accessed: Feb. 15, 2026. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0029549302000249>.
- [4] M. E. Bajgholi et al., “Reliability assessment of non-destructive testing (NDT) for the inspection of weld joints in the hydroelectric turbine industry,” en, *The International Journal of Advanced Manufacturing Technology*, vol. 128, no. 9, pp. 4223–4233, Oct. 2023. Accessed: Feb. 15, 2026. [Online]. Available: <https://doi.org/10.1007/s00170-023-12176-5>.
- [5] D. Chauveau, “Review of NDT and process monitoring techniques usable to produce high-quality parts by welding or additive manufacturing,” en, *Welding in the World*, vol. 62, no. 5, pp. 1097–1118, Sep. 2018. Accessed: Feb. 15, 2026. [Online]. Available: <https://doi.org/10.1007/s40194-018-0609-3>.
- [6] V. P. Bhatia et al., “History of Ultrasound,” en, in *Practical Urological Ultrasound*, P. F. Fulgham et al., Eds., Cham: Springer International Publishing, 2021, pp. 1–11. Accessed: Feb. 16, 2026. [Online]. Available: https://doi.org/10.1007/978-3-030-52309-1_1.
- [7] T. Y. Park et al., “Method to optimize the placement of a single-element transducer for transcranial focused ultrasound,” *Computer Methods and Programs in Biomedicine*, vol. 179, p. 104982, Oct. 2019. Accessed: Feb. 16, 2026. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0169260719301506>.
- [8] H. Taheri et al., “Nondestructive Ultrasonic Inspection of Composite Materials: A Comparative Advantage of Phased Array Ultrasonic,” en, *Applied Sciences*, vol. 9, no. 8, Apr. 2019. Accessed: Feb. 15, 2026. [Online]. Available: <https://www.mdpi.com/2076-3417/9/8/1628>.
- [9] B. W. Drinkwater et al., “Ultrasonic arrays for non-destructive evaluation: A review,” *NDT & E International*, vol. 39, no. 7, pp. 525–541, Oct. 2006. Accessed: Jan. 20, 2026. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0963869506000272>.
- [10] C. Holmes et al., “Post-processing of the full matrix of ultrasonic transmit–receive array data for non-destructive evaluation,” *NDT & E International*, vol. 38, no. 8, pp. 701–711, Dec. 2005. Accessed: Jan. 20, 2026. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0963869505000721>.
- [11] E. Jasiuniene et al., “3D Ultrasonic Non-destructive Evaluation of Spot Welds Using an Enhanced Total Focusing Method,” en, *Journal of Materials Engineering and Performance*, vol. 24, no. 2, pp. 825–831, Feb. 2015. Accessed: Apr. 8, 2026. [Online]. Available: <https://doi.org/10.1007/s11665-014-1348-2>.
- [12] F. J. Humphreys, “Review Grain and subgrain characterisation by electron backscatter diffraction,” en, *Journal of Materials Science*, vol. 36, no. 16, pp. 3833–3854, Aug. 2001. Accessed: Apr. 17, 2026. [Online]. Available: <https://doi.org/10.1023/A:1017973432592>.
- [13] W. Wang et al., “Metallic material microstructure grain size measurements from backscattering signals in ultrasonic array data sets,” en, *NDT & E International*, vol. 149, p. 103251, Jan. 2025. Accessed: Jan. 27, 2026. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0963869524002160>.

