Automated classification of subsurface impact damage in thermoplastic composites using depth-resolved terahertz imaging and deep learning

Dicky J. Silitonga [a,b,*], Pascal Pomarède [b], Niyem M. Bawana [c], Haolian Shi [a,c], Nico F. Declercq [a,d], D.S. Citrin [a,c], Fodil Meraghni [b], Alexandre Locquet [a,c]

[a] Georgia Tech-CNRS IRL2958, Georgia Tech–Europe, 2 Rue Marconi, Metz, 57070, France
[b] Arts et Métiers Institute of Technology, CNRS, Université de Lorraine, LEM3-UMR 7239 CNRS, 4 rue Augustin Fresnel, Metz, 57078, France
[c] School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, 30332-0250, USA
[d] George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA, 30332-0250, USA

A B S T R A C T

Reliable detection of barely visible impact damage is critical to ensure the structural integrity of composite components in service, particularly in safety-critical applications such as pressure vessels and transportation systems. This study presents a solution for detecting such damage in woven glass fiber-reinforced thermoplastic composites using terahertz (THz) time-of-flight tomography and convolutional neural networks. THz provides non-contact, non-ionizing, high-axial-resolution imaging of subsurface and back-surface damage, addressing key limitations of surface-based inspection methods. While THz imaging alone may not always permit conclusive damage identification, we bridge this gap by training neural network classifiers on depth-resolved THz B-scan images using ground truth from co-located X-ray micro-computed tomography. Among several pretrained architectures tested via transfer learning, DenseNet-121 exhibits the highest accuracy. The model remains robust even when trained on truncated B-scans excluding surface indentation features, confirming its ability to detect structural anomalies located internally or on the back surface. This is particularly relevant for applications where back-side access is not feasible. Experimental validation is performed on impacted glass-fiber-reinforced thermoplastic coupons prepared in accordance with ASTM D7136, with damage severity quantified through force–displacement data and micro-tomographic analysis. Labeling for supervised learning conforms to acceptance criteria from industrial standards for composite pressure vessels (ASME BPVC Section X, CGA C-6.2), ensuring regulatory alignment and enabling deployment in quality control workflows. The proposed method minimizes the need for expert interpretation or secondary validation and offers direct applicability to in-service inspection and manufacturing quality control.

1. Introduction

Terahertz (THz) time-of-flight tomography (TOFT), also known as THz pulsed imaging, has emerged as a promising nondestructive evaluation (NDE) technique for assessing damage in composite materials. TOFT offers high axial resolution (∼10 μm), thanks to its sub-picosecond time resolution and broad spectral bandwidth (>1 THz), making it well suited for surface and near-surface characterization of multilayer materials. Compared to ultrasound, which typically offers greater penetration but lower spatial resolution, THz imaging is contact-free and does not require coupling media. In contrast to X-ray imaging, THz employs nonionizing radiation, offering a safer alternative for both operators and sensitive composite structures. These features position THz TOFT as a powerful tool in NDE, especially where high-resolution depth profiling is critical [1].

Previous studies have demonstrated the utility of THz TOFT in detecting and characterizing failure modes in composite specimens, including delamination [2,3], fiber breakage [4], and matrix cracking [5,6]. Impact damage presents a complex condition as it can involve simultaneous presence of those failure modes [7]. Moreover, in low-velocity impact, damage may be subtle and easily obscured by the morphology of woven fibers in the composite itself. Consequently, the interpretation of THz TOFT measurements typically requires cross-validation with complementary techniques, such as X-ray micro-computed tomography (μCT), to definitively confirm the presence of damage and identify its failure modes [8].

A method that enables standalone interpretation of THz images, without reliance on μCT or expert judgement, is essential for scalable,

* Corresponding author at: Georgia Tech-CNRS IRL2958, Georgia Tech–Europe, 2 Rue Marconi, Metz, 57070, France.
E-mail addresses: djsilitonga@gmail.com, dicky.silitonga@ensam.eu (D.J. Silitonga).

in-service NDE. While initial validation against μCT is necessary, the ultimate goal is an autonomous system suitable for real-time inspection.

To address those issues, this study integrates deep-learning (DL) techniques with THz TOFT to enable automated detection of low-velocity impact-induced damage in glass-fiber-reinforced polymer (GFRP) laminates. Low-velocity impact can induce barely visible impact damage (BVID), which may be accompanied by subsurface damage compromising structural integrity without obvious surface indications, making reliable detection critical for ensuring long-term performance and safety. Undetected damage of this type can act as an initiation site for delamination growth, stiffness degradation, or buckling under cyclic loading, which are critical failure modes in structural composite applications. A convolutional neural network (CNN), is employed to perform binary classification of THz B-scans, distinguishing between damaged and undamaged specimens. This approach facilitates machine-vision-based inspection and reduces dependence on expert interpretation.

Although the use of artificial intelligence (AI) in nondestructive evaluation techniques for detecting BVID in composites has gained traction, important limitations remain. For instance, Hasebe et al. [9] used surface profilometry, and Tabatabaeian et al. [10] trained DL models on surface photographs. In both cases, predictions are based solely on the condition of the exposed surface, without access to internal or back-surface features. However, since impact damage in laminated composites often initiates on the back surface and propagates towards the front (impacted) face, incorporating subsurface information is essential for reliable classification. THz TOFT addresses this gap by capturing depth-resolved views of internal and back-side conditions. This through-thickness visibility is particularly advantageous in applications where direct access to the rear surface is not feasible, such as the interior of pressure vessels.

While previous studies have applied machine learning (ML) to THz A-scans [11–13] or C-scans [14] for delamination or debonding detection, these approaches often miss complex, multi-depth impact damage. This work addresses that gap by introducing a B-scan-based DL framework to detect structural damage beyond what planar or single-point scans can resolve.

Impact damage is particularly challenging to identify due to its distributed and multi-layered nature. A-scans, representing single-point time-domain signals, often fail to reveal such complex patterns, while C-scans provide a planar view of defects that requires the preselection of specific time delays and therefore lack flexibility for automated inspection workflows.

In this study, THz B-scans are employed as the basis for training the CNN model to classify damaged and undamaged composite. Unlike A-scans, B-scans offer a depth-resolved cross-sectional view of the interior of the sample, capturing structural and morphological variations both in the axial and transverse directions. While DL-based classification using B-scans has been explored with ultrasonic imaging datasets [15,16], the present work leverages the advantages of THz TOFT to generate images with enhanced feature visibility. This approach is expected to improve the accuracy and reliability of the damage detection classifier.

The present study concentrates on binary classification, providing a proof of concept for deep learning–based THz inspection to differentiate undamaged from damaged specimens. Although this task may appear simple, it has immediate industrial significance for pass–fail inspection [17,18], where decisions rely primarily on the detection of damage or structural anomalies [19], without requiring identification of the precise defect type or morphology. The binary framework presented here forms a foundation for subsequent developments towards multiclass classification, including estimation of damage depth and recognition of defect types. At this stage, model training is performed on a dataset derived from a single composite material, GFRP/polyamide 66 (PA66), to facilitate controlled dataset generation. Future work will extend this framework by transferring the learned representations and fine-tuning the model on larger, more diverse datasets, thereby enabling broader applicability and more advanced classification capabilities.

The proposed workflow involves scanning composite specimens using THz imaging, comparing these scans with X-ray μCT images as ground truth to accurately label the data, and subsequently creating a large, annotated dataset for training. For the binary classification task, DenseNet121 [20], a CNN pretrained on ImageNet [21], is employed. The B-scan images are labeled by comparison with X-ray μCT data, which serve as ground truth. Transfer learning is used to adapt the model to the domain-specific task of THz image classification, thereby reducing training time and data requirements. The objective is to establish a framework capable of performing binary classification, damaged versus undamaged, that could be integrated into inspection protocols for composite components. This has direct relevance to manufacturing quality control and in-service inspection of composite structures such as hydrogen storage pressure vessels.

When aligned with industrial standards such as ASME BPVC Section X (Fiber-Reinforced Plastic Pressure Vessels) [22] and CGA C-6.2 (Standard for Visual Inspection and Requalification of Fiber Reinforced High Pressure Cylinders) [23], the proposed method offers a scalable and automatable solution for structural health assessment. It supports integration into engineering inspection workflows, including periodic requalification and life-cycle management of safety-critical composite components. In this context, damage annotations used for model training adhere to the acceptance and rejection criteria defined in these standards, ensuring that the classifier's outputs are consistent with codified inspection practices.

Within the field of composite-materials engineering, AI has been employed for diverse applications, including material design and the analysis of response under loading. For instance, a ML-based tool has been developed to assist material designers in exploring various laminate configurations and enable accurate estimation of material properties through virtual testing [24]. Another example is the application of multiscale thermodynamics-informed neural networks (MuTINN) to facilitate the study of the complex behavior of composites subjected to different loading conditions [25]. The integration of DL into NDE technique for damage inspection, as demonstrated in this paper, further expands the role of AI across a broad range of aspects in the composite material lifecycle.

The remainder of the paper is organized as follows. In Section 2 we discuss the THz TOFT procedure, including specimen preparation, scanning setup, signal processing, and the selection of B-scan images for training. Section 3 contains the presentation of the DL framework, including the use of transfer learning with DenseNet and our modification applied to it, as well as training methodology. In Section 4, we describe the dataset generation process, including labeling based on X-ray μCT and image augmentation strategies. Section 5 presents and discusses the classification results. Finally, Section 6 concludes this paper by summarizing the findings and outlining the limitations and future directions of this work.

2. Experimental data acquisition

B-scans were acquired from a series of GFRP laminate coupons, some of which were subjected to low-velocity impact, using THz TOFT in reflection mode raster scanning across regions of interest centered around the impact site, recording time-resolved waveforms at each spatial coordinate. The raw time-domain signals were preprocessed using bandpass filtering and wavelet denoising to enhance signal-to-noise ratio. B-scans were then reconstructed, offering depth-resolved cross-sectional representations of the internal laminate structure. Each B-scan was annotated by comparison with colocated X-ray μCT scans, which served as the ground truth to verify internal damage presence and enable supervised training of the classifier. The X-ray μCT investigation were carried-out with an EasyTom (Nano) from RX solutions. A voxel resolution of 0.015 mm was achieved.

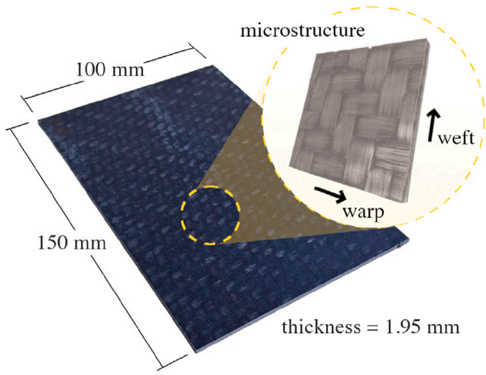

Fig. 1. Photograph of a GFRP laminate coupon prior to impact testing.

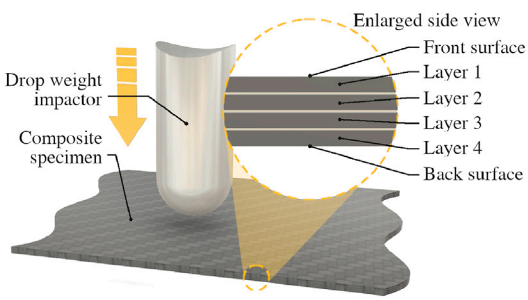

Fig. 2. Schematic of drop-weight impact setup and ply numbering convention.

GFRP laminates were prepared as standard impact-test coupons with dimension 150 mm × 100 mm (6" × 4") in accordance with ASTM D7136: Standard Test Method for Measuring the Damage Resistance of a Fiber-Reinforced Polymer Matrix Composite to a Drop-Weight Impact Event [26]. The material studied consists of a reinforced semi-crystalline thermoplastic polymer (Evolite XA1481) supplied by Solvay. According to manufacturer specifications, the composite consists of a 2 × 2 twill-woven glass-fiber fabric embedded in a polyamide 66 resin matrix, with a nominal fiber content of 50% by volume (70% by weight). The tested coupon have been cut from 500 mm × 1200 mm composites plates manufactured by thermo-compression process. Each laminate comprises four plies, arranged in a balanced stacking sequence $[0°/90°]_2$, with a total nominal thickness of 1.95 mm. A representative coupon is shown in Fig. 1.

2.1. Specimen preparation

To mitigate the influence of moisture on the mechanical response, all specimens were preconditioned to reach equilibrium moisture content before testing, following the procedure outlined in [27]. The conditioning was performed in a relative-humidity chamber (Memmert HCP 246) at 50% relative humidity (RH 50%) and at a temperature of 65 °C. The mass evolution, induced by the water uptake, of each coupon was recorded daily. The conditioning process was deemed complete when weight variation remained below a defined threshold for three consecutive days, indicating moisture saturation.

Once conditioned, the specimens were subjected to low-velocity impact using a drop-weight impact testing system (Instron CEAST 9350) at room temperature. In our tests, the setup uses a drop weight of 5 kg and a tup (impactor) diameter of 16 mm. The machine features an antirebound system that prohibits multiple impacts on the sample. For consistency, the face impacted by the drop weight is designated as the front surface, and plies are numbered from the front (Layer 1) to the back (Layer 4). The impact configuration and layer numbering convention are illustrated in Fig. 2.

Table 1
Impact test parameters and material specifications of the samples. The impactor's mass and diameter are 5 kg and 16 mm, respectively.

