


# 3-Dimensional Residual Neural Architecture Search for Ultrasonic Defect Detection

Shaun McKnight<sup>1\*</sup>, Christopher MacKinnon<sup>1</sup>, S. Gareth Pierce<sup>1</sup>, Ehsan Mohseni<sup>1</sup>, Vedran Tunukovic<sup>1</sup>, Charles N. MacLeod<sup>1</sup>, Randika K. W. Vithanage<sup>1</sup>, Tom O'Hare<sup>2</sup>

**Abstract**—This study presents a deep learning methodology using 3-dimensional (3D) convolutional neural networks to detect defects in carbon fiber reinforced polymer composites through volumetric ultrasonic testing data. Acquiring large amounts of ultrasonic training data experimentally is expensive and time-consuming. To address this issue, a synthetic data generation method was extended to incorporate volumetric data. By preserving the complete volumetric data, complex preprocessing is reduced, and the model can utilize spatial and temporal information that is lost during imaging. This enables the model to utilize important features that might be overlooked otherwise.

The performance of three architectures were compared. The first architecture is prevalent in the literature for the classification of volumetric datasets. The second demonstrated a hand-designed approach to architecture design, with modifications to the first architecture to address the challenges of this specific task. A key modification was the use of cuboidal kernels to account for the large aspect ratios seen in ultrasonic data. The third architecture was discovered through neural architecture search from a modified 3D Residual Neural Network (ResNet) search space. Additionally, domain-specific augmentation methods were incorporated during training, resulting in significant improvements in model performance, with a mean accuracy improvement of 22.4% on the discovered architecture. The discovered architecture demonstrated the best performance with a mean accuracy increase of 7.9% over the second best model. It was able to consistently detect all defects whilst maintaining a model size smaller than most 2-dimensional (2D) ResNets. Each model had an inference time of less than 0.5 seconds, making them efficient for the interpretation of large amounts of data.

**Index Terms**— Ultrasonic Testing, Volumetric, Image Processing and Computer Vision, Neural Architecture Search, Non-Destructive Evaluation, Synthetic Data



## I. INTRODUCTION

Composites are versatile materials that are widely used in many industries due to their superior mechanical properties such as corrosion resistance, specific strength, and specific stiffness. Carbon Fiber Reinforced Polymer (CFRP) is a widely used composite in the aerospace industry making up over 50 wt% for the two most recent long-range aircraft, the Airbus A350 and the Boeing 787, and up to 70-80 wt% for private jets and helicopters [1]. CFRP is fabricated by layering multiple carbon ply sheets and curing them with a thermoset polymer.

The anisotropic nature of these composites can be engineered to meet specific structural requirements, making them ideal for high-performance applications [2]–[10]. However, the manufacturing process can introduce defects in the composite, compromising its integrity and performance [2], [3], [5], [8], [9], [11], [12]. Defects can range from delamination and cracks to foreign object inclusions, fiber distortions, and porosity [7], [12]. As the use of composites in safety-critical parts continues to rise, the detection, characterization, and quantification of defects become increasingly important [5]. Non-Destructive Evaluation (NDE) refers to a suite of techniques employed to inspect components without causing any damage. Radiography, thermography, electromagnetic methods, and ultrasound are among the most widely used NDE techniques. These methods allow inspection of components with varying levels of complexity and size. The choice of NDE technique depends on the nature of the

<sup>1</sup>Sensor Enabled Automation, Robotics, and Control Hub (SEARCH), Centre for Ultrasonic Engineering (CUE), Electronic and Electrical Engineering Department, University of Strathclyde, Glasgow, UK

<sup>2</sup>Spirit AeroSystems, Belfast, UK

### Highlights

- **Neural Architecture Search is leveraged to discover an effective architecture for detection in ultrasonic volumetric data.**
- **The discovered architecture was able to achieve 100% accuracy on the experimental test data, a 7.8% improvement over the next best performing architecture.**
- **This work demonstrates a deep-learning defect detection technique for volumetric ultrasonic data. It provides a solution to the challenge of automated ultrasonic inspection.**

component and the defects to be detected. The application of these NDE techniques has significantly improved the

Ultrasonic Testing (UT) is a versatile technique that can be used to inspect components made of various materials, including composites, and is based on the transmission, propagation, and reception of ultrasonic waves. UT has become widely adopted and standardized for volumetric inspection in the aerospace industry due to its relatively easy and hazard free implementation compared to radiography and its ability to detect a wide range of volumetric defects [3], [7], [10], [11]. In UT, sound waves are excited on the surface of a component, and the reflected/scattered wave from internal scatterers can provide valuable information about the volumetric discontinuities of the component. Currently, phased arrays are the preferred technology for generating the initial sound wave owing to their operational flexibility. Phased arrays employ independently controllable UT elements that enable more complex electronic scanning and imaging, such as beam steering, dynamic depth focusing, and variable sub-apertures [9]. By controlling each individual element (or sub-aperture of elements) of a linear phased array, depth-wise sectional images (B-scans) can be created in a single scan (Figure 1, Figure 2 (b)). When combined with mechanized scanning perpendicular to the length of a linear phased array, complete 3-dimensional (3D) volumetric scan data of components can be generated by stacking multiple individual B-scans together at known positions (Figure 2 (b)). This technique is highly valuable for assessing the structural integrity of large and complex components and has significant implications for the reliability and safety of aerospace structures.

UT data is commonly visualized as images, either by selecting a B-scan directly, or as an amplitude or time of flight C-scan; where either the maximum response amplitude or the time index of the maximum response amplitude within the volume is imaged to produce a top-down section view across the sample [13].




Figure 1: Demonstration of how individual probe elements can make up a linear phased array which can produce B-scan and C-scan images.

reliability and safety of various structures and components across numerous industries.

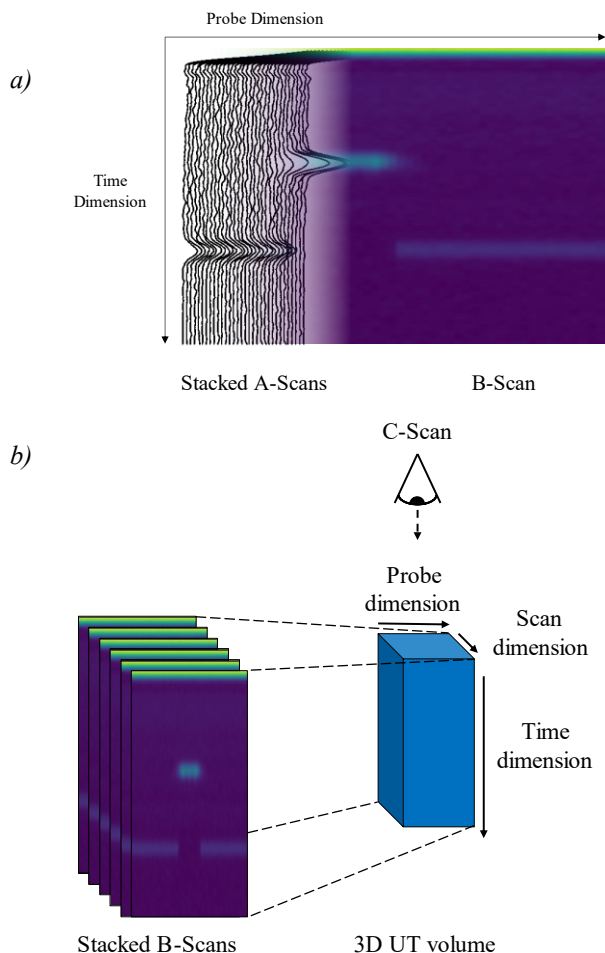


Figure 2: a) Representation of how A-scans are stacked to form B-scans. b) How B-scans are stacked to create a full UT volume.

The integration of robotics into NDE has revolutionized large-scale inspection processes by enabling efficient automated inspection of large components [14]. While robotic scanning offers greater flexibility and significantly reduces scan time compared to manual scanning, the interpretation of results remains a tedious and time-consuming task in industry. To interpret results in line with existing standards, there is a requirement for highly trained and qualified operators [10], [15]–[19]. Despite the significant improvements brought about by robotic NDE, the need for expert human interpretation of results persists. This highlights the need for further research and development of automated data interpretation techniques that can supplement or even replace human interpretation, to improve the efficiency and reliability of NDE in various industries. By reducing the dependence on

human interpretation, automation can potentially enhance the consistency, repeatability, and traceability of the NDE processes, while reducing inspection time and costs. The interpretation of UT scan results by human operators presents two significant drawbacks, namely, poor time efficiency and the risk of human error [17]. Low levels of automation for data interpretation is feasible for mass-produced parts with precisely known geometries, but this approach typically relies on hard-coded features such as predefined time-gating, filtering, and amplitude thresholding, which may not be adequate for complex tasks such as changes in manufacturing conditions, variations in geometry, or defect characterisation [19]. There is a clear necessity for an automated approach to interpret UT data, which could leverage Deep Learning (DL). This approach should seamlessly integrate with robotic inspection systems, to substantially enhance the quality and efficiency of large component inspection. Using DL would lead to shorter signal interpretation time and faster UT automation uptake in aerospace and other industries, where DL has been identified as a key requirement for transitioning from low to high levels of industrial automation [19].

Despite the potential benefits of applying DL techniques to ultrasonic signal analysis for composite components, its uptake has been limited [19]. Shortage of training data is one of the main challenges that hinders research developments in this area. This shortage, combined with industrial concerns about the interpretability and compliance with standards of DL models, has presented challenges for the effective use of DL techniques. As a result, the adoption of DL in UT signal analysis for composite components has been slow, despite its promising potential to enhance the accuracy and efficiency of defect detection and characterization.