- [14] P. Fulton, "Design and Evaluation of a Rotational Scanning System for Enhanced 3D Ultrasonic Imaging," en, 2025.
- [15] T. N. I. Ho, "Accelerated Ultrasonic Non-Destructive Evaluation with Total Focusing Method Image Stitching," en, 2025.
- [16] N. Hansen et al., "The strain and grain size dependence of the flow stress of copper," *Acta Metallurgica*, vol. 30, no. 2, pp. 411–417, Feb. 1982. Accessed: Apr. 17, 2026. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0001616082902218>.
- [17] Q. Huang et al., "A Review on Real-Time 3D Ultrasound Imaging Technology," en, *BioMed Research International*, vol. 2017, no. 1, p. 6 027 029, 2017, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1155/2017/6027029>. Accessed: Jan. 22, 2026. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1155/2017/6027029>.
- [18] R. W. Prager et al., "Three-dimensional ultrasound imaging," EN, *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, vol. 224, no. 2, pp. 193–223, Feb. 2010. Accessed: Jan. 22, 2026. [Online]. Available: <https://doi.org/10.1243/09544119JEIM586>.
- [19] *An Introductory Guide to Multi-Criteria Decision Analysis (MCDA) – Government Analysis Function*. Accessed: Mar. 11, 2026. [Online]. Available: <https://analysisfunction.civilservice.gov.uk/policy-store/an-introductory-guide-to-mcda/>.
- [20] D. Kyne, *Pairwise Comparison: Definition, Methods, Tools, Examples // OpinionX — Free Stack Ranking Surveys*, en-GB. Accessed: Mar. 11, 2026. [Online]. Available: <https://www.opinionx.co/blog/pairwise-comparison>.
- [21] J. L. LaGuardia, "COMPUTER AIDED PARAMETRIC SCREW DESIGN AND ANALYSIS USING Pro/ENGINEER SOLID MODELING SOFTWARE," en,
- [22] L. Sandanamsamy et al., "A comprehensive review on fused deposition modelling of polylactic acid," en, *Progress in Additive Manufacturing*, vol. 8, no. 5, pp. 775–799, Oct. 2023. Accessed: Apr. 26, 2026. [Online]. Available: <https://link.springer.com/10.1007/s40964-022-00356-w>.
- [23] G. Tian et al., "Experimental investigations of process parameters on dimensional accuracy for FDM printed CF/PETG parts," en, *Polymers and Polymer Composites*, vol. 33, p. 09 673 911 251 379 825, Sep. 2025. Accessed: Apr. 26, 2026. [Online]. Available: <https://journals.sagepub.com/doi/10.1177/09673911251379825>.
- [24] P. Bean et al., "Numerical Modeling and Experimental Investigation of Effective Elastic Properties of the 3D Printed Gyroid Infill," en, *Applied Sciences*, vol. 12, no. 4, p. 2180, Feb. 2022. Accessed: Apr. 23, 2026. [Online]. Available: <https://www.mdpi.com/2076-3417/12/4/2180>.
- [25] Ultimaker, *3D printing infill density: Optimizing strength and speed*, en, Sep. 2025. Accessed: Apr. 26, 2026. [Online]. Available: <https://ultimaker.com/learn/3d-printing-infill-density-optimizing-strength-and-speed/>.
- [26] K. Sathyanarayana et al., "Metallic lightweight materials: Properties and their applications," en, in *Lightweight and Sustainable Composite Materials*, Elsevier, 2023, pp. 47–67. Accessed: Apr. 20, 2026. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/B9780323951890000032>.
- [27] C. Design, *Circuit Designer Tutorials*, en. Accessed: Apr. 26, 2026. [Online]. Available: <https://docs.circuitdesigner.com>.
- [28] R. G. Budynas et al., *Shigley's Mechanical Engineering Design*, 10th ed. New York: McGraw-Hill Education, 2015.
- [29] T. E. Toolbox, "Solids and Metals - Speed of Sound," 2004. [Online]. Available: https://www.engineeringtoolbox.com/sound-speed-solids-d_713.html.
- [30] A. Simons et al., "Design and Development of a Delta 3D Printer Using Salvaged E-Waste Materials," en, *Journal of Engineering*, vol. 2019, pp. 1–9, Nov. 2019. Accessed: Apr. 26, 2026. [Online]. Available: <https://www.hindawi.com/journals/je/2019/5175323/>.
- [31] *What is G-code: Definition, Function, Types & Uses*, en-US, Section: CNC Machining, Jul. 2023. Accessed: Apr. 26, 2026. [Online]. Available: <https://www.3erp.com/blog/g-code/>.
- [32] C. Mineo et al., "Robotic path planning for non-destructive testing – A custom MATLAB toolbox approach," en, *Robotics and Computer-Integrated Manufacturing*, vol. 37, pp. 1–12, Feb. 2016. Accessed: Apr. 26, 2026. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0736584515000666>.
- [33] *Everything You Need to Know About Load Cells — RS*. Accessed: Apr. 26, 2026. [Online]. Available: <https://uk.rs-online.com/web/content/discovery/ideas-and-advice/load-cells-guide>.
- [34] O. Djuric et al., "Signal-processing-G2066," *GitHub Repository*, 2026, <https://github.com/bibi848/signal-processing-G2066>. [Online]. Available: <https://github.com/bibi848/signal-processing-G2066>.