Sample #	Drop height (mm)	Impact velocity (m/s)	Real impact energy (J)	Composite type	Thick. (mm)
1	60	1.27	4.39	GFRP, PA 66 matrix 50% vol Glass Fiber woven fabric, 2 × 2 twill, 4 plies ($[0°/90°]_2$)	1.95 mm
2	81	1.42	5.42		
3	101	1.51	6.10		
4	121	1.64	7.16		
5	141	1.76	8.21		
6	161	1.87	9.24		
7	202	2.07	11.15		

The selection of impact energies was guided by preliminary trials aimed at identifying a range that would reliably produce damage consistent with the BVID category. An initial impact test at 25 J, selected as a first approximation of the laminate's impact resistance, resulted in extensive damage far exceeding the BVID definition. The energy was then reduced to 7.16 J, which produced only a minute indication of fiber disruption on the back surface, discernible only under favorable viewing conditions. From this level, the impact energy was increased in increments of approximately 1 J, up to 11 J, at which point the damage became clearly visible. Additional step-down tests were subsequently conducted from 7.16 J in similar increments, down to 4.39 J, where even the front-surface indentation was barely perceptible under any viewing conditions.

In the experiments, the nominal impact energy was prescribed on the drop-weight system by specifying the drop height and impactor mass. The values reported in this paper correspond to the actual energies, obtained by integrating the area under the recorded force–displacement curve for each test. To confirm that the impacted specimens indeed contained BVID, which by definition is not readily observable to the naked eye, all samples were examined using X-ray μCT. As summarized later in Table 4, every impacted specimen exhibited subsurface damage qualifying as BVID, with the exception of the lowest-energy case (4.39 J), which produced only a slight surface indentation without internal structural compromise.

With the systematic iteration approach described above, seven impact energies were selected to generate dataset of expectedly damaged samples with varying severity levels in a series of coupons. Those impact energy levels were 4.39, 5.42, 6.10, 7.16, 8.21, 9.24, and 11.15 J. The corresponding calculated areas of the force–displacement curve, obtained from the drop weight test system, are indicated as shaded regions under each series plot in Fig. 3(a).

The peak force upon sample deformation increases as a function of impact energy, except that the maximum force reached at 9.24 J is comparable to that of 8.21 J although with a steeper rate (Fig. 3(b)) and greater deformation (Fig. 3(c)). This can be attributed to the onset of damage, such as fiber breakage or matrix cracking, leading to the yielding of a certain part/layer of the laminates. A similar indication is evident in the 11.15 J graph as a slight drop of force around 2350 N, before continuing to increase as the remaining undamaged laminates bear the progressing load up to the peak.

Table 1 summarizes the parameters of the drop-weight impact experiments, along with the specifications of the composite specimens tested.

2.2. THz TOFT

A commercial Terahertz Time Domain Spectroscopy (THz-TDS) system (TeraPulse Lx, TeraView Ltd., Cambridge, UK) was employed for the THz TOFT experiments to perform raster scanning in reflection mode. The system uses an ultrafast near-infrared fiber laser to generate broadband, pulsed THz radiation spanning ~0.1–6 THz. Quasi-single-cycle THz pulses are emitted from a photoconductive antenna (PCA)

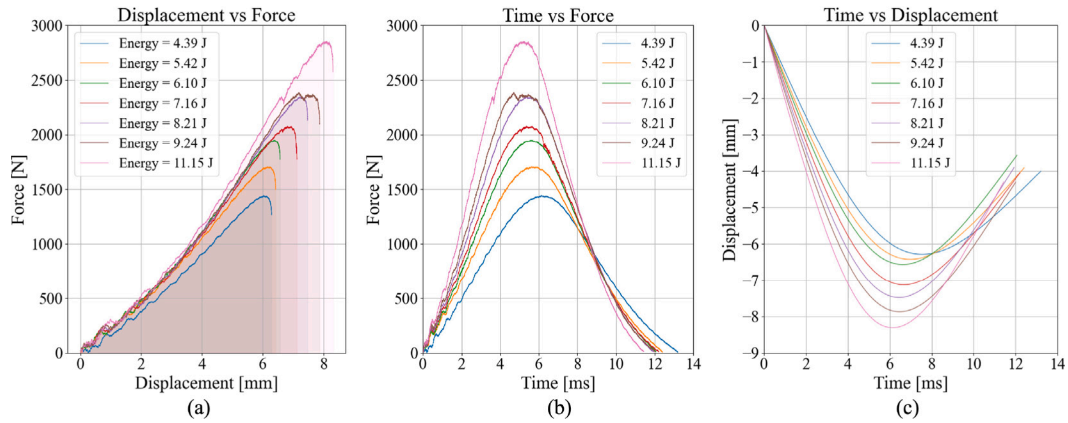

Fig. 3. (a) Force–displacement curves with shading under each plot to indicate the area integrated for calculating the impact energy; Time evolution of (b) force and (c) displacement during impact.

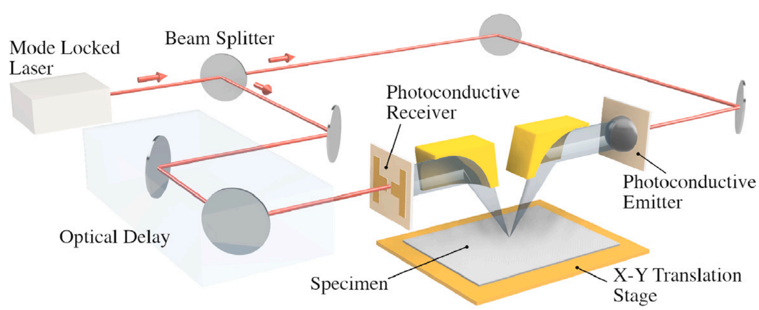

Fig. 4. Schematic diagram of the THz system for THz TOFT.

Table 2
System parameters of the THz-TDS employed in the current work.

Parameter	Value
Waveform rate	15 Hz
Sampling interval	0.02 ps
Laser pulse duration	sub-100 fs
Laser wavelength	780 nm
Power of THz radiation	~1 μW

and detected by another antenna. The system directs the THz pulses onto the sample at near-normal incidence, and the reflected signals are captured and analyzed for both magnitude and phase using a time-gated optical probe. A schematic diagram of the system is shown in Fig. 4 and the system parameters are presented in Table 2.

Prior to measurement, a reference signal was acquired by recording the reflection from a bare metal disk, which served as a reflector across the THz band. The reference signal thus was proportional to the signal produced by the apparatus.

Raster scanning was conducted via a motorized (x, y) translation stage. During all scans, the measurement chamber was purged with N_2 gas to suppress water-vapor absorption along the THz beam path. The scanning process generated a three-dimensional dataset comprising x- and y-coordinates and time delay, forming a matrix that encapsulates both spatial and temporal information. This dataset can be post-processed to extract A-scans at individual points, as well as to generate B- (cross-sectional) and C-scans (planar images) using amplitude contrast mechanisms, as detailed in Section 2.5.

All specimens were scanned from the impacted side in reflection mode. Scanning from the impacted side was chosen as it represents more likely real-world inspection scenarios, such as in the case of pressure vessels, where the impact typically originates from the external direction, and the outer surface is more accessible for inspection.

The scanned area covered a 30 mm × 30 mm square region centered around the indentation site (for impacted specimens), with a transverse spatial resolution (step size) of 0.3 mm along both x- and y-axes set on the scanner controller software. To enhance signal quality and reduce noise, each spatial point acquisition was averaged over five THz pulses.

2.3. Rationale for selecting THz imaging for dataset generation

The choice of THz imaging as the basis for the working dataset is justified through quantitative and qualitative comparisons with X-ray μCT and ultrasonic imaging. These comparisons are summarized in Fig. 5 and Table 3. In particular, ultrasonic inspection at the high frequencies required to obtain detailed stratigraphic information in GFRP laminates is subject to strong attenuation, which limits the retrieval of reliable subsurface features [28].

Measurements were performed on a 30 mm × 30 mm area of an impacted specimen (7.16 J) to provide a consistent basis for comparison. Observations indicate that ultrasonic imaging provides rapid scans but limited subsurface detail as noted above, whereas X-ray μCT reveals fine internal structures at the expense of significantly longer measurement duration. Typically, the measurement time provides an indirect estimate of man-hour cost comparison, although the actual rate also depends on the operator's competence level and other factors. THz imaging occupies an intermediate position with moderate resolution and clear indications of subsurface features, including layer layup and damage cues.

Unlike X-ray μCT, where axial (depth) resolution is defined by the measurement voxel size (0.015 mm), the axial resolution of both THz and ultrasonic imaging is not straightforward to quantify, as it depends on temporal signal peak separation, material properties, morphology, and signal quality. *A priori*, the axial resolution of the three methods employed in this work, from highest to lowest, is X-ray μCT, followed by THz and ultrasonic imaging [29]. In the current case, the specimen

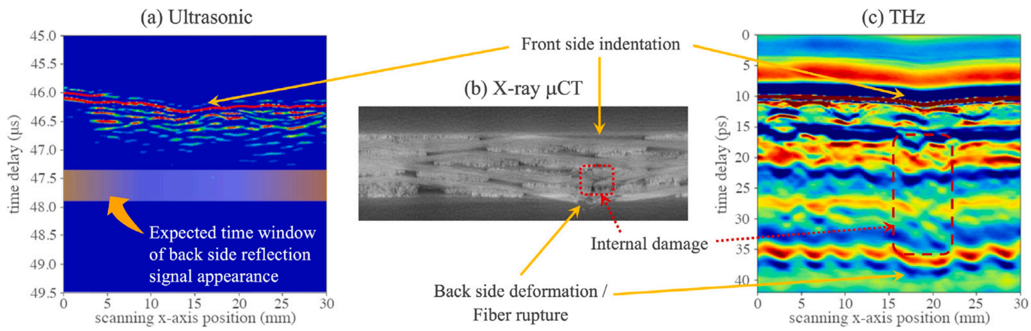

Fig. 5. Cross-sectional views (B-scan) of the specimen impacted at 7.16 J obtained using different techniques: (a) Ultrasonic, (b) X-ray μCT, and (c) THz. Time axes in the ultrasonic and THz images are adjusted for optimal visibility.

Table 3
Measurement characteristics observed on experiments related to Fig. 5.

Characteristics	Ultrasonic	X-ray μCT	THz
Measurement time	12 min	110 min	31 min
Lateral resolution	0.3 mm	0.015 mm	0.3 mm
Cross-sectional visibility	Layup unresolved; backside signal absent	High resolution (0.015 mm); internal damage clearly visible	Layup resolved; internal damage cues indicated

has a nominal thickness of 2 mm composed of four layers, giving an average layer thickness of 0.5 mm. Within this spatial dimension, THz imaging clearly visualizes the layup structures as shown in Fig. 5(c) [see also Fig. 8(a) and (b)].

In contrast, the ultrasonic B-scan image obtained from our measurement cannot resolve the layup, and the backside signal is absent due to attenuation, as exhibited in Fig. 5(a). As is commonly understood, attenuation depends on frequency: higher frequencies improve axial resolution but suffer greater signal loss, while lower ones increase penetration depth at the cost of resolution due to pulse overlap. A 30 MHz immersion transducer was used in this study to produce images with best contrast after experimenting with lower frequencies scans down to 1 MHz, which resulted in poorer amplitude contrasts (results not shown here for conciseness).

X-ray μCT provides the highest internal detail but requires significantly longer measurement times, more complex instrumentation, and additional operational safety measures. These observations demonstrate that THz imaging is an effective technique, offering non-contact operation, moderate lateral resolution, and sensitivity to subsurface features. Nevertheless, owing to its detailed representations, X-ray μCT images, like in Fig. 5(b), were used as reference or ground truth to verify the THz imaging results during the labeling procedure, as will be described in the subsequent sections.

2.4. THz signal processing

To improve signal quality and suppress artifacts, raw THz TOFT time-domain signals underwent preprocessing prior to image dataset generation. The signal processing pipeline is illustrated in Fig. 6 for a representative THz A-scan.

In the first step of the signal processing, the raw signal shown in Fig. 6(a) is subjected to Gaussian bandpass filtering between cutoff frequencies of 0.5 and 1.5 THz, with a standard deviation of 0.05 THz. In this study, THz signal components below 0.5 THz correspond to wavelengths greater than 300 μm in the material. Such long-wavelength components reduce axial resolution and can obscure features at smaller dimensional scales [30], and were therefore removed. At the higher end of the spectrum, THz-TDS systems based on photoconductive antennas (PCA) exhibit decreasing detection efficiency [31,32], with the measured high-frequency content dominated by noise rather than useful information. Although the manufacturer specifies a nominal bandwidth up to 6 THz, frequencies above 1.5 THz in our experiments displayed low signal-to-noise ratios. This portion of the spectrum was therefore excluded to avoid noise overwhelming meaningful structural features.

The band pass filtering operation enhances the definition of primary peaks by removing low-frequency baseline drift and high-frequency noise, as evident in Fig. 6(b). Nevertheless, small-amplitude fluctuations, such as the residual ripples between 2–5 ps, persist in the filtered signal. The ripples introduced random features unrelated to the physical state of the material, with inconsistent amplitudes and temporal positions across signals. Such inconsistencies complicate the pattern recognition process of the deep learning model and therefore needed to be mitigated to ensure that only peaks corresponding to physical features were emphasized. The denoising step not only improved the smoothness of the reconstructed images but also supported the classifier in more reliably identifying structural patterns in the investigated specimens.

To further improve feature clarity, wavelet denoising is applied. By projecting the signal onto a wavelet basis that resembles the reference signal, this step attenuates incoherent components while preserving sharp transients such as reflection peaks. The wavelet denoising uses the stationary wavelet transform with a 4th-order Symlet wavelet (sym4). The resulting waveform is shown in Fig. 6(c).