Synthetic datasets are widely used in Machine Learning (ML) to augment small training datasets [20] and they have been successfully implemented for UT of composites with encouraging results for 2-dimensional (2D) classification of C-scan images [21]. Part of this work builds upon the work in the previous paper to extend one of the synthetic data generation methods to make it applicable for full 3D volumes. The synthetic datasets are based on simulations from semi-analytical physics based software that has been shown to produce experimentally accurate defect responses [22], [23]. This software offers a less computationally expensive alternative to Finite Element Analysis (FEA), allowing for the simulation of composite responses based on bulk material properties [24].

When ML is used to interpret UT NDE data in literature, it is typically applied to interpret A-scan time traces or 2D images constructed from A-scans [25]–[31]. Compared to B-scans, A-scans lack all spatial information and nowadays, they are rarely used alone to characterize defects by human operators since the introduction of phased arrays. C-scans preserve detailed spatial information, however constructing the 2D image from the volumetric data necessitates the compression of temporal information. Whilst C-scans excel in capturing intricate spatial details, their need for temporal compression results in minimal representation of through-depth features. Compression of A-Scans to C-Scans often removes useful

features such as the backwall response, which can be important when detecting defects with a low reflective index such as porosity [32]. Furthermore, to produce C-scan images appropriate gating must be applied to remove the front wall surface response. This can be challenging when trying to detect near-surface defects. In the aerospace industry, operators typically start with a C-scan to gain a complete picture of defect responses and then move to analysis of B-scans for further information about the nature of the responses [33], [34]. Whilst current ML approaches in literature make use of data in formats that are easily interpreted by humans (images or time-traces), ML algorithms are not limited to image-level analysis and have proved very capable at interpreting 3D volumetric data [35], [36]. By implementing algorithms capable of volumetric interpretation, we retain all spatial and depth information, this gives the algorithms more relevant features to learn from and removes the need for image pre-processing and gating.

Convolutional Neural Networks (CNNs) have been used effectively for decades in a wide variety of image and volumetric analysis tasks with models such as ResNet typically having tens of millions of parameters [37], and are still widely used as backbones or standalone architectures [38]. However, these networks are typically applied to data of similar dimensions, or data which has been scaled to give even dimensionality of each axis. UT data has extreme aspect ratios due to the difference in requirements of sample rate in the spatial and time dimensions. Compressing the data in the time dimension to match the spatial dimension, normally dictated by the sub-aperture pitch and the scan acquisition rate, would result in a substantial loss of depth information. Alternatively, the spatial dimensions could be upsampled to match the number of samples in the time dimension, but this is highly inefficient, creating data instances that would require large amounts of memory, and would make training intractable. Therefore, retaining the original dimensionality and aspect ratio of the UT data is highly preferable. Using CNNs to interpret images with high dimensionality is not new and the use of rectangular kernels instead of square kernels in CNNs has given positive results for classification of speech signals, which have high aspect ratios when represented as spectrogram images [39]. This paper makes use of a similar approach for volumetric data.

Network architecture design is a key component of effectively leveraging machine learning techniques. Traditionally, network design heuristics and 'rules of thumb' would be used, in tandem with domain expert knowledge to construct a specific architecture. Automatic architecture design or Neural Architecture Search (NAS) is a development on this approach where a practitioner can leverage compute to aid the process of architecture selection. This process, which can be considered a subset of hyperparameter optimization, generally involves an iterative process of selecting, training, and evaluating architectures. In its simplest form, a 'Random Search' involves repeating the above process until some threshold or limit in terms of performance or computation budget is reached. More complex approaches to NAS often focus on efficient model evaluations, making use of proxy evaluation methods [40], [41] or efficient sampling algorithms

[42], [43] attempting to make the largest improvement with each evaluation.

This paper presents a comparative analysis of the performance achieved from three separate architectures for defect detection in volumetric ultrasonic data. The first, VoxNet [44], is prevalent in the literature for volumetric classification problems, the second architecture presents modifications to VoxNet for this task using a traditional network design approach, and finally a discovered architecture from NAS. VoxNet is a 3D CNN initially proposed for classification of LiDAR, RGBD and CAD data. It has since been used as a backbone for different volumetric classification tasks [45]. Additionally, notable contributions of this study to knowledge in the field encompass the introduction of domain-specific augmentations, which exert a substantial impact on the classification performance. Furthermore, synthetic data generation techniques are leveraged from prior 2D work to generate 3D UT datasets from semi-analytical simulations, effectively addressing one of the prominent challenges encountered in the application of deep learning for NDE: the scarcity of effective training data.

This work presents a novel DL architecture designed to process volumetric UT data. In contrast to prior methods relying on time-series data or 2D image-based approaches, which diminish spatial or temporal features, whilst often requiring additional processing. The main contributions of this work are:

- Interpretation of volumetric UT data, instead of images or time signals. This reduces preprocessing requirements and allows the model to learn from greater features.
- Introduction of two domain specific methods for data augmentation, helping with the domain transfer from synthetic to experimental data.
- Discovery of a novel 3D CNN architecture through NAS.

### 1.1. Pipeline overview

In this work, the automated data interpretation is simplified by inspecting the complete volumetric data, eliminating image processing steps like gating to remove front and back wall responses, while preserving all spatial and temporal information. Whilst the models are trained on synthetic data, they are tested using experimentally collected UT data from samples with manufactured defects that aim to mimic delamination's. Manufactured defects are commonly used in literature to act as test cases and qualify NDE techniques and operators where naturally occurring defects are not always available [6], [27], [28]. An overview of the simulation and deep learning pipeline is presented in Figure 3. Figure 3 also shows how NAS can fit into this process, with Figure 15 providing a more detailed overview of the NAS pipeline.

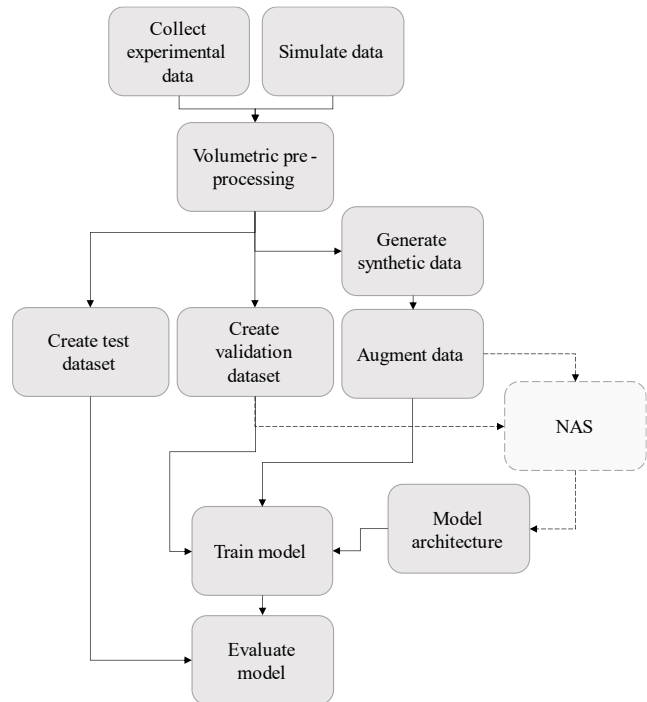


Figure 3: Overview of the pipeline for automated volumetric UT classification.

## 2. ACQUISITION OF EXPERIMENTAL TEST DATA AND SYNTHETIC TRAINING DATA

### 2.1. Experimental data collection

Experimental ultrasonic data was acquired from CFRP samples, both with and without artificially introduced defects, to serve as test data. To imitate delamination defects, which are one of the most common defects in composites [31] and a significant life-limiting failure mode [46], flat-bottom holes were drilled from the backside of the samples. Prior to introducing defects, clean scans of each sample were taken to form a defect-free test set. The use of the same CFRP base sample ensured that the trained models learned defect-specific responses rather than the underlying properties of different composite samples.

Composite samples measuring  $254.0 \times 254.0 \times 8.6$  mm (W  $\times$  D  $\times$  H) were provided by Spirit AeroSystems and were manufactured to the BAPS 260 specification with woven fabric, and Cycom 890 resin using a vacuum assisted resin transfer molding process. Of the three samples, two contained defects. The first sample contained 15 flat-bottom holes measuring 3.0, 6.0, and 9.0 mm in diameter, with each defect drilled to depths of 1.5, 3.0, 4.5, 6.0, and 7.5 mm from the front surface. The different defect sizes were spaced 30 mm apart with different depth defects spaced 35 mm apart. The second sample contained 25 flat-bottom holes, drilled to the same depths as the first sample but with additional defect diameters of 4.0 and 7.0 mm, as shown in Figure 4. All defects were manufactured to tolerances in depth of  $\pm 0.3$  mm and in diameter of  $\pm 0.2$  mm.

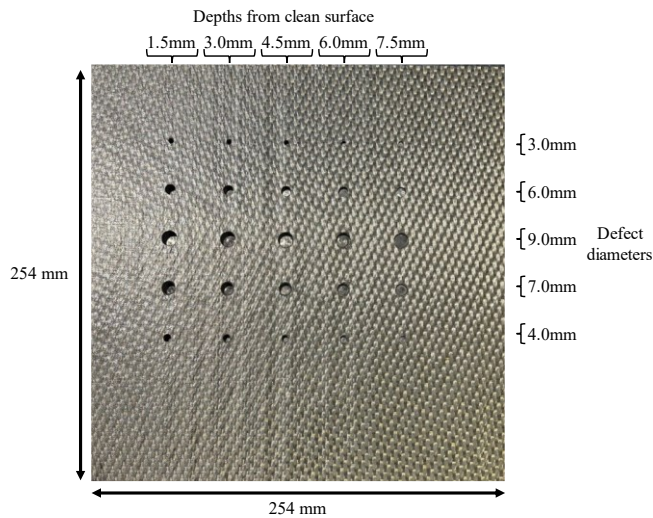


Figure 4: The composite test sample showing 25 Flat-Bottom Holes.