- [35] N. Alagha et al., “Generalized raised-cosine filters,” *IEEE Transactions on Communications*, vol. 47, no. 7, pp. 989–997, Jul. 1999. Accessed: Apr. 11, 2026. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/774849>.
- [36] F. R. Kschischang, “The Hilbert Transform,” Oct. 2006.
- [37] L. Dagum et al., “OpenMP: An industry standard API for shared-memory programming,” *IEEE Computational Science and Engineering*, vol. 5, no. 1, pp. 46–55, Jan. 1998. Accessed: Apr. 16, 2026. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/660313>.
- [38] C. Gong et al., “Efficient Embarrassingly Parallel on Graphics Processor Unit,” in *2010 2nd International Conference on Education Technology and Computer*, ISSN: 2155-1812, vol. 4, Jun. 2010, pp. V4–400–V4–404. Accessed: Apr. 16, 2026. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/5529656>.
- [39] J. Driver, *Volumetric Imaging Using Ultrasound Arrays*, 2024.
- [40] A. Alba et al., “Phase Correlation Based Image Alignment with Subpixel Accuracy,” vol. 7629, Nov. 2012.
- [41] B. Zitová et al., “Image registration methods: A survey,” en, *Image and Vision Computing*, vol. 21, no. 11, pp. 977–1000, Oct. 2003. Accessed: Apr. 26, 2026. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0262885603001379>.
- [42] K. Aghajani et al., “A robust image registration method based on total variation regularization under complex illumination changes,” en, *Computer Methods and Programs in Biomedicine*, vol. 134, pp. 89–107, Oct. 2016. Accessed: Apr. 26, 2026. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0169260716301961>.
- [43] F. J. Margetan et al., “Backscattered microstructural noise in ultrasonic toneburst inspections,” en, *Journal of Nondestructive Evaluation*, vol. 13, no. 3, pp. 111–136, Sep. 1994. Accessed: Apr. 24, 2026. [Online]. Available: <https://doi.org/10.1007/BF00728250>.
- [44] L. W. Schmerr, *Fundamentals of Ultrasonic Nondestructive Evaluation: A Modeling Approach* (Springer Series in Measurement Science and Technology), en. Cham: Springer International Publishing, 2016. Accessed: Apr. 24, 2026. [Online]. Available: <http://link.springer.com/10.1007/978-3-319-30463-2>.
- [45] J. E. Gubernatis et al., “The Born approximation in the theory of the scattering of elastic waves by flaws,” *Journal of Applied Physics*, 1977. Accessed: Apr. 24, 2026. [Online]. Available: <https://sci-hub.su/10.1063/1.324142>.
- [46] J. Zhang et al., “Defect detection using ultrasonic arrays: The multi-mode total focusing method,” *NDT & E International*, vol. 43, no. 2, pp. 123–133, Mar. 2010. Accessed: Jan. 22, 2026. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0963869509001480>.
- [47] R. B. Thompson et al., “Scattering of elastic waves in simple and complex polycrystals,” en, *Wave Motion*, vol. 45, no. 5, pp. 655–674, Apr. 2008. Accessed: Apr. 24, 2026. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0165212507001023>.
- [48] R. Quey et al., “Large-scale 3D random polycrystals for the finite element method: Generation, meshing and remeshing,” en, *Computer Methods in Applied Mechanics and Engineering*, vol. 200, no. 17–20, pp. 1729–1745, Apr. 2011. Accessed: Apr. 24, 2026. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S004578251100003X>.
- [49] X. Zheng et al., “Modeling of Polycrystalline Material Microstructure with 3D Grain Boundary Based on Laguerre–Voronoi Tessellation,” en, *Materials*, vol. 15, no. 6, p. 1996, Jan. 2022. Accessed: Apr. 8, 2026. [Online]. Available: <https://www.mdpi.com/1996-1944/15/6/1996>.
- [50] J. Pitkänen, “The Effect of Grain Size on the Defect Detectability in Copper Components in Ultrasonic Testing,” en,
- [51] F. E. Stanke et al., “A unified theory for elastic wave propagation in polycrystalline materials,” *The Journal of the Acoustical Society of America*, 1984. Accessed: Apr. 24, 2026. [Online]. Available: <https://sci-hub.su/10.1121/1.390577>.
- [52] J. Krautkrämer et al., *Ultrasonic Testing of Materials*, en. Berlin, Heidelberg: Springer Berlin Heidelberg, 1990. Accessed: Apr. 24, 2026. [Online]. Available: <http://link.springer.com/10.1007/978-3-662-10680-8>.
- [53] R. Demirli et al., “Model-based estimation of ultrasonic echoes. Part I: Analysis and algorithms,” en, *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 48, no. 3, pp. 787–802, May 2001. Accessed: Apr. 24, 2026. [Online]. Available: <http://ieeexplore.ieee.org/document/920713/>.
- [54] F. J. Massey, “The Kolmogorov–Smirnov Test for Goodness of Fit,” en,
- [55] R. Wagner et al., “Statistics of Speckle in Ultrasound B-Scans,” *IEEE Transactions on Sonics and Ultrasonics*, vol. 30, no. 3, pp. 156–163, May 1983. Accessed: Apr. 24, 2026. [Online]. Available: <https://ieeexplore.ieee.org/document/1539480/>.