2.5. Selection of B-scans for training dataset generation

From the processed THz TOFT data, images can be constructed to enable inspection in either transverse (C-scan) or cross-sectional (B-scan) views. The complex architecture of the woven GFRP, with its inherent spatial nonuniformities across multiple layers, along with the irregular morphology of impact damage, renders single-point inspection (A-scan) inadequate. C-scans are widely used in NDE to map surface or subsurface features at fixed depth slices and are effective in highlighting damage-affected areas.

In the case of woven composite structures, C-scans facilitate the analysis of layer-wise features, including not only cracks or punctures, but also more complex phenomena such as fiber orientation [33] and internal laminate deformation under lateral tensile stress [34]. C-scan imaging, however, requires the selection of a specific contrast mechanism, typically a fixed time delay, to visualize features at a certain depth. For example, Fig. 7(b) shows a C-scan constructed at a time delay corresponding to the front surface reflection, clearly revealing the surface indentation caused by impact, as seen in the photograph in Fig. 7(a). Nevertheless, a C-scan constructed at another arbitrarily selected time delay, Fig. 7(c), may fail to indicate the presence of damage. Selecting an appropriate time delay typically involves scanning through the entire time axis or defining a time window, which is inefficient and unreliable when the damage spans multiple depths within the laminate.

Determining the optimal time delay in terahertz C-scan imaging is an iterative process. An initial estimate can be obtained by selecting

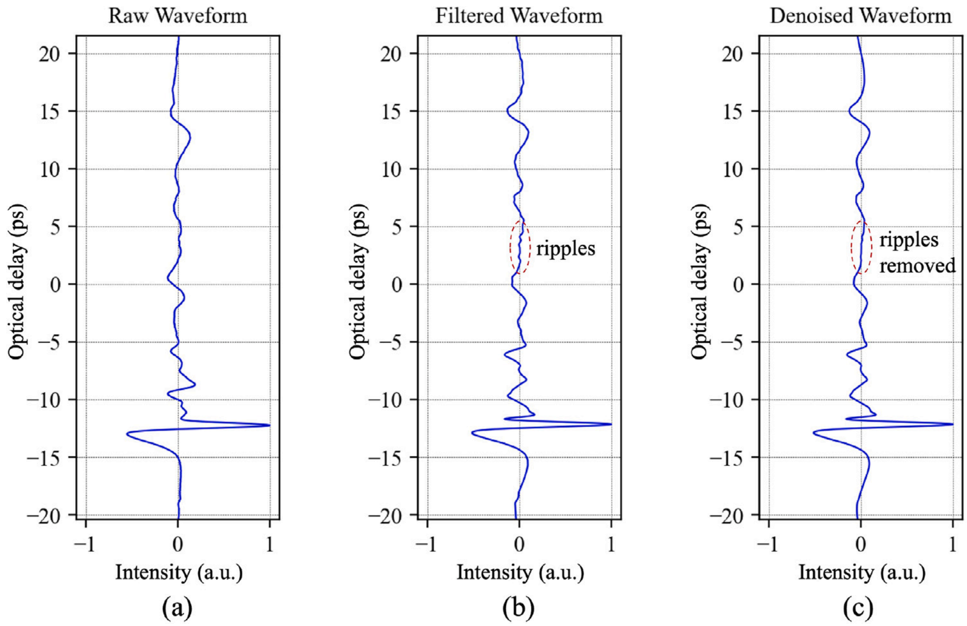

Fig. 6. THz TOFT A-scan signal processing: (a) raw signal; (b) after bandpass filtering (0.1–1.5 THz); (c) after wavelet denoising.

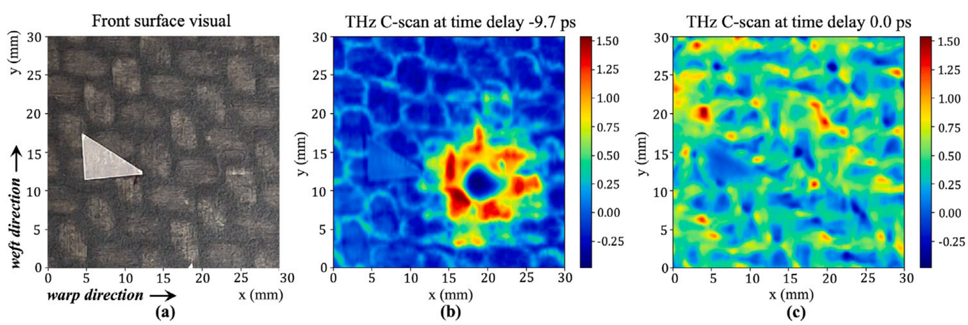

Fig. 7. (a) Photograph of impacted surface of a coupon subjected to low velocity impact. THz C-scan at (b) time delay corresponding to the front surface (−9.7 ps) clearly shows the damage, but at (c) a suboptimal time delay (e.g., 0 ps) the damage is not visible, illustrating C-scan limitations in depth discrimination.

time stamps corresponding to peaks in a representative A-scan, but refinement around this approximation is required to identify the slice that provides maximum image contrast for revealing features or defects. The definition of optimum contrast in C-scan imaging is partly subjective, depending on the inspector's judgement; the principal criterion is that the resulting image clearly highlights the relevant features. These challenges in consistently extracting the most informative slices limit the practicality of C-scan imaging for dataset generation in the present study.

B-scans offer solution to this issue by providing depth-resolved cross-sectional views, capturing both axial and transverse structural variations. This makes them particularly effective for analyzing distributed damage patterns, enabling simultaneous geometric and material assessment in a format well suited for ML input. B-scans constructed over a sufficient spatial extent are more informative, capturing the overall pattern of material response rather than isolated features from A-scan and C-scan.

In this work, B-scans were extracted along single spatial axes near the impact site, with offsets selected to maximize coverage of the damaged region. Each B-scan was manually compared to a co-located X-ray μCT slice for annotation.

Unlike conventional fixed-plane B-scan analysis, we introduced lateral and angular slicing to capture the anisotropic character of damage and diversify the training set. This enabled the model to learn invariant features in spatial orientations and diversify the training set. To enrich the data set, additional B-scans were generated by lateral slicing (offsets up to ± 2 mm) and angular slicing (rotation around the impact center). This diversified the dataset by introducing different perspectives of the same damage site, thereby improving the model's ability to generalize across variations in damage morphology and orientation.

The dimensional accuracy of B-scans is illustrated in Fig. 8, which shows the correspondence between a THz B-scan and a colocated X-ray μCT image. Layer thicknesses can be estimated from the time delay Δt using

$$d = \frac{\Delta t}{2} \frac{c}{n_{\mathrm{GFRP}}} \tag{1}$$

where c is the speed of light in vacuum and n_{GFRP} is the assumed representative refractive index of the composite material. Small variations in n_{GFRP} due to fiber orientation or resin-rich regions are neglected here.

A-scans are useful for estimating ply count and thickness. Four plies are resolved in Fig. 8(a), and the overall calculated thickness (1.84 mm) agrees with the nominal value (1.95 mm), validating the use of a representative refractive index. Nevertheless, A-scans do not provide information on the lateral extent or continuity of internal features. In contrast, the B-scan shown in Fig. 8(b) reveals both through-thickness layering and transverse structural patterns. While B-scans are already established for detecting delamination [21], this work explores their applicability for identifying more complex impact-induced damage modes.

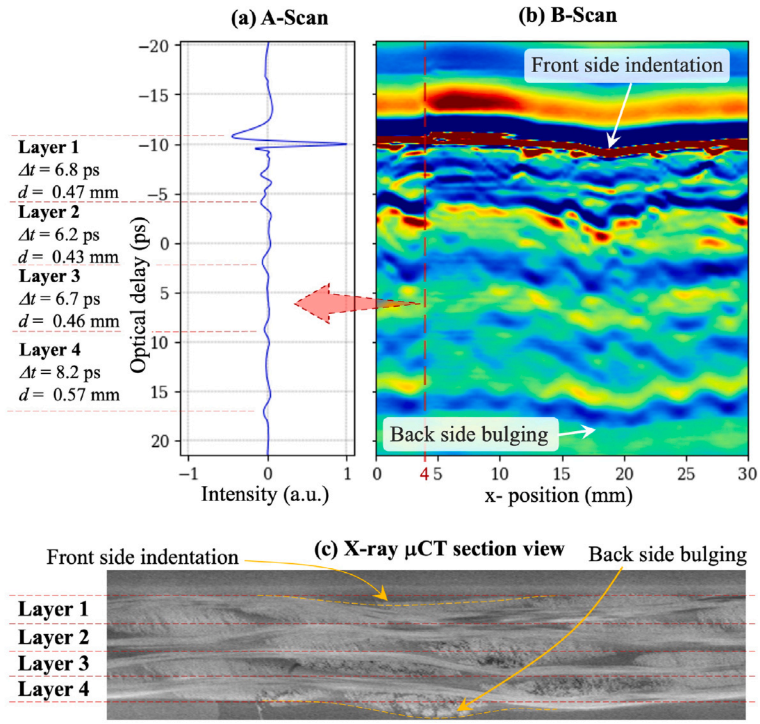

Fig. 8. Relationship between THz B-scan and X-ray μCT image of the same cross-section of an impacted coupon. (a) THz A-scan at x-position = 4 mm indicated on the adjacent plot of (b) THz B-scan of the cross section along y = 11 mm in Fig. 7. (c) X-ray μCT image. The x- and y-positions correspond to the scanning envelope shown in Fig. 7.

The ability to visualize anomalies in both axial and transverse directions makes B-scans well-suited for capturing the irregular, multi-depth nature of impact damage. As such, B-scan images form the basis of the training dataset used in this study. To ensure accurate labeling, each THz B-scan was compared against an X-ray μCT image of the same cross section [Fig. 8(c)], which served as ground truth for classification. This cross-validation step enabled supervised training of the neural network model on reliably annotated data. Fig. 8(c) also confirms that the drop-weight impact experiments carried out in this work indeed produces the expected damage typical of low-velocity impact, characterized by an indentation on the impact side and a localized crack on the opposite side [35]. This localized damage acts as an initiation site for damage propagation under further loading, such as fatigue caused by the pressurization cycles of pressure vessels.

3. Transfer learning with DenseNet for THz B-scan image classification

Traditional machine-learning approaches assume a consistent and structured feature distribution, thus their effectiveness is restricted for handling rather simpler tasks. For complex situations involving unstructured data and patterns, deep learning with CNN is necessary. Nevertheless, deep learning requires training on large dataset to achieve a robust model with a high level of generalization. When trained on a limited dataset, the model's ability to generalize to new data distributions is compromised, leading to performance degradation. This implies a typically time- and data-intensive process to build a robust deep learning model.

Transfer learning addresses a key challenge in DL: the need for large labeled datasets [36]. Originating in cognitive science, the idea is to apply knowledge gained in one context to another. Rather than training models from scratch for each new task, it allows the reuse of representations learned on related source domains to improve performance in data-scarce target domains. This is particularly valuable in cases where annotated data is limited. By leveraging the learned hierarchical features (e.g., edges, textures, local patterns), CNN models pre-trained on large datasets like ImageNet can be fine-tuned on the target-domain dataset, reducing both data and computational requirements.

In composite NDE and analysis, transfer learning has yielded notable classification accuracy scores for complex images [10] and under limited data conditions [37]. CNNs like VGGNet, ResNet, and DenseNet have been applied successfully to damage detection in fiber-reinforced composites [38], metals [39], and additively manufactured polymers [40]. Transfer learning has also supported hybrid approaches, including physics-informed neural networks (PINNs) for modeling interfacial damage [41].

Among CNN architectures, DenseNet stands out for its dense connectivity: each layer receives inputs from all preceding layers and contributes to all subsequent ones [42]. This architecture demonstrates an improved learning of subtle patterns that is critical for identifying complex or low-contrast defects in heterogeneous materials, such as BVID in multilayered composites. The THz B-scans that capture cross-sectional features of multilayer structures, including internal damage within the laminates, benefit from these DenseNets' characteristics.

In this study, a DenseNet model pretrained on ImageNet is fine-tuned to classify THz B-scan images of GFRP laminates as damaged or undamaged. The network learns from depth-resolved time-domain measurements, combining the structural sensitivity of THz imaging with the feature extraction capabilities of deep learning. The following sections detail the model architecture and training strategy.

3.1. DenseNet architecture

While CNNs have proven effective in computer vision due to their hierarchical feature learning, deeper networks can suffer from vanishing gradients, redundant features, and parameter inefficiency. DenseNet [42] addresses these issues by introducing dense connectivity: each layer receives the concatenated outputs of all preceding layers and passes its output to all subsequent ones within the block. This design improves feature reuse and reduces the number of trainable

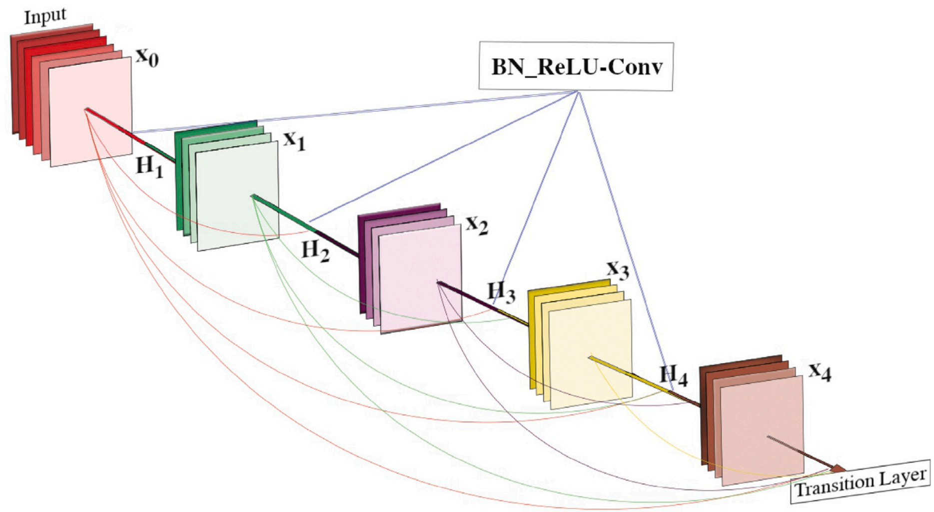

Fig. 9. Schematic diagram of a 5-layer dense block with a growth rate of k = 4. Each layer takes all preceding feature maps as input.