The ultrasonic data was acquired at room temperature using a robotically deployed unfocused linear phased array. The array used was an Olympus Inspection Solutions RollerFORM-5L64 [47], which had a central frequency of 5 MHz and was made up of 64-elements with a pitch of 0.8 mm and elevation of 6.4 mm. The elements were driven at 100 V with a receiver gain of 22.5 dB. The sample rate was 100MHz. The pulse repetition frequency was set to collect a B-scan every 0.8 mm with a scan speed of 10 mm/s which was controlled using a fully automated robotic system built around a KUKA KR 90 R3100 extra HA industrial robot (Figure 5) [48]. Robotic scanning enabled the concatenation of encoded B-scans to form volumetric datasets. To ensure a steady coupling of the roller-probe to the surface of the component and consistent transfer of acoustic wave energy into the sample at different scanning positions, Force-Torque compensation was used to control the distance from the samples surface with feedback from the force axis perpendicular to the sample. This was accomplished with integration of a Schunk GmbH & Co. FTN-GAMMA-IP65 SI-130-10 Force-Torque sensor, mounted between the robot's flange and the roller-probe. This ensured a constant scanning force of 70 N which maintained consistent tyre compression throughout the scan. Water was used as an acoustic couplant in the scanning process. This data acquisition setup is widely used in industry and has been employed for data collection on large composite aerospace components [49].




Figure 5: Overview of the experimental scan setup of KUKA KR90, Force-Torque sensor, and ultrasonic roller probe used for data acquisition.

## 2.2. CIVA Simulations

Due to the lack of available experimental training data, a simulated dataset was constructed for training. This was done using CIVA, a semi-analytical physics-based commercial NDE simulation software [50]. CIVA has the ability to accurately model wave propagation and its interaction with defects, which has been experimentally validated for UT [22], [23]. Additionally, the software is computationally efficient when compared to other alternatives such as Finite Element Analysis (FEA). The full control of the simulated domain enabled the modelling of similar defects and material properties to the experimental domain. However, the use of semi-analytical software instead of FEA had limitations in that the software was unable to model responses from ply interactions and lacks noise seen in experimental data. As a result, differences existed between the simulations and measured experimental responses, leading to the use of the synthetic data generation steps discussed in Section 2.4 to reduce the differences between simulated and experimental domains. For further information on the difference between the simulated and experimental domain and the need for accurate synthetic data, please refer to the previous work on this topic [21].

To set up the simulation, the individual layers of composite were constructed and used to generate equivalent homogeneous material properties of the experimental CFRP samples. A single ply layer was constructed and alternated repeatedly with 8 layers at orientations of 0, 45, -45, and 90 degrees to match the experimental sample as closely as possible. The resulting multilayer structure was homogenized to be consistent with a homogeneous medium having mechanical properties equivalent to those of the multi-ply composite, with the fiber density set to 50% best match the experimental sample's density of 1440 kg/m<sup>3</sup>.

To simulate the waveform, a sinusoidal wave of 5 MHz was employed, accompanied by a Hanning filter that provided a bandwidth of 66% at 12 dBs.

For running multiple, sequential simulations, a parametric study was set up, using the composite bulk properties previously calculated and varying the diameter and depth of defects. Flat bottom hole defects were simulated with diameters from 3.0 mm to 15.0 mm, increasing every 0.5 mm, with varying depths from 1.5 mm to 7.0 mm from the surface, in increments of 1.5 mm. A defect-free simulation was also run to provide the basis for defect-free synthetic data. Both the front and back wall surface reflections were included in the model. The full simulations took less than 15 hours on a desktop computer with a 24-Core 3.79 GHz CPU and 128 GB of memory.

## 2.3. Signal processing and dataset generation

The resolution of the UT data in the array dimensions was constrained by the element pitch, and the scan width was restricted by the number of elements in the array. This limited the inspection data to 64 voxels in the array dimension. To

match this, 64 B-scans were selected in the scan dimension to create cuboidal datasets (Figure 2(b)). The distance between the elements was 0.8 mm, and the robotic scanning speed was regulated with the pulse repetition frequency to ensure a B-scan offset of 0.8 mm. This enabled the generation of volumes with square voxels in the spatial domains, along both the probe and scan directions. By utilizing this approach, the study was able to achieve a standardized volumetric resolution that was consistent throughout the dataset.

The data collected in both experimental and simulated domains was in the form of radio frequency A-scans, also known as amplitude scans. In order to obtain 3D volumetric datasets from these sources, a number of signal processing steps were performed. Initially, the A-scans were centered at zero amplitude and enveloped by taking the absolute of the Hilbert transform, as shown in Figure 6 (a). The Hilbert transform was used to obtain the analytical signal, which is useful for calculating the instantaneous response of a time series. This approach is a standard signal processing technique used when generating C-scan images from time series ultrasonic data. Subsequently, each volumetric dataset was normalized between 0 and 1 by dividing by its maximum peak amplitude.

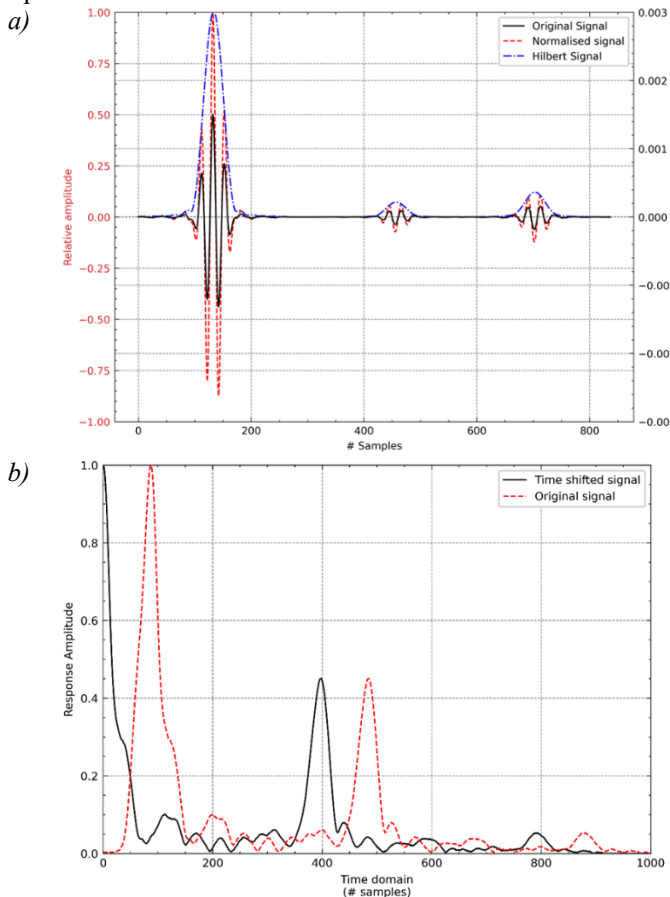


Figure 6: a) Example of relative amplitude response from simulations, normalized signal, and Hilbert transform, applied to the original signal. b) Demonstration of how individual A-scans are time shifted to the front wall response.

Once the data was normalized, the offset in the time domain was compensated for by aligning the peak front wall response to the origin. This made sure that features were correctly

aligned in the time domain and helped to account for any variability in the acoustic path length between individual transducers and the surface of the sample. Figure 6 (b) Shows how the time shifting was done for an individual A-scan with the Hilbert transform applied. Figure 7 shows the effect of this on the complete ultrasonic volume.

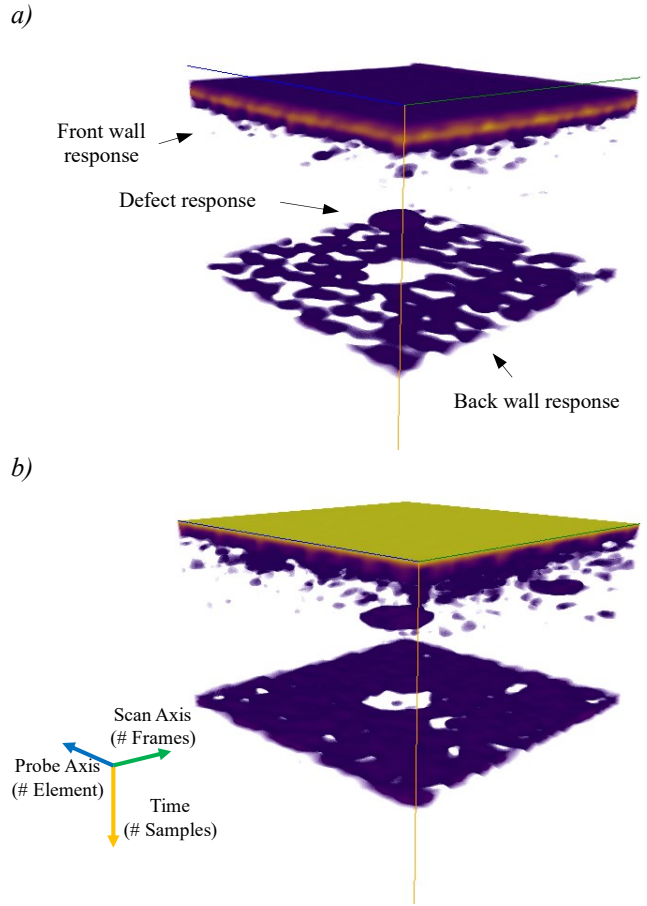


Figure 7: a) Volumetric data with Hilbert transform applied only. b) Volumetric data with time shifting to the central response of the front wall peak. Both figures have been thresholded to remove the lowest 10% of amplitudes to aid in visual clarity.