- [56] J. M. Almira et al., “Fast normalized cross-correlation for template matching with rotations,” en, *Journal of Applied Mathematics and Computing*, vol. 70, no. 5, pp. 4937–4969, Oct. 2024. Accessed: Apr. 27, 2026. [Online]. Available: <https://link.springer.com/10.1007/s12190-024-02157-6>.

A Appendix A

Supplementary figures and tables.




Figure 46: Time domain (left) and frequency domain (right) for filtering applied to defect detection in pure aluminium at 5MHz.




Figure 47: Comparison of TFM images for an aluminium sample using a 128 element, 10MHz array: (a) TFM image obtained from the unfiltered data, (b) Hilbert transform applied to the image, and (c) normalised image in decibel scale.




Figure 48: Copper scanning sweep using 3D printed jig for a 2D ultrasonic array.




Figure 49: CAD models of the smaller parts. (a) Motor housing for X-axis, (b) motor housing for Z-axis, and (c) bearing slot for X-axis.

B Appendix B

Product design specification.

Table 12: Supplementary product design specifications.

Group	Category	Requirement	Priority	Assessment
All	Performance	System must detect and resolve features at the expected scale (defect or microstructure backscattering).	MUST	Use samples with known feature sizes.
All	Performance	System SNR must be sufficient to observe microstructural backscatter.	MUST	Visually inspect scans performed on pure samples to confirm that material backscatter is present.
All	Performance	Repeat scans of the same region must produce consistent spatial alignment.	MUST	Perform repeated scans and measure variation in reconstructed feature positions.
All	Performance	Total scan + processing time must be within 30 minutes.	WANT	Measure time from scan start to final stitched output.
All	Performance	System should support repeated scanning without data loss or errors.	WANT	Run extended scans and verify no failed data.
All	Validation	System performance must be robust to small variations in coupling conditions.	MUST	Vary coupling layer thickness and “lift-off” to evaluate changes in image quality.
All	Validation	System output must correlate with experimental or simulated results.	MUST	Compare reconstructed images with synthetic data or known samples.
All	System	All acquired data must be stored without corruption or loss.	MUST	Validate dataset completeness after acquisition and check for missing entries.
All	System	Data must be traceable to acquisition parameters and scan location.	MUST	Verify metadata logging.
All	System	System must fail safely without corrupting data or hardware.	MUST	Interrupt scans mid-operation and verify safe shutdown and data integrity.
All	Constraints	System performance must not be significantly affected by environmental noise or vibration.	WANT	Operate system under varying environmental conditions and measure effect on signal quality.
All	Constraints	System must operate within laboratory constraints (space, power, safety).	MUST	Verify compliance with laboratory requirements.
Code Team	Performance	Imaging of 400x400 TFM image must be under 30 seconds for all personal computers in sub-group.	MUST	Timing functions used to measure execution time on all member’s computers.