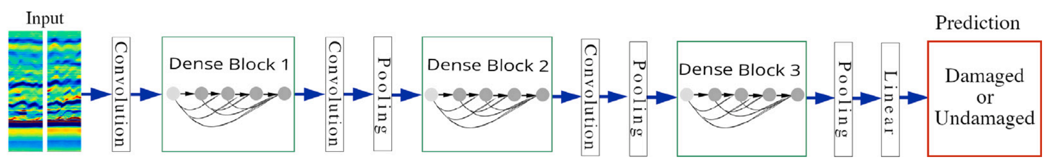

Fig. 10. Schematic diagram of the process of classifying B-scans. A deep DenseNet with three dense blocks.

parameters, thus helping the network learn compact representations from fewer examples.

DenseNets are parameterized by a growth rate k, which controls how many new feature maps each layer contributes. Thus, the input to layer ℓ consists of $k_0 + k(\ell - 1)$ channels, where k_0 is the number of input channels. The DenseNet formulation defines each layer ℓ as a composite transformation $H_\ell(\cdot)$ applied to the concatenation of all preceding feature maps:

$$x_\ell = H_\ell([x_0, x_1, \ldots, x_{\ell-1}]),$$

where $[x_0, x_1, \ldots, x_{\ell-1}]$ denotes the channel-wise concatenation. Each transformation $H_\ell(\cdot)$ typically consists of batch normalization, an activation function, and a 3×3 convolution. Fig. 9 shows a basic dense block with five layers and a growth rate $k = 4$.

To manage dimensionality and enable downsampling, transition layers comprising 1×1 convolutions and 2×2 average pooling are inserted between dense blocks, preserving compactness and feature reuse. This architecture, as employed in the current work, is illustrated in Fig. 10.

The activation function affects the model's ability to learn nonlinear features. ReLU remains the standard choice for its simplicity and efficiency, but alternatives such as Leaky ReLU, ELU, and GELU are also widely used depending on task requirements [43]. More recent variants, including *SoftPlus* (a smooth approximation of ReLU with a tunable β) and *SoftMax* (commonly used in classification tasks to produce probability distributions) [44], are increasingly integrated into modern architectures.

3.2. Modified DenseNet architecture

We employed DenseNet-121 as the backbone architecture owing to its efficient feature reuse and parameter economy, characteristics that are advantageous for THz B-scan classification where subtle structural variations may occur across multiple depth layers [42]. The dense connectivity pattern, in which each layer receives feature maps from all preceding layers, supports the preservation of fine-grained details that might otherwise be diminished in more conventional architectures. Such properties are beneficial for identifying barely visible impact damage in composite laminates [45].

DenseNet-121 comprises an initial convolutional layer, followed by four dense blocks with interleaved bottleneck and transition layers. The complete architecture includes:

- A 7×7 convolution (64 filters, stride 2),
- A 3×3 max pooling layer,
- Four dense blocks with [6, 12, 24, 16] layers,
- Three transition layers (1×1 conv + 2×2 avg pooling),
- A global average pooling layer and a modified classifier.

To meet the objectives of our classification task, the model architecture was adapted through specific modifications. The standard ImageNet classifier, designed for 1000-class recognition, was replaced with a task-specific head suitable for binary classification of THz B-scans. The modified classifier consists of a dropout layer followed by a linear transformation. This adjustment addressed several considerations: (i) the relatively small size of our domain-specific dataset, requiring regularization to mitigate overfitting; (ii) the need to adapt features learned from natural images to the characteristics of THz B-scans; and (iii) the binary nature of the classification task.

The model was implemented using PyTorch [46], retaining the pretrained ImageNet weights in the feature-extraction backbone to leverage learned low-level representations such as edges and textures that remain relevant to THz images. The classifier head is configured as:

```
nn.Sequential(
  nn.Dropout(0.5),
  nn.Linear(in_features=1024, out_features=2)
)
```

The dropout rate of 0.5 provided regularization appropriate for the dataset size (2708 images), while the linear layer mapped the 1024 feature dimensions from DenseNet-121 to the two output classes. This configuration preserved the feature extraction capacity of DenseNet while tailoring the decision-making component to the binary classification task.

The choice of DenseNet-121 over alternatives such as ResNet or VGG is further supported by previous studies in medical imaging, where datasets are often limited and fine-grained feature discrimination is required [47]. These conditions are comparable to those encountered in THz B-scan analysis.

3.3. Training strategies

The dataset of 2708 THz B-scans was split 70:30 into training and testing sets using a fixed random seed. Details on dataset generation and distribution are discussed in Section 4.2.

The transfer-learning approach is particularly useful when dataset size is limited, as in the present study. The models were pretrained on ImageNet, a publicly available dataset containing over 14 million labeled images [48], to learn general visual feature representations. The knowledge acquired during this pretraining was then transferred to the present classification task, with fine-tuning performed on the smaller, domain-specific dataset of THz B-scan images. Similar strategies have been applied in other DL-based NDT studies where annotated datasets are difficult to obtain. For example, investigations of railroad defects using ultrasonic B-scan images reported satisfactory performance with datasets of approximately 1000 images or fewer [49–51] While performance generally improves with larger datasets, the results presented in this work demonstrate the efficacy of transfer learning.

In total, 2708 B-scan images were used: 1492 from undamaged specimens and 1216 from damaged ones, corresponding to a 55:45 class distribution. Although the imbalance is relatively minor, a `WeightedRandomSampler` was applied during training to ensure equal representation of both classes. Sampling weights were assigned inversely to class frequencies, giving proportionally greater weight to the minority (damaged) class.

Given the dataset size, class balancing, and use of an ImageNet-pretrained model, a single train-test split, combined with data augmentation (random flips, rotations, and color jitter), was sufficient to ensure generalization while keeping training efficient.

The data augmentation procedure was applied on-the-fly during training, thereby increasing dataset variability without generating additional image files. The total number of images therefore remained 2708, while random variations were introduced at each epoch. Augmentation was implemented in PyTorch using the `transforms` module. Horizontal and vertical flipping were applied with `RandomHorizontalFlip` and `RandomVerticalFlip`. Images were randomly rotated within $\pm 30°$ range using `RandomRotation(30)`. Additional variability was introduced through color adjustments by applying `ColorJitter (brightness = 0.2, contrast = 0.2,saturation = 0.2, hue = 0.2)`.

Training was performed using the Adam optimizer [52] with an initial learning rate of 1×10^{-4} and weight decay of 1×10^{-4}. A scheduler reduced the learning rate by a factor of 0.1 every 10 epochs. The loss function was cross-entropy [53], appropriate for binary classification.

We applied early stopping with a patience of 10 epochs and capped training at 100 epochs. Training loss was monitored at each epoch to ensure convergence and avoid overfitting.

4. Dataset generation and labeling strategy

4.1. Overview of the implemented strategy

As a supervised learning task, the classification of THz B-scans required a labeled dataset in which each sample was annotated according to its actual structural condition. For this purpose, each B-scan was compared with an X-ray μCT image acquired at the same cross section to determine whether the region contains visible damage. The X-ray μCT images served as ground truth for supervised training and evaluation.

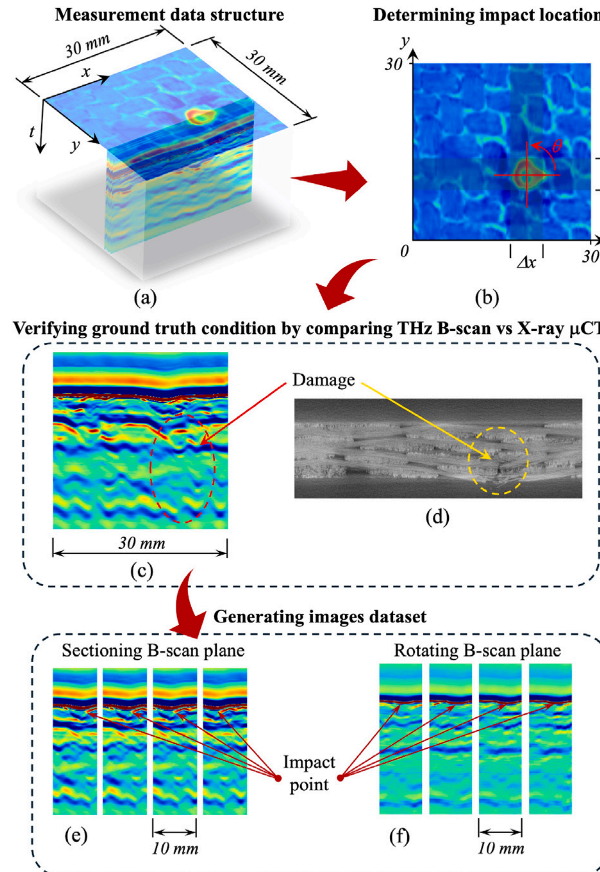

Fig. 11. Dataset generation process: (a) THz scanning data structure; (b) near-surface C-scan used to localize the impact site; (c) extracted B-scan image; (d) corresponding X-ray μCT image used for labeling; (e) additional B-scans generated by lateral sectioning; (f) B-scans generated by angular slicing around the impact location.

The labeling of B-scan images was conducted in accordance with established codes and guidelines for composite pressure vessel inspection. Images were annotated as damaged if the corresponding X-ray μCT cross-section revealed features such as cracks, delaminations, fiber ruptures, or surface protrusions (pimples), as these violate the acceptance criteria outlined in ASME BPVC Section X [22]. In addition, B-scans displaying dent-like indications on the side opposite the impact were also labeled as damaged, as such features warrant rejection under the CGA C-6.2 standard for high-pressure composite cylinders [23]. All other images without those features were labeled as undamaged.

To generate a large amount of data from a single measured specimen, B-scans were created by extracting cross sections within a narrow range around the impact center. Subsequently, the B-scans were sectioned at different intervals, provided that the damage was included within the B-scan section. This overall procedure is illustrated in Fig. 11.

The THz scanning data form a data cube consisting of two spatial dimensions (x, y) and time delay, as illustrated in Fig. 11(a). From this dataset, both C- and B-scans can be reconstructed. Given that low-velocity impact generates a damage zone of finite spatial extent, and that the lateral resolution is defined by the scan step size (here, 0.3 mm; transverse optical resolution is similar), multiple B-scans can be extracted that intersect the damaged volume.

Using the C-scan at the front surface Fig. 11(b), the impact location is estimated. To maximize the likelihood of including the damaged region, cross-sectional planes are selected within a narrow offset window around this center point, denoted Δx and Δy.

Each selected B-scan was compared by eye to the corresponding X-ray μCT cross section to ensure accurate labeling, as shown in Fig. 11(c) and (d). To further increase data diversity and representation, smaller sections of each B-scan–centered on the region of interest–were extracted [Fig. 11(e)]. This approach removed irrelevant areas far from the impact center and yielded additional training samples per specimen.

In addition to lateral slicing, B-scans were also generated by rotating the cross-sectional plane about a point near the impact location, as indicated by the angle θ in Fig. 11(b). This angular slicing, illustrated in Fig. 11(f), provided an alternative angular view of the internal structure and damage, potentially improving generalization in the classification model.

4.2. Quantitative description of B-scan slicing

As outlined in the previous section, B-scan images were generated using two approaches: lateral slicing and angular slicing.

In lateral slicing, B-scans were obtained on planes along the x- and y-axes, referring to the axis designation shown in Fig. 11(a). The cross-sectional planes that contained damage were those in the vicinity of the impact point, annotated as Δx and Δy in Fig. 11(b). These widths were determined by the size of the perceptible damage. Since the scanning resolution was uniform across all measurements (0.3 mm), the number of B-scans with visible damage indications depended on the extent of the damage.

The scanning area was 30 mm × 30 mm, as stated earlier in Section 2.2, thus each B-scan spanned 30 mm [Fig. 11(c)]. However, the impact only affected a limited region within that span. Therefore, it was sufficient to isolate a smaller section of the B-scan that included the damaged area. In this work, we used 10-mm sections of B-scans as training samples. Furthermore, the impact location did not necessarily need to be centered for an image to qualify for the damaged class. In other words, an image was assigned to the damaged class if the impact signature was present, regardless of its position. Consequently, the 10-mm sections could be shifted sideways, generating positional variations while still encompassing the impact point. An example of this sectioning technique is illustrated in Fig. 11(e), which shows several 10 mm sections with lateral offsets, all belonging to the damaged class since they contain impact indications.

Through this procedure, multiple images of the damaged class were extracted from a single sample. For instance, referring to Table 4, based on the measured size of the indentation crater in the sample impacted at 5.42 J, a width of 1.8 mm was taken for Δx and Δy. With a scanning resolution of 0.3 mm, that width covered 7 cross-sectional planes, hence 14 planes for both directions. From each full-span (30 mm) B-scan, 8 smaller 10 mm sections were extracted, as described earlier. Thus, 112 images (8 sections from each of 14 full-span B-scans) of the damaged class were obtained from this sample. When the induced damage was larger, as in the case of the 9.24 J sample, Δx and Δy were wider (3.3 mm), yielding 24 planes and thereby 192 (24 × 8) images. For the nonimpacted samples, 10-mm sections were extracted from random locations, but at least 20 mm away from the edges to avoid cutting artifacts. These were considered nondefective and labeled as the undamaged class.