#### 2.4. Synthetic data generation method


Our previous studies have shown that semi-analytical simulated data alone is not representative enough of the experimental domain [21]. Therefore, there is a need for methods of translating the simulated domain closer to the experimental domain. Fully statistical methods of generating noise are advantageous as they can be re-sampled continuously to keep generating unique noise profiles which are in line with experimental data. In this paper, we extend previous work in generating 2D synthetic images, and propose a new approach for adding noise to complete volumetric UT data.

The previous study concluded [21] that A-scan level noise was the best fully generative statistical method for adding noise. Additionally, all the other approaches, except for the simulated A-scan noise, introduced noise at an image level, which is intractable for volumetric data. To adapt the methodology described in the original paper for the analysis of full volumetric data, unique noise

profiles for each A-scan were generated and subsequently summed with the simulated responses past the front wall. To temporally align the responses, the time shift of the front wall response was performed when generating the noise distributions from experimental data, and when combining the simulated responses with the generated noise profiles (Figure 6 b)).

Figure 8 shows an example of the addition of noise on simulated data at an A-scan level and Figure 9 demonstrates this for a complete ultrasonic volume. The statistical noise distributions of the A-scans were calculated from a separate hold out sample with the same layout and thickness as the test samples. For further details on building up the noise profiles, please refer to the previous work [21].

A)



b)

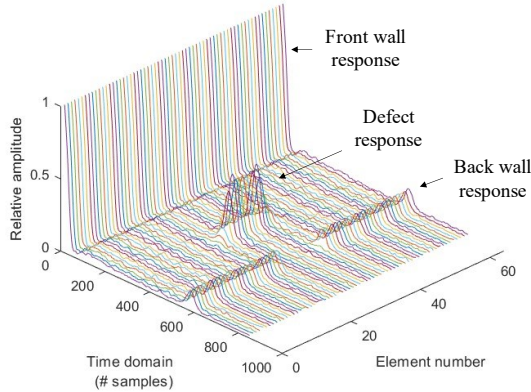



Figure 8: a) A frame of 64 simulated A-scans for a simulated defect response. b) the corresponding A-scans with synthetically added noise for the same defect response.

a)



b)

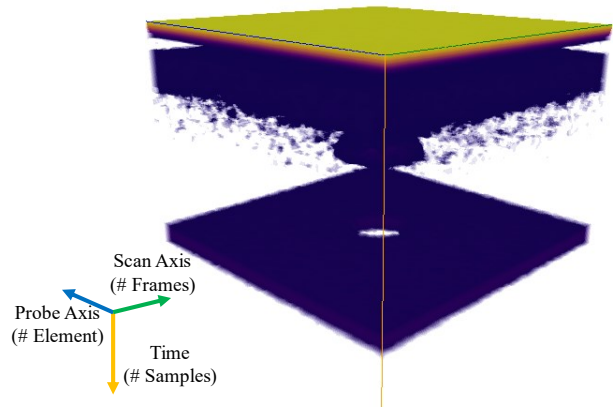


Figure 9: a) Complete ultrasonic volume of simulated A-scans for a defect response. b) the corresponding synthetically noised volume for the same defective response. Both figures have been thresholded to remove the lowest 10% of amplitudes to aid in visual clarity.

A summary of the datasets generated from the experimental and synthetic data is given in Table 1.

Table 1: Summary of the datasets produced.

Data source	Dataset	Number of datapoints
Simulated defect responses (300 Flat-Bottom Holes)	Synthetic defective train	300
Simulated defect free response	Synthetic defect free train	300
Experimental defect reference sample (15+25 Flat-Bottom Holes)	Defect test (70%)	25
	Defect validation (30%)	15
Experimental defect free reference sample	Defect free test (70%)	25
	Defect free validation (30%)	15

### 2.5. Augmentation methods for synthetic training data

The generalizability of ML models is a critical aspect of their performance. One approach to improve generalizability is to augment the training data. Augmenting the training data makes the task more challenging by adding noise at the training stage, reducing the likelihood of overfitting, and often improves performance in the target domain. This is particularly important when the target (experimental) domain is different from the training (synthetic) domain.

As demonstrated in Figure 9 (a), there is little variation between simulated A-Scans. However, this is not the case for

experimentally acquired data. Received amplitudes are affected by the sensitivity of individual elements, variations of couplant on the surface of the sample, and roughness of surface finishes (particularly for manufactured defects). The anisotropy of CFRP can also result in variations in attenuation, which impact effect received amplitudes. Surface roughness and local changes in fiber density due to the materials inherent anisotropy produce small changes in time of flight. Traditional augmentation methods such as those used for images (e.g. crop, mix-up, flipping etc.) do not model these variations well and can produce unrealistic examples. Therefore, in this study, we introduce two types of augmentation that were generated online for each minibatch during training. These augmentations aim to mimic the inter-element response variability observed within the UT probes used for data collection.

The first type of augmentation is related to the magnitude of response measured by the UT elements, which varies due to many factors not included in the simulation, such as manufacturing tolerances of the sample and the UT array probe, wear and tear of the probe and electrical wires/connections, or inter-layer multiple scattering of the sound waves. To mimic these whilst preserving the correct normalization, each A-scan was scaled by a constant past the front wall. The scale factor was sampled from a uniform distribution to give a scale factor between 80-120%. An example of this is given in Figure 10 (a).

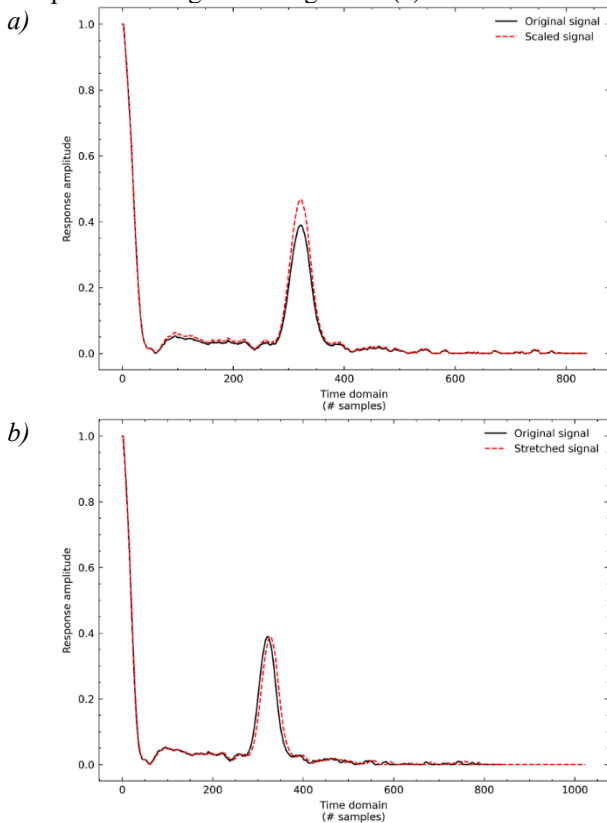


Figure 10: a) Example of how scaling augmentation is done on an individual A-scan. b) Example of how dilation augmentation and padding is completed for an individual A-scan.

The second type of augmentation mimics any changes in ultrasonic travel time seen by different elements. This can be

caused by a variety of factors, such as variations in component sound speed due to the anisotropic nature of composites, departure from central frequency for certain elements, etc. To simulate this 1-D interpolation was used to randomly stretch or compress the signal in the time domain. The dilation amount was randomly sampled from a uniform distribution for each A-scan up to  $\pm 15$  samples. An example of this is given in Figure 10 (b).

By introducing these augmentation methods, we aim to improve the generalizability of the models to the experimental domain. The online nature of these augmentations means that they can be easily incorporated into the training process without the need for additional data collection or pre-processing steps. To ensure consistent length of data in the time domain, each A-scan was padded with zeroes to a length of 1024 samples during training.

### 3. NETWORK ARCHITECTURES

In this paper we investigated the performance of three different 3D CNN architectures for binary classification of 3D defect and defect free UT data with extreme aspect ratios. The first 3D CNN, VoxNet, was designed for similar volumetric classification tasks and acts as a baseline architecture. For low aspect ratios CNNs (such as VoxNet) typically make use of square or cuboidal kernels which are appropriate for their equal (or near equal) aspect ratios. The use of CNNs on data with more extreme aspect ratios is less common and is particularly extreme for UT data between the time and the spatial domains, with an aspect ratio of 16. To overcome this challenge, a task-specific architecture was hand-crafted by adapting VoxNet in a manner that follows the traditional approach to architecture design. This custom network is specifically designed to tackle the extreme aspect ratio problem and enhance overall classification performance. As an alternative to traditional architecture design, neural architecture search was employed to develop a third architecture for comparison. For each model Adam optimizer [51] was used with a constant learning rate of 0.001,  $\beta_1$  of 0.9 and  $\beta_2$  of 0.999. A batch size of 8 was utilized in the training process. The chosen loss function for this model was binary cross-entropy, with a sigmoid activation function applied to the final layer to facilitate classification.

Due to the small amounts of experimental test data, there was a likelihood of noisy results during both training and testing phases. To mitigate this, each model was trained ten times with varying random initializations, and their individual results were averaged across the performance metrics. This gives a better representation of the model's performance by averaging out any noisy results due to the small datasets. During the training phase, a fixed validation set comprising 30% of the total test data was randomly selected from each class of experimental data. This set was used to monitor the model's performance and minimize the risk of overfitting. The models were trained with a patience of 10 epochs, where the training process monitored binary cross entropy loss on the validation data, for improvement. If there was no enhancement for a consecutive period of 10 epochs, the training process was halted. The model parameters with the lowest validation loss were used to evaluate the classification



performance on the test set. This approach ensured that the final model's defect detection performance was evaluated using the parameters that had the best ability to generalize to the target domain, as opposed to the model that had overfit to the synthetic domain.