Group	Category	Requirement	Priority	Assessment
Code Team	Performance	Filtering must preserve phase information required for TFM reconstruction	MUST	Compare TFM images with/without filtering and verify no spatial shift in defect positions.
Code Team	Performance	Radon transform algorithm to combine rotated 2D images into full 3D volumes.	MUST	Using radon transform on synthetic data to test accuracy prior to use on experimental data.
Code Team	Performance	Speed of sound calculations for pixel size calculations in TFM	MUST	Use of low frequency pulses (5MHz) on pure samples to calculate the time of flight between back wall reflections.
Code Team	Performance	Dynamic range of processed signals must be sufficient to capture weak backscatter from microstructure.	MUST	Confirm visibility of low-amplitude scattering in pure samples.
Code Team	Performance	TFM must avoid numerical artefacts such as aliasing.	MUST	Ensure pixel sizes are at most half of the wavelength of the ultrasound.
Code Team	Performance	Stitching must be robust to intensity variations between scans.	MUST	Apply artificial gain variations to images and verify that correct alignment still occurs.
Code Team	Validation	Development of a synthetic data generator which can replicate 2D TFM images.	MUST	Synthetic images compared to experimentally collected data.
Code Team	Validation	Synthetic data pipeline must produce statistically consistent outputs.	MUST	Run multiple realisations and compare distributions.
Code Team	Validation	Simulated backscatter must match experimental amplitude distributions	MUST	Compare amplitude distributions between synthetic and experimental data.
Code Team	System	Full pipeline must run without manual intervention.	MUST	End-to-end test on full dataset.
Code Team	System	Intermediate outputs must be accessible for debugging.	WANT	Verify ability to export intermediate data.
Code Team	System	Code must handle invalid or corrupted input data efficiently.	WANT	Inject incomplete or noisy data and verify error handling.
Build Team	Performance	The scanning platform operates automatically during scanning process.	MUST	Operational test to verify that the scanning platform completes a full 1D or 2D scan cycle without human intervention.
Build Team	Performance	The scanning platform shall perform the operation without mechanical failure.	MUST	Repeated operational testing with observation of platform behavior to confirm reliable continuous operation.
Build Team	Performance	The translational positioning accuracy should be within ± 1 mm	MUST	Percentage error of calibration tests must be within a millimetre of programmed distances
Build Team	Performance	The scans overlap with each other	MUST	Analysis from the collected data based on step size.
Build Team	Performance	The array holding mechanism can incorporate a rotational mechanism.	MUST	Operational test to verify that the array holder can rotate.
Build Team	Performance	The platform rig can be utilised for other projects.	WANT	Design review to assess the ease of re-configuration for other applications.
Build Team	Structural	The platform rig can be utilised without high knowledge required	WANT	Practical testing where a user follows basic instructions to operate the platform.

Group	Category	Requirement	Priority	Assessment
Build Team	Geometry	The dimension of the platform rig is 580mm×750mm×400mm.	MUST	Measurement of the overall diensions of the scanning platform after assembly to verify that the dimensions remain within specified geometry.
Build Team	Manufacturing	Maximum of 3 items to be made by the technical services.	MUST	Veriifications through manufacturing records.
Build Team	Safety	The platform rig shall not present sharp edges or exposed fasteners that could cause injuries.	MUST	Review design process and visual inspection of the accessible surfaces of the final model.
Build Team	Cost	The cost of the items to be bought must be $\leq$ £900.	MUST	Calculations of total components costs against the budget given.
Build Team	Aesthetic	The scanning platform has a clean and professional appearance	WANT	Visual inspection of the final assembled platform.

A complete copy of the full PDS is available on request.