Rotational slicing was performed by taking a 10-mm section plane with its centerline passing through the impact point. The plane was then rotated 180° about the centerline at intervals of 2°, generating 90 images of the damaged class for each impacted sample. This way, the resulting images contained the impact indication at or near the center. Images of the undamaged class were obtained in the same manner, except that the center points (rotation axis positions) were chosen at arbitrary locations on the non-impacted samples, provided they were at least 20 mm away from the edges.

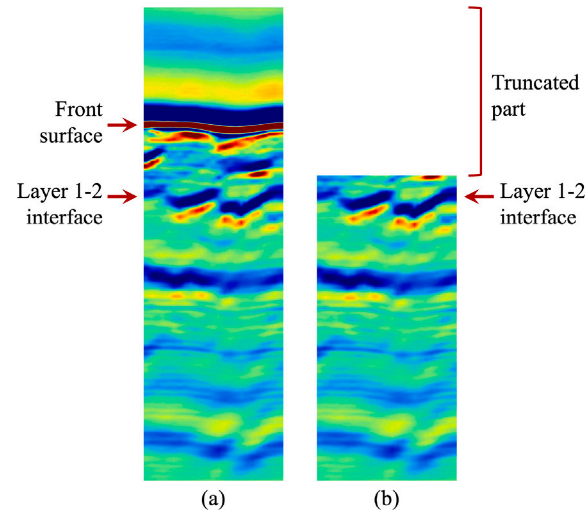

Fig. 12. Truncation of B-scans to exclude the front surface reflection: (a) original (full image) B-scan; (b) truncated image retaining only subsurface structure.

In this work, a total of 3108 images were generated using the above methods. Of these, 2708 images were used for training the model, and 400 were reserved as a held-out dataset for final independent evaluation. The held-out dataset was completely set aside during training and used only after the model was fully developed. Thus, the held-out images served as new data that had never been seen by the model, preventing biased predictions due to data leakage.

4.3. Exclusion of front surface reflections in B-scan images

A common feature in the B-scan obtained by THz TOFT is a high-amplitude reflection corresponding to the front surface. In the impacted region, this may be manifested as an indentation profile that could serve as a strong damage indicator. Such surface features, however, can also arise from unrelated phenomena, such as manufacturing artifacts or surface contamination. Since this study focuses on detecting structural damage that may not affect the front surface, such as subsurface or backside deformation, it is important to examine the classifier's performance without relying on the front surface profile.

To this end, an alternative dataset was created by spatially truncating each B-scan to exclude the front surface reflection. The resulting truncated images begin at the estimated interface of the second layer, emphasizing back-surface and internal features. This process is illustrated in Fig. 12 for several samples. This dual dataset, comprising full and truncated B-scans, enables evaluation of the classifier's ability to detect damage based on internal features alone. It also reflects a realistic use case for nondestructive inspection of regions that are not externally accessible, such as the inner wall of a composite pressure vessel.

5. Classification results and discussion

5.1. Model performance evaluation

We first evaluate the model's learning dynamics using training loss curves (Fig. 13) for both full and truncated B-scan datasets. For both datasets, the training loss decreases rapidly in the initial epochs. The truncated-image dataset, however, exhibits slower convergence, requiring more epochs to reach the early stopping criterion, defined as no improvement in loss for 10 consecutive epochs.

Table 4
Summary of dataset composition with lateral and rotational slicing.

Impact energy (J)	Width Δx, Δy (mm)	Lateral slicing images	Rotational slicing images	Training images	Held-out images	Class label
No impact	–	850	540	1340	50	Undamaged
4.39	1.8	112	90	152	50	Undamaged
5.42	1.8	112	90	152	50	Damaged
6.10	2.1	128	90	168	50	Damaged
7.16	2.7	160	90	200	50	Damaged
8.21	3.0	176	90	216	50	Damaged
9.24	3.3	192	90	232	50	Damaged
11.15	3.6	208	90	248	50	Damaged
Total images				**2708**	**400**	

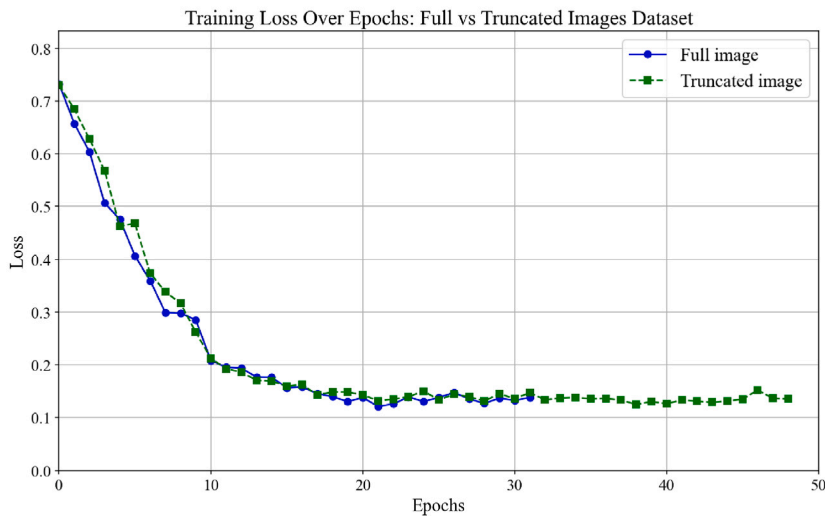

Fig. 13. Training loss evolution over epochs for full and truncated image datasets.

This behavior is attributed to the absence of the indentation profile in the truncated images, which likely reduces the salience of impact-related features during training. The full-image dataset contains the front surface reflection and indentation pattern, which may serve as a dominant visual cue for impact damage. In contrast, the truncated images dataset removes this feature, forcing the model to rely instead on structural anomalies in internal or back-surface regions. This increases the complexity of the learning task and affects the convergence rate, though not substantially degrading final performance.

The confusion matrices in Fig. 14 demonstrate effective classification for both datasets with DenseNet-121 model. Each matrix represents an example taken from a testing instance, and the results remain consistent over multiple trials as reflected in the statistics of model's performance shown in Fig. 15.

High true positive (TP) and true negative (TN) rates are observed, indicating that the model reliably distinguishes between damaged and undamaged regions. Although training on truncated images results in slower convergence, the model retains strong classification performance, suggesting robustness to the removal of superficial features.

To quantitatively assess performance, the model is evaluated using standard classification metrics: accuracy, precision, recall, and F1-score. These are computed based on the confusion matrix, where TP represents correctly classified damaged samples, TN denotes correctly classified undamaged samples, FP (false positive) indicates undamaged samples incorrectly classified as damaged, and FN (false negative) corresponds to damaged samples incorrectly classified as undamaged. The metrics for both datasets are shown in Fig. 15, showing the mean values of the performance parameters of 5 test runs with error bars that represent the corresponding standard deviation. The chart shows that there is only a small difference in all performance metrics between the full B-scan dataset and the truncated one.

Accuracy, defined as

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \quad (2)$$

reflects the model's overall ability to correctly classify both damaged and undamaged samples. High accuracy is observed for both datasets, with a slight decrease (0.3%) for the truncated images. This result confirms that the CNN can identify structural damage using internal and back-surface features alone, a desirable property for practical applications such as *in-situ* inspection of inaccessible regions.

Precision, given by

$$\text{Precision} = \frac{TP}{TP + FP}, \quad (3)$$

measures the proportion of correctly identified damaged samples among all samples predicted as damaged. High precision is important in applications where false positives may result in unnecessary repairs or part rejection. Precision values are 97.0% for full images and 95.5% for truncated images, showing minimal degradation when the front-surface features are excluded.

Recall, or sensitivity, is defined as

$$\text{Recall} = \frac{TP}{TP + FN}, \quad (4)$$

and quantifies the model's ability to correctly identify all damaged samples. This metric is particularly important in safety-critical applications, where missed detections (false negatives) may result in undetected structural degradation and subsequent failure in service. In NDE, especially for pressure vessels and safety-critical composite components, maintaining a high recall reduces the likelihood of overlooking subsurface or backside damage that may not present visible surface indications. In this study, the model achieved a recall of 97.2% on full images and 96.6% on truncated images, corresponding to a slight decrease of 0.6% when the front surface reflection is excluded.

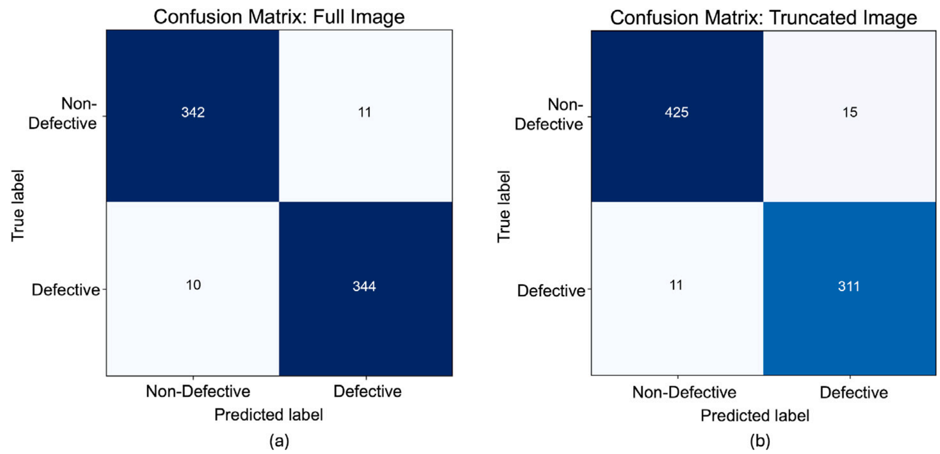

Fig. 14. Confusion matrix from a testing instance on (a) full B-scan images and (b) truncated B-scan images.

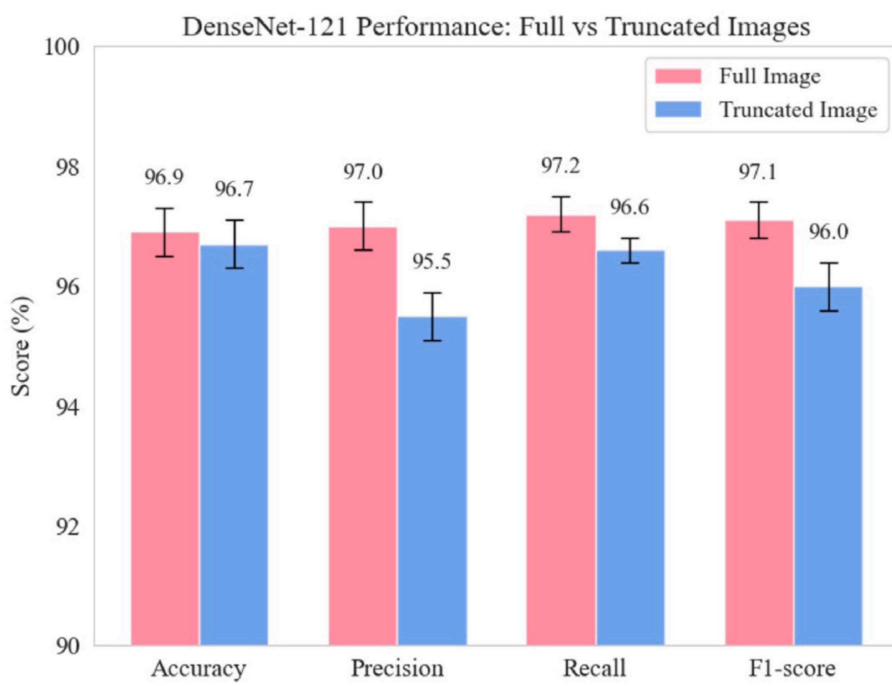

Fig. 15. Classification performance metrics for full and truncated B-scan datasets.

This result demonstrates that even when deprived of the most visually prominent damage cue (the indentation), the classifier remains highly sensitive to internal structural anomalies, confirming its relevance for inspecting inaccessible regions.

F1-score, defined as

$$\text{F1-score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}, \tag{5}$$

provides a single metric that balances precision and recall, offering insight into the model's ability to avoid both false positives and false negatives. This is especially useful when considering the dual cost of misclassification; false positives may result in unnecessary repair or part rejection, while false negatives could allow critical damage to go undetected. Even in a dataset with relatively balanced class distribution, the F1-score offers a robust measure of the classifier's reliability under operational constraints. In this study, the F1-scores for full and truncated datasets are both high and closely aligned, reflecting the model's ability to generalize across different types of image input and its robustness in identifying damage signatures that are not limited to superficial features.

Considering that the specimens were subjected to different impact energies, it is important to evaluate how damage severity influences classifier performance. To this end, inference tests were conducted using 100 B-scan images for each impact energy level. These images were entirely unseen during model development. The tests were performed on full B-scan images; however, as discussed earlier, comparable performance is expected for truncated images. The results, grouped by impact energy, are shown in Fig. 16. They indicate that model accuracy increases with higher impact severity, consistent with the presence of more pronounced and distinguishable indications. In addition, the 100% recall observed for the 0 J (non-impacted) group confirms the validity of the control condition, since no false negatives could occur in this baseline case, i.e. the dataset contained no damaged samples that could be misclassified as undamaged.