### 3.1. Evaluation metrics

To quantitatively assess the binary classification performance of each network, average mean accuracy, precision, recall and F1 scores were calculated according to Equations 1-4.

$$Accuracy = (TP + TN) / (TP + TN + FP + FN) \quad 1$$

$$Precision = TP / (TP + FP) \quad 2$$

$$Recall = TP / (TP + FN) \quad 3$$

$$F1 = (2 \times Precision \times Recall) / (Precision + Recall) \quad 4$$

Where TP is true positive, TN is true negative, FP is false positive, and FN is false negative, with positives being the presence of a defect. Each result was individually averaged using a simple mean across the 10 training cycles.

### 3.2. VoxNet: Baseline Architecture

Introduced by Maturana and Scherer, VoxNet [44] is a 3D CNN designed to tackle classification problems of 3D data that can be represented as voxels to form an occupancy grid. Originally tested on LiDAR, RGBD and CAD data it has since been used as the backbone for methods tested on ModelNet40 [45].

While the data from UT for this task differs from the datasets previously employed with VoxNet, the process of converting data into voxel-based format within the VoxNet pipeline is well aligned to the 3D representation of UT data. As a result, VoxNet was employed to establish baseline model performance metrics for this task.

VoxNet is constructed using two 3D convolutional layers with cuboidal kernels, followed by a pooling layer and two fully connected layers (Figure 11). For further details on the model please refer to the original paper. VoxNets total number of parameters is 235M.

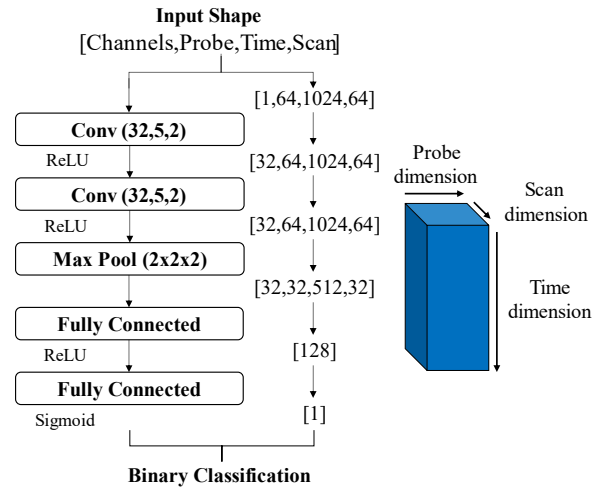


Figure 11: The VoxNet architecture. Where Conv ( $f,d,s$ ) indicates the number of filters  $f$ , filter size  $d$ , and stride  $s$ , of the convolutional layer.

### 3.3. Hand Designed Architecture

The second architecture, referred to as CustomNet, demonstrates a conventional approach to architectural design. In this context, adaptations to VoxNet have been implemented to contemporize and enhance its performance specifically for the given task.

Our UT dataset stands out to previous VoxNet datasets due to its higher dimensionality coupled with notable differences in spatial and temporal dimensions. To effectively handle these unique attributes, adjustments were made to the models' architecture. Specifically, the number of convolutional layers was increased to enhance the extraction of meaningful features from the complex dataset. Additionally, cuboidal kernels with non-uniform dimensions were employed in the initial four blocks (refer to constant blocks in Figure 12) of the model. This approach aimed to address the uneven dimensionality inherent in the data, ultimately equalizing the dimensions and contributing to a more robust feature representation throughout the network (Figure 12). After this a feature block with cube kernels of equal dimensionality could be used (refer to feature block in Figure 12).

In the process of updating VoxNet, we incorporated convolutional layers for pooling instead of the previously employed max-pooling layers. Additionally, ReLU was substituted with LeakyReLU, and batch normalization was introduced. To mitigate overfitting, dropout and global average pooling were employed to reduce the number of features for classification, avoiding the use of large fully connected layers. These modifications are geared towards improving the model's performance by incorporating contemporary practices that have shown substantial performance benefits, as highlighted in previous studies [52]. The final architecture is given by the diagram in Figure 12. The total parameter size of the network was estimated to be 1.28 M parameters.

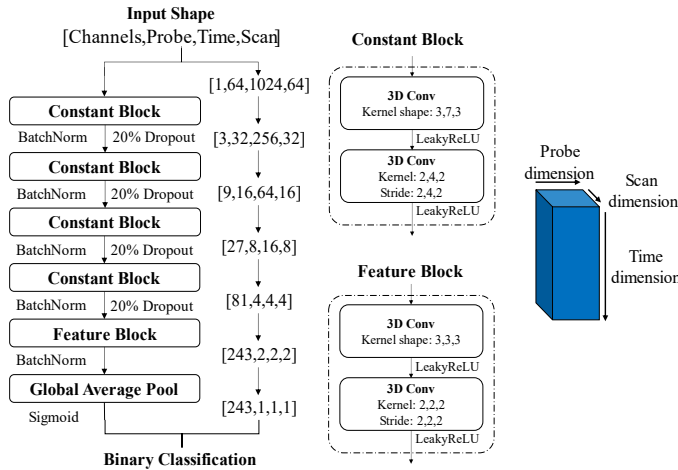


Figure 12: Network architecture for the CustomNet.

### 3.4. NAS Discovered: 3D ResNet based Neural Architecture Search

#### 3.4.1. Neural Architecture Search

The final architecture was developed through NAS of a modified ResNet search space to account for 3D convolutions and operations. One of the challenges in applying NAS to a new domain task is the design of the search space. For this task, a new search-space framework which utilizes a novel search space based on a ResNet-like structure is introduced. A fixed stem was used to down sample the data by a factor of 4 in the spatial dimensions and a factor of 8 in the time dimension whilst aiming to retain information through increasing the channels to 64. A further down sample block with average pooling followed by two to four residual blocks were all searched individually. An overview of the structure can be seen in Figure 13. The residual blocks and bottleneck features of the ResNet architecture are retained, whilst searching operations for each edge within the residual block. This provided a large diversity of architectures, which is key to attaining good performance in a novel application, whilst also ensuring that many networks conformed to successful design principles. Each residual block contained two fixed point-wise convolutions used to down and up sample the number of channels. Figure 14 shows an example of a residual block denoting the searched and fixed operations.

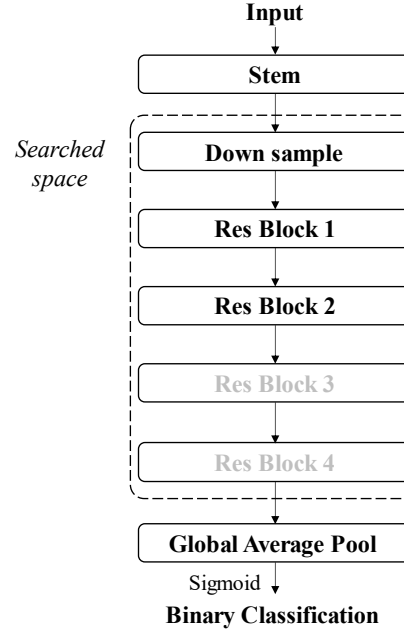


Figure 13: Representation of the ResNet style searched space.




Figure 14: Diagram of the searched residual block.

These blocks were then stacked in groups, with the resolution down sampled between groups. Equation 5 gives the probability of a new group being created for each residual block, otherwise they were added to the current group. This makes groups unlikely to be extremely long or short.

$$P(\text{newResBlockGroup}) = \frac{1}{\text{currentResBlockGroupSize}} \quad 5$$

The primitive operations of a search space are the list of operations which are assigned to the edges of a network architecture. The implemented approach incorporated a standard set of operations commonly found in the NAS literature. These operations comprised of convolutions, pooling, and skip connections, which are widely recognized and utilized within the field. These operations were all 3D due to the dimensionality of the data. In contrast to standard practice, which makes use of separable convolutions, the approach presented in this study deployed both depth-wise and point-wise convolutions as the fundamental convolutions within the search space. This significantly reduced the number of parameters in each operation of the architecture, greatly reducing the computational cost. Specifically, the depth-wise convolutions were applied with equidimensional cube kernels, of size 3, 5, or 7, coupled with dilation values that ranged

from 1 to 4. Skip connections, point-wise convolutions, as well as average and max pooling operations were also searched for. For the pooling operations, equidimensional cube kernels of size 3, 5, or 7, with a dilation value of one were employed. The search encompassed the exploration of GELU activation function and batch normalization, as well as the absence of activation and normalization operations. This allowed for architectures with fewer activation and normalization function which has been shown to be beneficial [52]. The searched down sample operation had a fixed kernel size with two in the spatial dimensions a four in the temporal dimension with a dilation of one. Throughout the relevant operations, a stride of one was employed.

A simple random search was applied to this search space for 80 iterations. Each model was evaluated using the validation dataset, with the lowest loss on validation across the training taken as the evaluation metric. For each searched architecture, a model was retrained with new initializations three times and the mean evaluation metrics were used when selecting the discovered architecture, this ensured a more accurate estimate of model performance. Cross validation was unable to be used as the combination of NAS and domain transfer would have resulted in data leak between the NAS stage and the final model test evaluation stage. Figure 15 provides an overview of the NAS process and demonstrates how separation of the validation and test set were maintained in context of the complete model pipeline, given in Figure 3.