In summary, the classifier achieves robust performance for both full and truncated B-scans. Although convergence is slower without the front-surface reflection, the final performance is comparable. This supports the use of internal and back-surface features for reliable damage detection and highlights the potential of the proposed approach

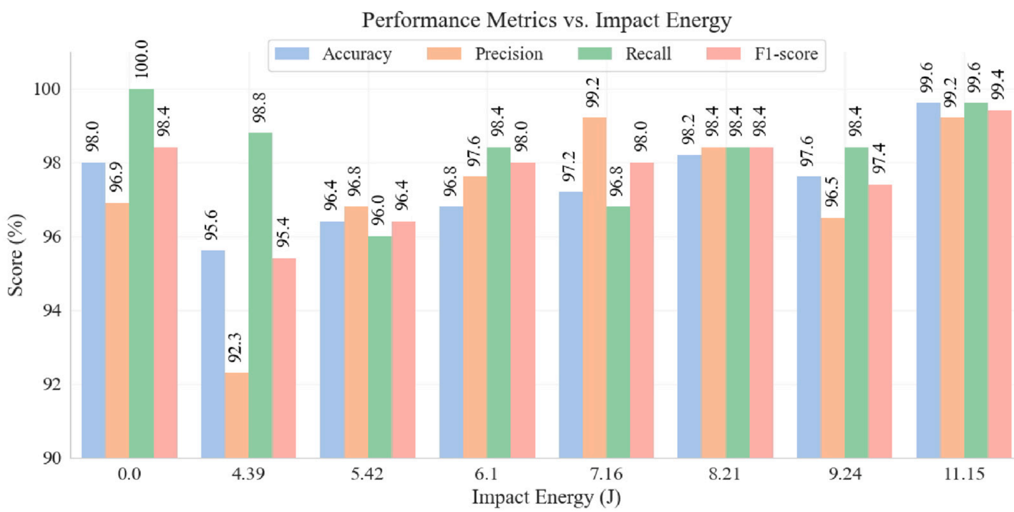

Fig. 16. Performance metrics for each impact energy level, tested on full B-scan images. Values are averaged over five runs, with standard deviations below 1% for all parameters. Accuracy increases with impact severity, and the 0 J group (non-impacted specimens) shows 100% recall, as expected for the control condition.

for inspecting regions that are not externally visible, such as the inner surfaces of pressure vessels.

5.2. Interpretation of classification results

As shown in the confusion matrices (Fig. 14), the DenseNet121-based classifier achieves consistent and accurate predictions across both the full- and truncated-image datasets. To further assess the model's decision-making process, a subset of B-scans from the test set is visualized alongside the corresponding X-ray μCT images in Figs. 17 and 18 for full images and truncated images, respectively. Each row displays the input B-scan, the model's binary prediction, and the X-ray μCT cross-section serving as ground truth. This comparison enables a direct qualitative evaluation of the model's predictions with respect to internal structural conditions.

Note that the comparison with X-ray μCT images here is only to validate the classifier's prediction results and confirm its accuracy. Once validated, the trained classifier enables reliable damage recognition solely based on THz B-scan images, eliminating the need for X-ray μCT imaging.

In all examples, the horizontal field of view spans 10 mm. The vertical axis of the X-ray μCT images is presented with the same aspect ratio as the horizontal axis, corresponding to the nominal composite thickness of 1.95 mm. In contrast, the vertical axis of the THz B-scans is stretched to improve visibility of time-domain reflections. Layer numbering throughout this analysis follows the convention that the impacted surface is designated as Layer 1 and the back surface as Layer 4.

Fig. 17(a) presents a B-scan obtained from a sample impacted at 11.15 J. The classifier correctly identifies the sample as damaged. The X-ray μCT image reveals severe damage including fiber breakage and crack propagation extending from the back surface towards the front. This failure mode, typical of low-velocity impact-induced bending, is well known to initiate from the tension side (back surface) and progress towards the point of load application [54]. The ability of the model to recognize this multilayered damage pattern confirms its effectiveness in capturing spatially distributed structural anomalies.

At an intermediate impact energy of 6.10 J (Fig. 17(b)), the classifier again successfully predicts the presence of damage. Here, the X-ray μCT image reveals localized deformation at the back surface and delamination below the fourth ply. Although the damage is more limited in extent, it still violates structural integrity criteria due to the presence of noticeable deformation on the back surface. This back surface, in the practical case of a pressure vessel subject to external impact, would be on the inner side that is not easily accessible for visual inspection. The classifier's success in detecting this type of damage demonstrates its sensitivity to single-layer damage.

Fig. 17(c) presents a case where the specimen, impacted at 4.39 J, shows no internal damage according to the X-ray μCT image. Although a surface indentation is present, no structural discontinuities are detected in the interior. The classifier correctly labels this sample as undamaged, consistent with the inspection guidelines, as there are no fiber ruptures and debond of the laminate [22], as well as internal (back-side) deformation [23]. It highlights the classifier's ability to distinguish between benign surface features and internal failures. In another example, Fig. 17(d), a B-scan extracted 22 mm away from the impact site is also correctly classified as undamaged, as confirmed by the X-ray μCT image. These cases demonstrate the model's low FP rate and its capacity to learn spatial localization of impact damage.

Similar observations are made for the classification results obtained using the truncated image dataset, as illustrated in Fig. 18. In these cases, the reflection peak from the front layer–presumably a prominent feature used by the classifier to detect anomalies–is absent. The exclusion of this feature alters the input representation, potentially affecting how the model identifies damage. It is important to note that surface indentation on the front layer does not always correlate with internal structural damage. This is exemplified in Fig. 17(c), corresponding to a low impact energy of 4.39 J, where a visible minor indentation on the front surface is present in the B-scan image, yet no subsurface damage is observed.

These results confirm that the classifier does not rely solely on surface reflections or the presence of indentation. While the front surface may provide useful cues, it can also be misleading in cases where surface features do not correlate with internal damage. The model instead learns to identify the characteristics of spatiotemporal patterns associated with impact-induced failure mechanisms, such as delamination, back surface deformation, and interlayer discontinuities, thus offering a more reliable indicator of structural integrity.

In summary, the successful classification of both full and truncated B-scan images highlights the robustness of the proposed method. The model demonstrates strong generalization to different damage depths, supporting its applicability for NDE of critical composite components, including regions not accessible to conventional visual methods. In addition, once the model is fully developed, running prediction on a single image takes less than one second on a standard PC (e.g., Intel® Pentium® Gold G5420 CPU @ 3.80 GHz with 16 GB RAM).

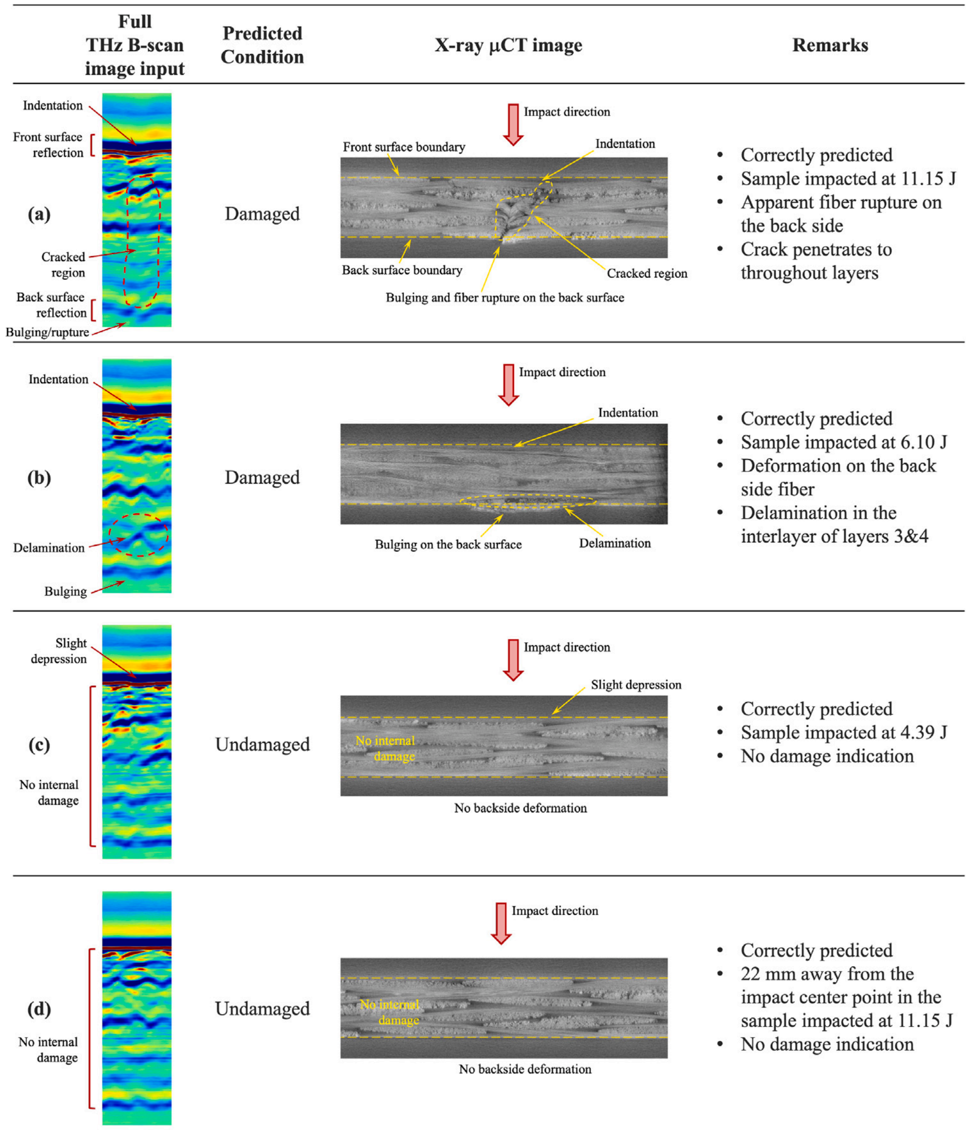

Fig. 17. Examples of correct predictions using full B-scan images. Each row shows the THz B-scan (first column), model prediction (second column), and ground truth from X-ray μCT (third column).

5.3. Comparison with other CNN architectures

It has been demonstrated in the preceding sections that DenseNet-121 achieves high classification accuracy. To establish the general performance figure of transfer learning for this work, other CNN architectures are evaluated, namely VGG-16, ResNet-50, GoogleNet, and Efficientnet. The handling, transformation, and augmentation of the datasets are maintained identical for all the training instances of those architectures.

While the performance metrics vary across architectures, all tested models yield convincing results, with accuracies around or in excess of 90%, as shown in Fig. 19(a) and (b) for training on full and truncated images, respectively. Each algorithm is tested for 5 runs and the scores shown above the bars indicate mean values of the obtained performance metrics. The standard deviations are between 0.3 and 0.5% in all cases, suggesting the repeatability and reliability of the models. A similar trend is observed in most architectures, where metric scores are slightly higher for full images compared to truncated ones. The exception is VGG-16, where the tendency is reversed. Nevertheless, the differences among the best-performing architectures are subtle, suggesting comparable effectiveness regardless of the presence of the first-reflection feature in the B-scan images.

Further optimization could certainly improve the performance of individual architectures. This comparison, however, is not intended to scrutinize each CNN algorithm, which is a task beyond the scope of this paper. Rather, it serves to validate that THz image classification via transfer learning is a robust approach, as evidenced by consistently high accuracy across diverse classifier architectures.

A directly comparable study on impact damage detection using THz B-scan images is not readily available in the literature. However, transfer learning has been applied successfully in other composite NDE tasks, yielding high predictive performance. For example, Gulsen et al. [55] investigated defect recognition in composites using ultrasonic C-scan surface images and reported high performance with VGG, DenseNet, and ResNet models, with DenseNet-121 achieving an accuracy of up to 98.8%. Such high accuracy is attainable partly because detecting surface defects is generally less complex than identifying subsurface defects, as attempted in the present work, where ultrasound has proven insufficient. The relatively stronger performance of DenseNet-121 compared to other models in that study is consistent with our findings on THz B-scan classification

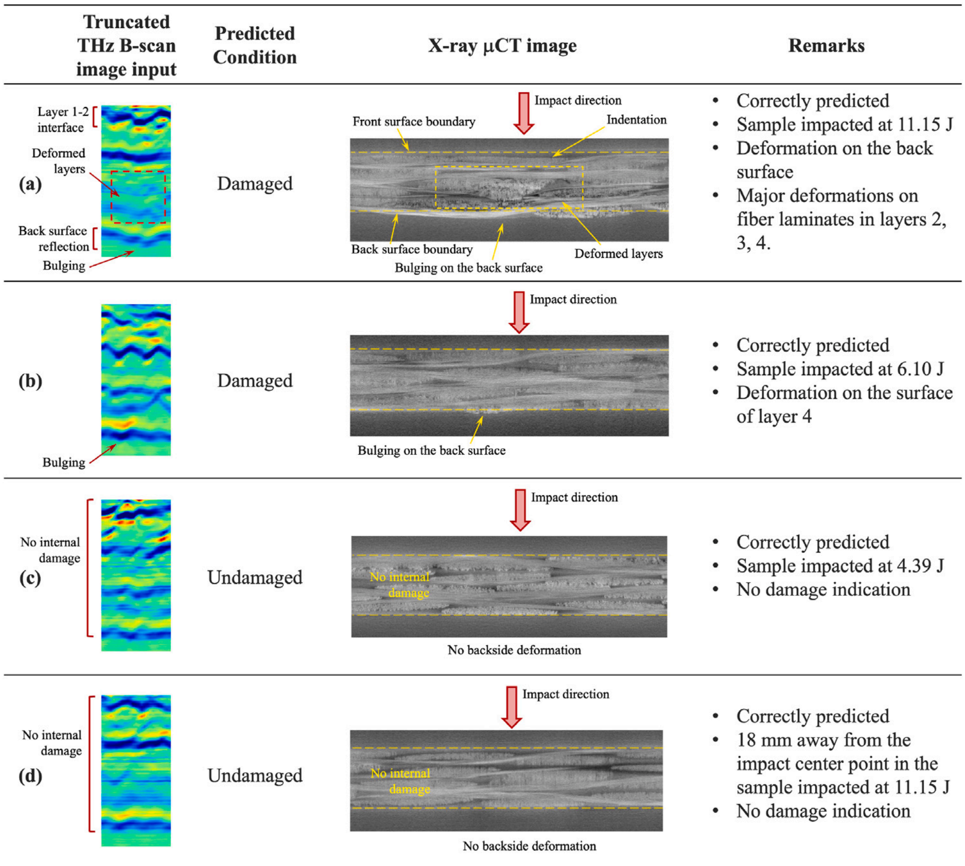

Fig. 18. Examples of correct predictions using truncated B-scan images.

A related study [56] examined delamination detection in composites using ultrasonic B-scan images, an imaging modality similar to our approach, where a ResNet-50–based model achieved an accuracy close to 99%. That investigation, however, addressed only a single defect mode, delamination, which is relatively well defined, whereas the present study focuses on impact damage that may involve multiple, morphologically complex defect modes. Furthermore, in [56] the investigated layup exhibited a minimum interlayer spacing of about 1 mm, making feature discrimination more straightforward than in the current material, where the nominal ply spacing is only 0.5 mm and requires correspondingly finer depth resolution.

Although no closely comparable case exists in the literature, the demonstrated effectiveness of transfer learning across other NDE modalities supports the validity of the present results and underlines the novelty of applying this approach to THz B-scan data for impact damage detection.

6. Conclusions

6.1. Concluding remarks

This study confirms that low-velocity impacts can produce barely visible damage in glass fiber-reinforced polyamide laminates, potentially compromising structural performance under subsequent loading. Reliable detection of such damage is essential in engineering applications where safety margins and operational reliability are critical. To address this challenge, we have presented an inspection method that combines terahertz time-of-flight tomography with deep learning-based classification of depth-resolved cross-sectional images.

The output of this approach is a binary assessment, either damaged or undamaged, directly applicable to pass–fail inspections for quality control and periodic maintenance. In such contexts, the presence of damage beyond a defined threshold may be sufficient to justify part rejection, regardless of the precise failure mode. This renders the method suitable for implementation in industrial workflows where speed, reliability, and automation are required.

The method was developed and validated on thermoplastic laminates composed of woven glass fibers embedded in a polyamide 66 matrix. These materials are of interest in structural applications due to their energy absorption capacity and processability. Although the experimental campaign focuses on flat specimens, the method is transferable to curved or tubular geometries typically encountered in components such as pressure vessels and pipelines. The approach is applicable to component- and structure-level inspection of cylindrical tanks, fuselage elements, or pressure housings, where internal access is restricted and depth-resolved information is required for reliable defect detection.

By leveraging the depth-resolved structural information provided by terahertz imaging, the proposed approach enables detection of internal damage patterns that span multiple plies and may be invisible to conventional surface-based inspection methods. Classification performance was validated using co-located X-ray μCT, which provided ground truth for training and evaluation. Once trained, the classifier reliably detected damage using THz imaging alone, eliminating the need for additional X-ray scans or expert interpretation during deployment. This facilitates integration into quality assurance pipelines without reliance on supplementary techniques and supports compliance with standards such as ASME BPVC Section X and CGA C-6.2 for pressure vessel requalification. The use of standardized criteria for damage annotation enhances confidence in the model's practical utility and promotes future adoption in regulated inspection workflows, while recognizing that penetration depth is material dependent and imposes application-specific limits. Notably, the classifier maintained strong performance

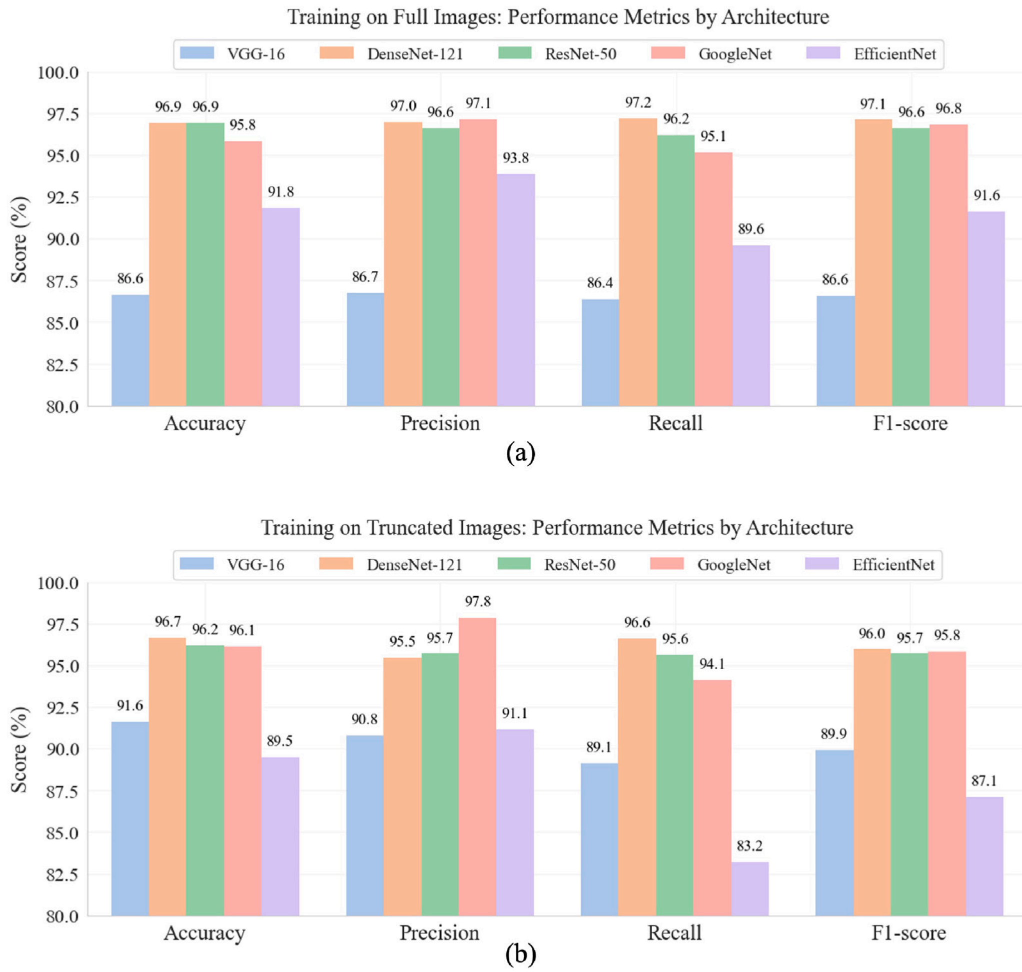

Fig. 19. Performance of various CNN architectures trained on (a) full images and (b) truncated images. The values shown are mean values from 5 testing runs, with standard deviation between 0.3 and 0.5%.

even when surface features were excluded from input, underscoring its effectiveness for detecting internal and back-side damage as key concerns in safety-critical applications.

These results demonstrate that combining terahertz imaging with data-driven classification enables robust and automatable inspection of composite structures. The method offers a practical alternative to traditional inspection techniques, particularly where access is limited or where non-contact evaluation is preferred.

6.2. *Future works*

The work presented in this paper focuses on flat-plate composite specimens, but the underlying concept is transferable to curved structures such as pipes and pressure vessels. This can be achieved by equipping the scanning system with a curve-following mechanism, for example by mounting the emitter and receiver on a rotation gantry moving coaxially with the object's curvature. Conventional rotational scanning systems can serve this purpose, maintaining the probe–object distance while acquiring measurements under quasi-normal incidence. In addition, current automation technologies enable similar measurements on objects with more complex geometries through the use of multi-degree-of-freedom robotic arms [57,58].

It should also be noted that the present approach is based on flat plates with uniform nominal thickness. In practice, many engineering components contain regions with intentional or fabrication-induced thickness variations that do not compromise structural integrity. Examples in pipelines and pressure vessels include grooves, seal seats, threads, fitting ends, and joints with added material for reinforcement. Special care is therefore required when interpreting results in such areas.

The findings and model performance reported in this study represent a proof of concept with potential for industrial implementation. The next step is to establish a standardized procedure and integrate the THz scanning system with a computer executing the classification task through a streamlined data feed. In this configuration, B-scans generated from each scan index can be transferred and analyzed as the measurement progresses, enabling real-time decision-making during inspections. The THz system used in this study is already of portable size, making it suitable for transport and on-site use. Portable THz imaging systems have previously been demonstrated in both industrial [59] and cultural heritage applications [60]. With further refinement of the workflow and integration of the trained classifier into the system, a practical route towards real-time, on-site inspection in industrial environments can be established.

While THz imaging offers distinct advantages, limitations remain. Penetration depth is strongly material dependent and must be considered to ensure suitability for the intended application. For GFRP, a practical upper limit of approximately 5 mm has been reported [61,62]. Although this thickness is sufficient for applications such as automotive panels, aircraft skins, and composite overwrapped pressure vessels, further developments in high-power sources and low-noise detection could extend the applicability of THz techniques to thicker materials.

Another challenge arises from the computational requirements of deep learning models. Training and inference demands increase with dataset size and image resolution, and may necessitate high-performance computing resources, adding complexity and cost. Continued optimization of model architectures will therefore be important to achieve efficient, low-resource classifiers and to support the practical deployment of THz-based inspection techniques.

Future work will focus on extending this framework to multiclass damage identification, adapting it to additional composite systems, integrating it with real-time scanning systems for inline inspection in manufacturing and maintenance environments, and optimizing model architectures to ensure efficient deployment.

CRediT authorship contribution statement

Dicky J. Silitonga: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Pascal Pomarède:** Writing – original draft, Resources, Project administration, Methodology, Formal analysis, Conceptualization. **Niyem M. Bawana:** Writing – original draft, Software. **Haolian Shi:** Writing – original draft, Methodology. **Nico F. Declercq:** Writing – review & editing, Funding acquisition. **D.S. Citrin:** Writing – review & editing, Funding acquisition, Conceptualization. **Fodil Meraghni:** Writing – review & editing, Resources, Funding acquisition, Conceptualization. **Alexandre Locquet:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization.

Declaration of competing interest

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

Acknowledgments

This work was supported partly by the French France 2030 program "Initiative d'Excellence Lorraine (LUE)", reference ANR-15-IDEX-04-LUE, and Institut Carnot ARTS. The High Performance Computing resources used in this work were provided by Georgia Tech Europe and Arts et Métiers Institute of Technology's Cassiopée HPC Center.

Data availability

Data will be made available on request.

References

[1] Guillet J-P, Recur B, Frederique L, et al. Review of terahertz tomography techniques. J Infrared Millim Terahertz Waves 2014;35:382–411. http://dx.doi.org/10.1007/s10762-014-0057-0.

[2] Hu J, Yang L, Qiao P, Zhou Y, Shi H, He Y. Research on internal defect detection of epoxy glass fiber composites based on terahertz imaging technology. Infrared Phys Technol 2023;133:104816.

[3] Wang X, Xu Y, Cui Y, Li W, Zhang L, Yan R, Chen X. DA-CNN-based similar terahertz signal identification for intelligent characterization of internal debonding defects of composites under high-resolution mode. Compos Struct 2023;322:117412.

[4] Kim D-H, Ryu C-H, Park S-H, Kim H-S. Nondestructive evaluation of hidden damages in glass fiber reinforced plastic by using the terahertz spectroscopy. Int J Precis Eng Manuf-Green Technol 2017;4(2):211–9. http://dx.doi.org/10.1007/s40684-017-0026-x.

[5] Latha AM, K K, Esampelly S, S B. Nondestructive testing of GFRP composites with compressive sensing-based terahertz single-pixel imaging. Nondestruct Test Eval 2025;1–18. http://dx.doi.org/10.1080/10589759.2025.2457575.

[6] Dong J, Locquet A, Declercq NF, Citrin DS. Impact damage characterization in hybrid fiber-reinforced composites using terahertz imaging in time and frequency domain. In: 2015 40th international conference on infrared, millimeter, and terahertz waves. IRMMW-THz, IEEE; 2015, p. 1–2. http://dx.doi.org/10.1109/irmmw-thz.2015.7327396.

[7] Tuo H, Lu Z, Ma X, Xing J, Zhang C. Damage and failure mechanism of thin composite laminates under low-velocity impact and compression-after-impact loading conditions. Compos Part B: Eng 2019;163:642–54. http://dx.doi.org/10.1016/j.compositesb.2019.01.006.

[8] Pomarède P, Miqoi N, Christophe D, Citrin D, Meraghni F, Locquet A. Terahertz pulsed imaging of low velocity impact damage in woven fiber composite laminates. J Compos Mater 2023;57:4719–30.

[9] Hasebe S, Higuchi R, Yokozeki T, Takeda S-i. Internal low-velocity impact damage prediction in CFRP laminates using surface profiles and machine learning. Compos Part B: Eng 2022;237:109844. http://dx.doi.org/10.1016/j.compositesb.2022.109844.

[10] Tabatabaeian A, Jerkovic B, Harrison P, Marchiori E, Fotouhi M. Barely visible impact damage detection in composite structures using deep learning networks with varying complexities. Compos Part B: Eng 2023;264:110907. http://dx.doi.org/10.1016/j.compositesb.2023.110907.

[11] Xu Y, Wang X, Zhou H, Hou Y, Wen B, Zhang L, Yan R, Chen X. Real-time terahertz characterization for composite delamination using a lightweight CPU adaptive network. Compos Part B: Eng 2022;247:110354. http://dx.doi.org/10.1016/j.compositesb.2022.110354.