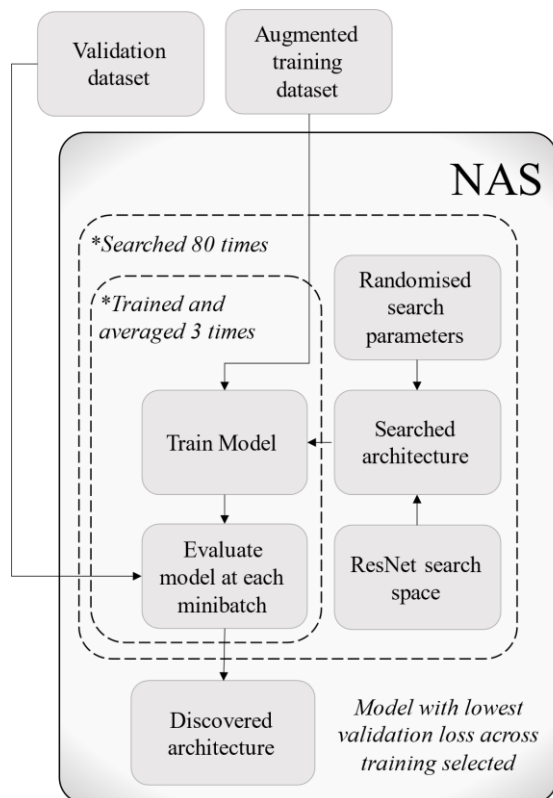


Figure 15: Overview of the process for NAS implementation

The final discovered architecture had 1.03 M parameters and is given in Figure 16, with the details of the residual blocks given in Figure 17.

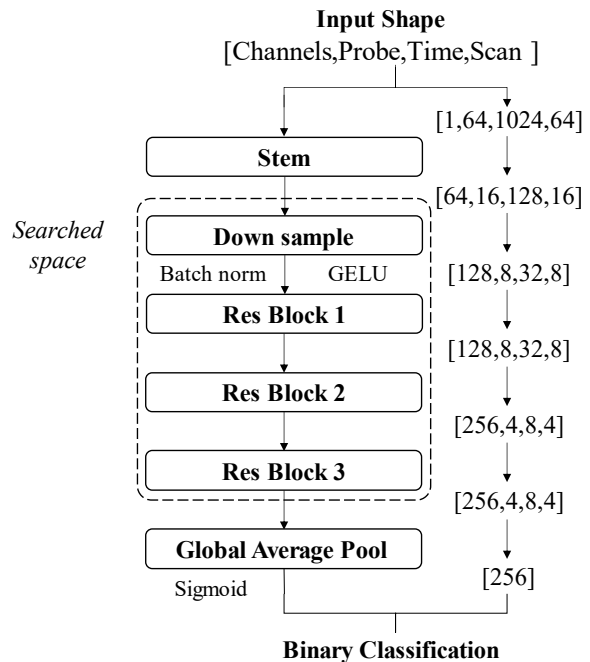
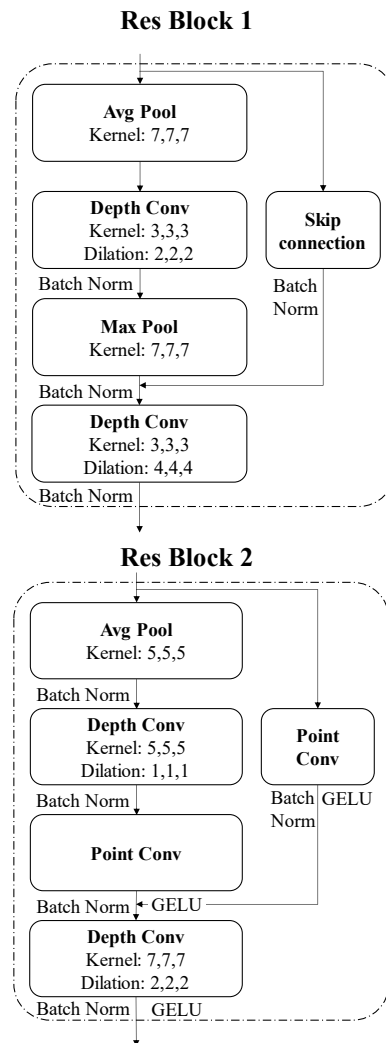


Figure 16: The overall structure of the discovered architecture.



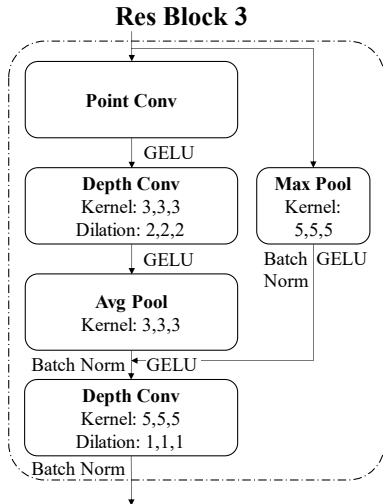


Figure 17: The details of each discovered residual block.

#### 4. RESULTS

Table 2: Average confusion matrices for VoxNet, CustomNet and the NAS discovered architecture.

	True	Predicted	
		No Defect	Defect
VoxNet	No Defect	15.8	9.2
	Defect	2.1	22.9
CustomNet	No Defect	24.3	0.7
	Defect	3.2	21.8
NAS Discovered	No Defect	25.0	0.0
	Defect	0.0	25.0

Table 3: Comparison of classification results across the different architectures. The means and standard deviations are presented as mean  $\pm$  std.

Model	VoxNet	CustomNet	NAS
Accuracy	0.774 $\pm$ 0.184	0.922 $\pm$ 0.095	<b>1.00</b> $\pm$ 0.00
F1	0.825 $\pm$ 0.114	0.904 $\pm$ 0.130	<b>1.00</b> $\pm$ 0.00
Precision	0.793 $\pm$ 0.219	0.975 $\pm$ 0.050	<b>1.00</b> $\pm$ 0.00
Recall	0.916 $\pm$ 0.073	0.872 $\pm$ 0.197	<b>1.00</b> $\pm$ 0.00

Table 2 provides average confusion matrixes for the test results of the VoxNet, CustomNet and NAS discovered models. Table 3 presents a summary of each method's performance, displaying the mean and standard deviation across various performance metrics.

The architecture discovered by NAS consistently produced ideal results when trained using data augmentation, with a mean classification accuracy of 1.00 and a standard deviation of 0.00 across the 10 separate training iterations, this demonstrated high confidence in the model's conclusions and robust design for the target domain.

Table 4: Comparison of the effects of data augmentation on the NAS discovered architecture. The means and standard deviations are presented as mean  $\pm$  std.

Augmentation	None	Scaling	Both
Accuracy	0.776 $\pm$ 0.178	0.806 $\pm$ 0.227	<b>1.00</b> $\pm$ 0.00
F1	0.830 $\pm$ 0.123	0.842 $\pm$ 0.168	<b>1.00</b> $\pm$ 0.00
Precision	0.745 $\pm$ 0.19	0.846 $\pm$ 0.227	<b>1.00</b> $\pm$ 0.00
Recall	0.972 $\pm$ 0.044	0.916 $\pm$ 0.182	<b>1.00</b> $\pm$ 0.00

Table 5: Comparison of model sizes and inference time for each architecture.

Model	VoxNet	CustomNet	NAS
Total Parameters (M)	235	1.28	<b>1.03</b>
Total Size (MB)	1779	557	<b>93</b>
Inference Time (seconds)	0.37	<b>0.03</b>	0.40

Table 4 demonstrates the impact of data augmentation on the best performing NAS model. Discarding data augmentation completely during training had a significant impact on the classification performance with a 22.4% drop in mean accuracy, along with a standard deviation increase in 17.8%, which demonstrated a significant reduction in statistical confidence. Whilst the addition of amplitude scaling augmentation improved the mean accuracy, it was only by 3%. This demonstrates the importance of using both augmentation methods in parallel for increased generalizability to the experimental domain.

Table 5 provides a summary of model sizes, and inference times for a single batch of test data. The NAS discovered architecture has a total size 16.7% and 5.2% smaller than CustomNet and VoxNet respectively. Whilst the CustomNet was 12 times faster at inference than the next closest, VoxNet.

#### 5. DISCUSSION

VoxNet demonstrated it was able to learn features from synthetic data and performed reasonably well on experimental data, with a mean F1 score of 0.825. However, its significant standard deviation in accuracies between training instances demonstrates that the architecture was not well optimized for the problem. The CustomNet improved on the accuracy of VoxNet substantially by 14.8%, whilst also reducing the standard deviation of results by 8.9% which indicated an increase in consistent generalizability to the experimental domain. This illustrates the benefits of tailoring architectural modifications to address the needs of specialized tasks. The experimental results demonstrated that the architecture discovered from NAS greatly outperformed the other two in terms of classification accuracy. Whilst all the models used in this paper are not large and are considerably smaller than typical sizes for 2D ResNet's and other CNNs [37], the NAS model was able to achieve the highest performance with a significantly lower model size, at only 5.2% the memory requirement of VoxNet. The black box nature of DL makes it difficult to specify which features lead to this improvement in

performance. This complexity is in fact a large motivator for NAS as the design space is too large for a human to efficiently find an optimal network architecture. The authors believe that the addition of skip connections and the ability to vary operations at different depths added by the NAS has a significant positive impact on performance. This demonstrated the importance of utilizing neural architecture search to optimize CNNs.

Due to the large fully connected layers, VoxNet results in a far greater number of total parameters than the other two networks. This results in a model which occupies far more memory. Whilst CustomNet and the NAS discovered architecture have a comparable number of parameters the discovered network is far smaller. This is a result of many of its operations being far more efficient, such as the separation of point and depth wise convolutions. Whilst the discovered architecture is the smallest, its inference time takes the longest due to the greater architectural complexity of the model and its operations. This said, all models have acceptable inference time and can process 8 samples in under half a second. However, CustomNet is notably twelve times faster at inference than the second fastest network, VoxNet, which could be an advantage in some industrial settings.

When trained without data augmentation the NAS model performed significantly worse. Furthermore, the performance was only slightly improved by adding amplitude scaling augmentation alone. For best performance, both augmentation methods were needed in combination. This indicates that despite accurate synthetic data generation, data augmentation still has a significant role in producing generalizable models to the experimental domain.