[12] Zhang J-Y, Yang X, Ren J-J, Li L-J, Zhang D-D, Gu J, Xiong W. Terahertz recognition of composite material interfaces based on ResNet-BiLSTM. Meas 2024;233:114771.

[13] Xu Y, Hao H, Citrin D, Wang X, Zhang L, Chen X. Three-dimensional nondestructive characterization of delamination in GFRP by terahertz time-of-flight tomography with sparse bayesian learning-based spectrum-graph integration strategy. Compos Part B: Eng 2021;225:109225.

[14] Pałka N, Kamiński K, Maciejewski M, Dragan K, Synaszko P. Automatic histogram-based defect detection in glass fibre reinforced polymer composites using terahertz time-domain spectroscopy reflection imaging. Opt Lasers Eng 2024;174:107859.

[15] Medak D, Posilović L, Subašić M, Budimir M, Lončarić S. Deep learning-based defect detection from sequences of ultrasonic B-scans. IEEE Sensors J 2022;22:2456–63.

[16] Tunukovic V, McKnight S, Pyle R, Wang Z, Mohseni E, Pierce SG, Vithanage RKW, Dobie G, MacLeod C, Cochran S, O'Hare T. Unsupervised machine learning for flaw detection in automated ultrasonic testing of carbon fibre reinforced plastic composites. Ultrasonics 2022;140:107313.

[17] Aramburu AB, da Cruz JA, Xavier da Silva AA, Acosta AP, Minillo LQ, de Avila Delucis R. Non-destructive testing techniques for pressure vessels manufactured with polymer composite materials: A systematic review. Meas 2025;246:116729. http://dx.doi.org/10.1016/j.measurement.2025.116729.

[18] Dahmene F, Yaacoubi S, El Mountassir M, Langlois C, Bardoux O. Towards efficient acoustic emission testing of copv, without felicity ratio criterion, during hydrogen-filling. Int J Hydrog Energy 2016;41(2):1359–68. http://dx.doi.org/10.1016/j.ijhydene.2015.11.065.

[19] Haselmann M, Gruber DP, Tabatabai P. Anomaly detection using deep learning based image completion. In: 2018 17th IEEE international conference on machine learning and applications. ICMLA, 2018, p. 1237–42. http://dx.doi.org/10.1109/ICMLA.2018.00201.

[20] Iandola FN, Moskewicz MW, Karayev S, Girshick RB, Darrell T, Keutzer K. Densenet: Implementing efficient ConvNet descriptor pyramids. 2014, ArXiv abs/1404.1869.

[21] Deng J, Dong W, Socher R, Li L-J, Li K, Fei-Fei L. ImageNet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. 2009, p. 248–55. http://dx.doi.org/10.1109/CVPR.2009.5206848.

[22] ASME boiler and pressure vessel code, section X: Fiber-reinforced plastic pressure vessels. New York, NY: American Society of Mechanical Engineers; 2017.

[23] Standard for visual inspection and requalification of fiber reinforced high pressure cylinders. Chantilly, VA, USA: Compressed Gas Association; 2019.

[24] El Fallaki Idrissi M, Pasquale A, Meraghni F, Praud F, Chinesta F. Advanced metamodeling framework combining machine learning and model order reduction towards real-time virtual testing of woven composite laminates in nonlinear regime. Compos Sci Technol 2025;262:111055. http://dx.doi.org/10.1016/j.compscitech.2025.111055.

[25] El Fallaki Idrissi M, Praud F, Meraghni F, Chinesta F, Chatzigeorgiou G. Multiscale thermodynamics-informed neural networks (MuTINN) towards fast and frugal inelastic computation of woven composite structures. J Mech Phys Solids 2024;186:105604. http://dx.doi.org/10.1016/j.jmps.2024.105604.

[26] ASTM International. Standard test method for measuring the damage resistance of a fiber-reinforced polymer matrix composite to a drop-weight impact event. 2012th Ed.. ASTM International; 2012, current version approved April 1, 2012. Published June 2020.

[27] Chaichanawong J, Thongchuea C, Areerat S. Effect of moisture on the mechanical properties of glass fiber reinforced polyamide composites. Adv Powder Technol 2016;27(3):898–902. http://dx.doi.org/10.1016/j.apt.2016.02.006, special issue of the 5th Int'l Conf. on the Characterization and Control of Interfaces for High Quality Advanced Materials (ICCCI2015).

[28] Dong J, Kim B, Locquet A, McKeon P, Declercq N, Citrin D. Nondestructive evaluation of forced delamination in glass fiber-reinforced composites by terahertz and ultrasonic waves. Compos Part B: Eng 2015;79:667–75. http://dx.doi.org/10.1016/j.compositesb.2015.05.028.

[29] Wang J, Zhang J, Chang T, Cui H-L. A comparative study of non-destructive evaluation of glass fiber reinforced polymer composites using terahertz, x-ray, and ultrasound imaging. Int J Precis Eng Manuf 2019;20(6):963–72. http://dx.doi.org/10.1007/s12541-019-00114-z.

[30] Zhu T, Chen H, Pickwell-MacPherson E, Chen X, Fang G. Accurate reconstruction of terahertz spectral images with enhanced spatial resolution via complex mapping. Opt Express 2024;32(18):31657. http://dx.doi.org/10.1364/oe.529139.

[31] Kohlhaas RB, Breuer S, Mutschall S, Kehrt M, Nellen S, Liebermeister L, Schell M, Globisch B. Ultrabroadband terahertz time-domain spectroscopy using iii-v photoconductive membranes on silicon. Opt Express 2022;30(13):23896. http://dx.doi.org/10.1364/oe.454447.

[32] Tadokoro Y, Takida Y, Kumagai H, Nashima S, Kobayashi A. Comparative study on THz time-domain spectroscopy using 780-nm 1.3-ps laser pulses with different detections of lt-gaas photoconductive antenna and znte electro-optic sampling. In: Vodopyanov KL, editor. Nonlinear frequency generation and conversion: materials, devices, and applications XII. vol. 8604, SPIE; 2013, p. 86040D. http://dx.doi.org/10.1117/12.2001966.

[33] Calvo-de la Rosa J, Pomarède P, Antonik P, Meraghni F, Citrin D, Rontani D, Locquet A. Determination of the process-induced microstructure of woven glass fabric reinforced polyamide 6.6/6 composite using terahertz pulsed imaging. NDT E Int 2023;136:102799. http://dx.doi.org/10.1016/j.ndteint.2023.102799.

[34] Pomarède P, Calvo-de la Rosa J, Antonik P, Meraghni F, Citrin DS, Rontani D, Locquet A. THz-based monitoring of the deformation of the inner woven structure of glass-fiber reinforced polymers. IEEE Trans Terahertz Sci Technol 2023;13(4):396–404. http://dx.doi.org/10.1109/TTHZ.2023.3275278.

[35] Marguères P, Meraghni F. Damage induced anisotropy and stiffness reduction evaluation in composite materials using ultrasonic wave transmission. Compos Part A: Appl Sci Manuf 2013;45:134–44. http://dx.doi.org/10.1016/j.compositesa.2012.09.007.

[36] Kojima Y, Hirayama K, Endo K, Harada Y, Muramatsu M. Transfer-learning-aided defect prediction in simply shaped CFRP specimens based on stress distribution obtained from finite element analysis and infrared stress measurement. Compos Part B: Eng 2025;291:111958. http://dx.doi.org/10.1016/j.compositesb.2024.111958.

[37] Zhou Q, Chen Z, Xu J. Autonomous damage recognition of C/SiC composite based on transfer learning. Mater Express 2022;12(2):327–36.

[38] Khan A, Khalid S, Raouf I, Sohn J-W, Kim H-S. Autonomous assessment of delamination using scarce raw structural vibration and transfer learning. Sensors 2021;21(18):6239.

[39] Awan MR, Chan C-W, Murphy A, Kumar D, Goel S, McClory C. Deep learning and image data-based surface cracks recognition of laser nitrided titanium alloy. Results Eng 2024;22:102003. http://dx.doi.org/10.1016/j.rineng.2024.102003.

[40] Kuriachen B, Jeyaraj R, Raphael D, Ashok P, Sundari PS, Paul A. Defect detection in fused deposition modelling using lightweight convolutional neural networks. Eng Appl Artif Intell 2025;141:109802. http://dx.doi.org/10.1016/j.engappai.2024.109802.

[41] Chen Q, Tu W, Wu J, He Z, Chatzigeorgiou G, Meraghni F, Yang Z, Chen X. Elasticity-inspired data-driven micromechanics theory for unidirectional composites with interfacial damage. Eur J Mech A Solids 2025;111:105506. http://dx.doi.org/10.1016/j.euromechsol.2024.105506.

[42] Huang G, Liu Z, Van Der Maaten L, Weinberger KQ. Densely connected convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2017, p. 4700–8.

[43] Nguyen A, Pham K, Ngo D, Ngo T, Pham L. An analysis of state-of-the-art activation functions for supervised deep neural network. In: 2021 international conference on system science and engineering. ICSSE, IEEE; 2021, p. 215–20. http://dx.doi.org/10.1109/icsse52999.2021.9538437.

[44] Dubey SR, Singh SK, Chaudhuri BB. Activation functions in deep learning: A comprehensive survey and benchmark. Neurocomputing 2022;503:92–108. http://dx.doi.org/10.1016/j.neucom.2022.06.111.

[45] Pleiss G, Chen D, Huang G, Li T, van der Maaten L, Weinberger KQ. Memory-efficient implementation of densenets. 2017, arXiv preprint arXiv:1707.06990. URL https://arxiv.org/abs/1707.06990.

[46] Imambi S, Prakash KB, Kanagachidambaresan G. Pytorch, programming with TensorFlow: solution for edge computing applications. 2021, p. 87–104.

[47] Rajpurkar P, Irvin J, Zhu K, Yang B, Mehta H, Duan T, Ding D, Bagul A, Langlotz C, Shpanskaya K, et al. Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. 2017, arXiv preprint arXiv:1711.05225.

[48] Zhu S, Ye Z, Ai Q, Liu Y. EEG-ImageNet: An electroencephalogram dataset and benchmarks with image visual stimuli of multi-granularity labels. 2024, http://dx.doi.org/10.48550/ARXIV.2406.07151, URL https://arxiv.org/abs/2406.07151.

[49] Hu G, Li J, Jing G, Aela P. Rail flaw B-scan image analysis using a hierarchical classification model. Int J Steel Struct 2024;25(2):389–401. http://dx.doi.org/10.1007/s13296-024-00927-3.

[50] Wang S, Yan B, Xu X, Wang W, Peng J, Zhang Y, Wei X, Hu W. Automated identification and localization of rail internal defects based on object detection networks. Appl Sci 2024;14(2):805. http://dx.doi.org/10.3390/app14020805.

[51] Li W, Wang J, Qin X, Jing G, Liu X. Artificial intelligence-aided detection of rail defects based on ultrasonic imaging data. Proc Inst Mech Eng Part F: J Rail Rapid Transit 2023;238(1):118–27. http://dx.doi.org/10.1177/09544097231214578.

[52] Kingma DP, Ba J. Adam: A method for stochastic optimization. In: 3rd international conference on learning representations, ICLR 2015, San Diego, CA, USA, May (2015) 7-9, conference track proceedings. 2015, p. 1–15, published as a conference paper at ICLR 2015. URL https://arxiv.org/abs/1412.6980.

[53] Mao A, Mohri M, Zhong Y. Cross-entropy loss functions: Theoretical analysis and applications. In: International conference on machine learning. PMLR; 2023, p. 23803–28.

[54] Yang F, Cantwell W. Impact damage initiation in composite materials. Compos Sci Technol 2010;70:336–42.

[55] Gulsen A, Kolukisa B, Ozdemir AT, Bakir-Gungor B, Gungor VC. Defect classification of composite materials using transfer learning methods. Nondestruct Test Eval 2024;40(9):4338–54. http://dx.doi.org/10.1080/10589759.2024.2422527.

[56] Ding F, Cai Y, Tu W, Chen Q, Peng M, Ouyang F, Cai H, Fan K, Zhou W. Research on ultrasonic non-destructive detection method for defects in gfrp laminates based on machine learning. Ocean Eng 2025;327:120972. http://dx.doi.org/10.1016/j.oceaneng.2025.120972.

[57] Kang L-H, Han D-H. Robotic-based terahertz imaging for nondestructive testing of a pvc pipe cap. NDT E Int 2021;123:102500. http://dx.doi.org/10.1016/j.ndteint.2021.102500.

[58] Nie H, Hao F, Wang L, Guo Q, Chen H, Ren J, Wang K, Dang W, Liang X, Ma W. Application of terahertz nondestructive testing technology in the detection of polyethylene pipe defects. ACS Omega 2023;8(30):27323–32. http://dx.doi.org/10.1021/acsomega.3c02701.

[59] Latha AM, Ranjan R. Development of a compact and portable terahertz imaging system for industrial applications. In: 2024 IEEE microwaves, antennas, and propagation conference. MAPCON, 2024, p. 1–4. http://dx.doi.org/10.1109/MAPCON61407.2024.10922892.

[60] Doria A, Gallerano G, Giovenale E, Greco M, Picollo M. A portable THz imaging system for art conservation. In: 2018 first international workshop on mobile terahertz systems. IWMTS, 2018, p. 1–5. http://dx.doi.org/10.1109/IWMTS.2018.8454687.

[61] Wang S, Zhang X-C. Pulsed terahertz tomography. J Phys D: Appl Phys 2004;37(4):R1–36. http://dx.doi.org/10.1088/0022-3727/37/4/r01.

[62] Zhong S. Progress in terahertz nondestructive testing: A review. Front Mech Eng 2018;14(3):273–81. http://dx.doi.org/10.1007/s11465-018-0495-9.