Whilst ideal classification was achieved consistently for the discovered architecture when trained with data augmentation, this was tested on detection of manufactured defects only. Specifically, back drilled holes which are perpendicular to the propagating sound wave and act as ideal reflectors. This makes them comparably easier to detect than other defects.

Whilst samples with naturally occurring defects are challenging to get access to, in future work the authors aim to expand the simulation scope and test the models on naturally occurring defects which will likely prove more challenging to detect. For more challenging detection and characterization tasks a more sophisticated search optimization algorithm could be employed to discover architectures more efficiently. The achieved classification results suggest that the synthetic data generation process is a viable approach for producing fully synthetic 3D UT volumetric datasets that closely map to the experimental domain and enable the development of effective classifiers. However, due to the substantial improvement in classification performance achieved through the implementation of data augmentation methods, it is important to acknowledge that disparities between the synthetic and experimental domains persist. This observation underscores the necessity for augmentation techniques to further enhance the generalizability of the model.

Nonetheless, it is worth noting that the data augmentation methods employed in this study proved to be highly effective in aiding not only generalizability but also in facilitating the transfer of knowledge across domains.

The key benefits for analyzing the complete 3-D volumetric data instead of processed images were the ability to learn from greater features, the reduction in pre-processing requirements, and the potential reduction in inference time by analyzing the complete volume all at once. The impact on inference time is challenging to quantify, however if comparing the compute required to process 64 B-scan images (the equivalent spatial scan data), without parallelization for equivalent 2D classifiers, there is the potential for up to 64 times saving in inference time for the same scan area. Despite these advantages there are still potential benefits to analyzing UT data as images. One of these is the many opportunities for detection of a single defect in multiple B-scans. It is likely that defects will span multiple B-scan images, and as such by analyzing each B-scan there are multiple chances to detect an individual defect. This means an individual defect can still be detected even if individual defective images are incorrectly classified. However, the opportunities for characterization and localization of defects are far greater when retaining the volumetric spatial information and this work opens future prospects for 3D classification and segmentation which would be much more challenging if using C-scans or B-scans alone. The research outcomes demonstrated the considerable potential of employing 3D-CNNs in conjunction with well-designed data augmentation techniques and optimized architecture search spaces to address challenging 3D classification tasks characterized by extreme aspect ratios, as observed in the context of UT. Insufficient utilization of data augmentation severely hampered the model's ability to generalize to experimental datasets, leading to suboptimal classification performance. Likewise, choosing an unsuitable model architecture could result in the failure to capture crucial features necessary for accurate classification. Consequently, it is imperative to thoroughly consider both aspects during the design of a classification model for 3D UT data to ensure optimal performance.

## 6. CONCLUSION

Deep learning has demonstrated prior success in ultrasonic non-destructive evaluation when applied to either time series or image data. However, analyzing only time series or image data can result in a significant loss of information in either the temporal or spatial domains. This paper proposes the use of 3D convolutional neural networks to classify complete volumetric ultrasound data without compression, retaining all spatial and temporal information. This approach not only reduced the need for accurate gating when constructing C-scan images but also decreased the amount of signal processing required. To train the models, synthetic data was generated from semi-analytical simulations, while experimentally collected ultrasonic responses from manufactured defects were used for testing. Two forms of data augmentation were implemented based on physical variations seen in experimental ultrasonic responses to improve the model's classification performance in the experimental domain. Furthermore, the performance of three different architectures; one existing in the literature, one hand-designed based on current practices, and one designed by

Neural Architecture Search (NAS) from a ResNet search space modified for 3D, were compared.

The first architecture, VoxNet, performed reasonably well on experimental data, achieving a mean F1 score of 0.825.

However, its notable standard deviation in accuracies during training suggests suboptimal architecture optimization for this task. CustomNet's greatly improved on VoxNet with an accuracy increase of 14.8%, whilst reducing the standard deviations in accuracy by 8.9%, hence demonstrating an architecture better optimized for this task.

The third architecture, designed by NAS, when trained with data augmentation, gave the best results, providing 100% classification accuracy. The impact of online domain specific augmentation was notable, leading to a 22.4% decrease in mean accuracy for the NAS model when augmentation was omitted.

Overall, this work demonstrated that it is possible to train successful Deep Learning models to classify full volumetric ultrasonic data for NDE. The issue of a lack of data in most NDE situations was addressed by successfully implementing synthetic data generation in 3D. The work highlighted the importance of appropriate architecture selection and effective data augmentation when translating between synthetic and experimental domains, with both factors essential in achieving high classification accuracy.

The focus of this work was on the use of volumetric datasets, and whilst 100% classification accuracy was achieved through effective NAS, the authors recognize that back drilled holes are generally simple defects to detect by human operators. Future work aims to increase the complexity of the task by detecting a wider range of more challenging defects and expanding the simulation scope to better cover naturally occurring defects, where performance can be measured against human operators in a more realistic industrial scenario. Further work also aims to extend the problem to defect classification and sizing.

## 7. ACKNOWLEDGMENT

This work was supported through Spirit AeroSystems/ Royal Academy of Engineering Research Chair for In-Process Non-Destructive Testing of Composites, RCSR 1920/10/32.

## 8. REFERENCES

- [1] J. Zhang, G. Lin, U. Vaidya, and H. Wang, 'Past, present and future prospective of global carbon fibre composite developments and applications', *Composites Part B: Engineering*, vol. 250, p. 110463, Feb. 2023, doi: 10.1016/j.compositesb.2022.110463.
- [2] B. Djordjevic, 'Non Destructive Test Technology for the Composite', p. 7, Jan. 2009.
- [3] C. Meola, S. Boccardi, G. M. Carlomagno, N. D. Boffa, E. Monaco, and F. Ricci, 'Nondestructive evaluation of carbon fibre reinforced composites with infrared thermography and ultrasonics', *Composite Structures*, vol. 134, pp. 845–853, Dec. 2015, doi: 10.1016/j.compstruct.2015.08.119.
- [4] A. M.-E. Dorado, 'Composite Material Characterization using Acoustic Wave Speed Measurements', p. 5.
- [5] Ley, O. and V. Godinez, 'Non-destructive evaluation (NDE) of aerospace composites: application of infrared (IR) thermography', doi: 10.1533/9780857093554.3.309.
- [6] A. Kokurov and D. Subbotin, 'Ultrasonic detection of manufacturing defects in multilayer composite structures', *IOP Conference Series: Materials Science and Engineering*, vol. 1023, p. 012013, Jan. 2021, doi: 10.1088/1757-899X/1023/1/012013.
- [7] D. K. Hsu, '15 - Non-destructive evaluation (NDE) of aerospace composites: ultrasonic techniques', in *Non-Destructive Evaluation (NDE) of Polymer Matrix Composites*, V. M. Karbhari, Ed., in Woodhead Publishing Series in Composites Science and Engineering. , Woodhead Publishing, 2013, pp. 397–422. doi: 10.1533/9780857093554.3.397.
- [8] F. Heinecke and C. Willberg, 'Manufacturing-Induced Imperfections in Composite Parts Manufactured via Automated Fiber Placement', *J. Compos. Sci.*, vol. 3, no. 2, p. 56, Jun. 2019, doi: 10.3390/jcs3020056.
- [9] I. Papa, V. Lopresto, and A. Langella, 'Ultrasonic inspection of composites materials: Application to detect impact damage', *International Journal of Lightweight Materials and Manufacture*, vol. 4, no. 1, pp. 37–42, Mar. 2021, doi: 10.1016/j.ijlmm.2020.04.002.
- [10] L. Séguin-Charbonneau, J. Walter, L.-D. Thérout, L. Scheed, A. Beausoleil, and B. Masson, 'Automated defect detection for ultrasonic inspection of CFRP aircraft components', *NDT & E International*, vol. 122, p. 102478, Sep. 2021, doi: 10.1016/j.ndteint.2021.102478.
- [11] S. Gholizadeh, 'A review of non-destructive testing methods of composite materials', *Procedia Structural Integrity*, vol. 1, pp. 50–57, 2016, doi: 10.1016/j.prostr.2016.02.008.
- [12] M. Jolly *et al.*, 'Review of Non-destructive Testing (NDT) Techniques and their Applicability to Thick Walled Composites', *Procedia CIRP*, vol. 38, pp. 129–136, Jan. 2015, doi: 10.1016/j.procir.2015.07.043.
- [13] S. Maack, V. Salvador, and S. David, 'Validation of artificial defects for Non-destructive testing measurements on a reference structure', *MATEC Web of Conferences*, vol. 199, p. 06006, Jan. 2018, doi: 10.1051/mateconf/201819906006.
- [14] P. Gardner *et al.*, 'Machine learning at the interface of structural health monitoring and non-destructive evaluation', *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 378, no. 2182, p. 20190581, Oct. 2020, doi: 10.1098/rsta.2019.0581.
- [15] 'Introduction to non-destructive testing', Aerospace Testing International. Accessed: Nov. 17, 2021. [Online]. Available: <https://www.aerospacetestinginternational.com/features/introduction-to-non-destructive-testing.html>
- [16] F. W. Margrave, K. Rigas, D. A. Bradley, and P. Barrowcliffe, 'The use of neural networks in ultrasonic

- flaw detection', *Measurement*, vol. 25, no. 2, pp. 143–154, Mar. 1999, doi: 10.1016/S0263-2241(98)00075-X.
- [17] J. Ye, S. Ito, and N. Toyama, 'Computerized Ultrasonic Imaging Inspection: From Shallow to Deep Learning', *Sensors (Basel)*, vol. 18, no. 11, p. 3820, Nov. 2018, doi: 10.3390/s18113820.
- [18] B. Valeske, A. Osman, F. Römer, and R. Tschuncky, 'Next Generation NDE Sensor Systems as IIoT Elements of Industry 4.0', *Research in Nondestructive Evaluation*, vol. 31, no. 5–6, pp. 340–369, Nov. 2020, doi: 10.1080/09349847.2020.1841862.
- [19] S. Cantero-Chinchilla, P. D. Wilcox, and A. J. Croxford, 'Deep learning in automated ultrasonic NDE -- developments, axioms and opportunities', *arXiv:2112.06650 [eess]*, Dec. 2021, Accessed: Jan. 12, 2022. [Online]. Available: <http://arxiv.org/abs/2112.06650>
- [20] A. Figueira and B. Vaz, 'Survey on Synthetic Data Generation, Evaluation Methods and GANs', *Mathematics*, vol. 10, no. 15, Art. no. 15, Jan. 2022, doi: 10.3390/math10152733.
- [21] S. McKnight *et al.*, 'GANs and alternative methods of synthetic noise generation for domain adaption of defect classification of Non-destructive ultrasonic testing'. arXiv, Jun. 02, 2023. doi: 10.48550/arXiv.2306.01469.
- [22] S. Lonne, S. Mahaut, and G. Cattiaux, 'EXPERIMENTAL VALIDATION OF CIVA ULTRASONIC SIMULATIONS', Jan. 2006.
- [23] M. Darmon *et al.*, 'VALIDATION OF AN ULTRASONIC CHARACTERIZATION TECHNIQUE FOR ANISOTROPIC MATERIALS: COMPARISON OF EXPERIMENTS WITH BEAM PROPAGATION MODELLING', p. 20.
- [24] K. Jezzine, D. Ségur, R. Ecault, and N. Dominguez, 'Simulation of ultrasonic inspections of composite structures in the CIVA software platform', p. 8.
- [25] N. Munir, H.-J. Kim, J. Park, S.-J. Song, and S.-S. Kang, 'Convolutional neural network for ultrasonic weldment flaw classification in noisy conditions', *Ultrasonics*, vol. 94, pp. 74–81, Apr. 2019, doi: 10.1016/j.ultras.2018.12.001.
- [26] P. Zacharis, G. West, G. Dobie, T. Lardner, and A. Gachagan, 'Data-Driven Analysis of Ultrasonic Inspection Data of Pressure Tubes', *Nuclear Technology*, vol. 202, no. 2–3, pp. 153–160, Jun. 2018, doi: 10.1080/00295450.2017.1421803.
- [27] M. Meng, Y. J. Chua, E. Wouterson, and C. P. K. Ong, 'Ultrasonic signal classification and imaging system for composite materials via deep convolutional neural networks', *Neurocomputing*, vol. 257, pp. 128–135, Sep. 2017, doi: 10.1016/j.neucom.2016.11.066.
- [28] B. Wang, Y. Li, Y. Luo, X. Li, and T. Freiheit, 'Early Event Detection in a Deep-learning Driven Quality Prediction Model for Ultrasonic Welding', *Journal of Manufacturing Systems*, vol. 2021, pp. 325–336, Jun. 2021, doi: 10.1016/j.jmsy.2021.06.009.
- [29] J. Ye and N. Toyama, 'Benchmarking Deep Learning Models for Automatic Ultrasonic Imaging Inspection', *IEEE Access*, vol. 9, pp. 36986–36994, 2021, doi: 10.1109/ACCESS.2021.3062860.
- [30] R. J. Pyle, R. L. T. Bevan, R. R. Hughes, R. K. Rachev, A. A. S. Ali, and P. D. Wilcox, 'Deep Learning for Ultrasonic Crack Characterization in NDE', *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 68, no. 5, pp. 1854–1865, May 2021, doi: 10.1109/TUFFC.2020.3045847.
- [31] D. Medak, L. Posilovic, M. Subasic, M. Budimir, and S. Loncaric, 'Automated Defect Detection From Ultrasonic Images Using Deep Learning', *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 68, no. 10, pp. 3126–3134, Oct. 2021, doi: 10.1109/TUFFC.2021.3081750.
- [32] T. B. Hudson, P. J. Follis, J. J. Pinakidis, T. Sreekantamurthy, and F. L. Palmieri, 'Porosity detection and localization during composite cure inside an autoclave using ultrasonic inspection', *Composites Part A: Applied Science and Manufacturing*, vol. 147, p. 106337, Aug. 2021, doi: 10.1016/j.compositesa.2021.106337.
- [33] Institute of Fundamentals of Machinery Design, Silesian University of Technology, Gliwice, 44-100, Poland, A. Wronkiewicz, A. Katunin, and K. Dragan, 'Ultrasonic C-Scan Image Processing Using Multilevel Thresholding for Damage Evaluation in Aircraft Vertical Stabilizer', *IJIGSP*, vol. 7, no. 11, pp. 1–8, Oct. 2015, doi: 10.5815/ijigsp.2015.11.01.
- [34] Z. Zhang, M. Liu, Q. Li, and Y. Ang, 'Visualized characterization of diversified defects in thick aerospace composites using ultrasonic B-scan', *Composites Communications*, vol. 22, p. 100435, Dec. 2020, doi: 10.1016/j.coco.2020.100435.
- [35] Y. Zhou *et al.*, 'Multi-task learning for segmentation and classification of tumors in 3D automated breast ultrasound images', *Medical Image Analysis*, vol. 70, p. 101918, May 2021, doi: 10.1016/j.media.2020.101918.
- [36] Y. Liu, '3D Image Segmentation of MRI Prostate Based on a Pytorch Implementation of V-Net', *J. Phys.: Conf. Ser.*, vol. 1549, no. 4, p. 042074, Jun. 2020, doi: 10.1088/1742-6596/1549/4/042074.
- [37] M. Leong, D. Prasad, Y. T. Lee, and F. Lin, 'Semi-CNN Architecture for Effective Spatio-Temporal Learning in Action Recognition', *Applied Sciences*, vol. 10, p. 557, Jan. 2020, doi: 10.3390/app10020557.
- [38] J. Chai, H. Zeng, A. Li, and E. W. T. Ngai, 'Deep learning in computer vision: A critical review of emerging techniques and application scenarios', *Machine Learning with Applications*, vol. 6, p. 100134, Dec. 2021, doi: 10.1016/j.mlwa.2021.100134.
- [39] A. M. Badshah *et al.*, 'Deep features-based speech emotion recognition for smart affective services', *Multimed Tools Appl*, vol. 78, no. 5, pp. 5571–5589, Mar. 2019, doi: 10.1007/s11042-017-5292-7.
- [40] H. Liu, K. Simonyan, and Y. Yang, 'DARTS: Differentiable Architecture Search'. arXiv, Apr. 23, 2019. doi: 10.48550/arXiv.1806.09055.

- [41] J. Mellor, J. Turner, A. Storkey, and E. J. Crowley, 'Neural Architecture Search without Training'. arXiv, Jun. 11, 2021. doi: 10.48550/arXiv.2006.04647.
- [42] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, 'Algorithms for Hyper-Parameter Optimization', in *Advances in Neural Information Processing Systems*, Curran Associates, Inc., 2011. Accessed: Jun. 16, 2023. [Online]. Available: [https://papers.nips.cc/paper\\_files/paper/2011/hash/86e8f7ab32cfd12577bc2619bc635690-Abstract.html](https://papers.nips.cc/paper_files/paper/2011/hash/86e8f7ab32cfd12577bc2619bc635690-Abstract.html)
- [43] S. Schrodi, D. Stoll, B. Ru, R. Sukthanker, T. Brox, and F. Hutter, 'Construction of Hierarchical Neural Architecture Search Spaces based on Context-free Grammars'. arXiv, Jun. 05, 2023. Accessed: Jun. 16, 2023. [Online]. Available: <http://arxiv.org/abs/2211.01842>
- [44] D. Maturana and S. Scherer, 'VoxNet: A 3D Convolutional Neural Network for real-time object recognition', in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2015, pp. 922–928. doi: 10.1109/IROS.2015.7353481.
- [45] Z. Liu *et al.*, 'VB-Net: Voxel-Based Broad Learning Network for 3D Object Classification', *Applied Sciences*, vol. 10, no. 19, Art. no. 19, Jan. 2020, doi: 10.3390/app10196735.
- [46] Y. Fu and X. Yao, 'A review on manufacturing defects and their detection of fiber reinforced resin matrix composites', *Composites Part C: Open Access*, vol. 8, p. 100276, Jul. 2022, doi: 10.1016/j.jcomc.2022.100276.
- [47] 'RollerFORM: Phased Array Wheel Probe'. Accessed: Jun. 06, 2023. [Online]. Available: <https://www.olympus-ims.com/en/rollerform/>
- [48] M. Vasilev *et al.*, 'Sensor-Enabled Multi-Robot System for Automated Welding and In-Process Ultrasonic NDE', *Sensors*, vol. 21, no. 15, Art. no. 15, Jan. 2021, doi: 10.3390/s21155077.
- [49] C. Mineo *et al.*, 'Flexible integration of robotics, ultrasonics and metrology for the inspection of aerospace components', *AIP Conference Proceedings*, vol. 1806, no. 1, p. 020026, Feb. 2017, doi: 10.1063/1.4974567.
- [50] 'EXTENDE, Experts in Non Destructive Testing Simulation with CIVA Software'. Accessed: Nov. 07, 2022. [Online]. Available: <https://www.extende.com/>
- [51] D. P. Kingma and J. Ba, 'Adam: A Method for Stochastic Optimization'. arXiv, Jan. 29, 2017. doi: 10.48550/arXiv.1412.6980.
- [52] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, 'A ConvNet for the 2020s'. arXiv, Mar. 02, 2022. Accessed: Jun. 16, 2023. [Online]. Available: <http://arxiv.org/abs/2201.03545>