



OPEN ACCESS

EDITED BY

Erfu Yang,
University of Strathclyde, United Kingdom

REVIEWED BY

Xin Lin,
Wuhan University of Science and Technology,
China
Ebrahim Oromiehie,
University of New South Wales, Australia

*CORRESPONDENCE

Homayoun Najjaran,
✉ najjaran@uvic.ca

RECEIVED 14 August 2023

ACCEPTED 08 February 2024

PUBLISHED 21 February 2024

CITATION

Ghamisi A, Charter T, Ji L, Rivard M, Lund G and Najjaran H (2024), Anomaly detection in automated fibre placement: learning with data limitations.
Front. Manuf. Technol. 4:1277152.
doi: 10.3389/fmtec.2024.1277152

COPYRIGHT

© 2024 Ghamisi, Charter, Ji, Rivard, Lund and Najjaran. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Anomaly detection in automated fibre placement: learning with data limitations

Assef Ghamisi ¹, Todd Charter ¹, Li Ji², Maxime Rivard³, Gil Lund⁴ and Homayoun Najjaran ^{1,5*}

¹Department of Electrical and Computer Engineering, University of Victoria, Victoria, BC, Canada, ²LlamaZOO Interactive Inc., Victoria, BC, Canada, ³Digital Technologies, National Research Council Canada, Ottawa, ON, Canada, ⁴Fives Lund LCC, Seattle, WA, United States, ⁵Department of Mechanical Engineering, University of Victoria, Victoria, BC, Canada

Introduction: Conventional defect detection systems in Automated Fibre Placement (AFP) typically rely on end-to-end supervised learning, necessitating a substantial number of labelled defective samples for effective training. However, the scarcity of such labelled data poses a challenge.

Methods: To overcome this limitation, we present a comprehensive framework for defect detection and localization in Automated Fibre Placement. Our approach combines unsupervised deep learning and classical computer vision algorithms, eliminating the need for labelled data or manufacturing defect samples. It efficiently detects various surface issues while requiring fewer images of composite parts for training. Our framework employs an innovative sample extraction method leveraging AFP's inherent symmetry to expand the dataset. By inputting a depth map of the fibre layup surface, we extract local samples aligned with each composite strip (tow).

Results: These samples are processed through an autoencoder, trained on normal samples for precise reconstructions, highlighting anomalies through reconstruction errors. Aggregated values form an anomaly map for insightful visualization. The framework employs blob detection on this map to locate manufacturing defects.

Discussion: The experimental findings reveal that despite training the autoencoder with a limited number of images, our proposed method exhibits satisfactory detection accuracy and accurately identifies defect locations. Our framework demonstrates comparable performance to existing methods, while also offering the advantage of detecting all types of anomalies without relying on an extensive labelled dataset of defects.

KEYWORDS

automated fibre placement, anomaly detection, computer vision, unsupervised learning, convolutional autoencoder

1 Introduction

Automated Fibre Placement (AFP) is an advanced composite manufacturing method for forming strong and lightweight components from strips of reinforced fibres known as tows. It is commonly used in quality-critical industries such as aerospace, where quality inspection and assurance are paramount (Palardy-Sim et al., 2019b; Böckl et al., 2023). Most existing inspection techniques are implemented via manual human examination strip by strip, which is time-consuming, and thus a major production bottleneck (Meister et al.,

2021b). To address this problem, recent research seeks to automate defect detection by using artificial intelligence (AI), computer vision (CV) and deep learning (DL) methodologies, reducing manual effort in AFP inspection, and expediting the production (Brasington et al., 2021).

Supervised learning with explicit labeling, a widely employed methodology in AFP inspection, relies on extensive labeled datasets of surface imaging. The surface images, crucial for training convolutional neural networks (CNNs), encompass various forms, including photographs, thermal images, and depth maps from profilometry sensors (Schmidt et al., 2019; Sacco et al., 2020; Zemzemoglu and Unel, 2022; Zhang et al., 2022). Photographs, captured using high-resolution cameras, provide detailed visual information of the composite part. Thermal imaging, on the other hand, exploits temperature variations on the surface, enabling the detection of defects based on their thermal signatures. This technique utilizes infrared sensors to capture and visualize the heat distribution across the composite material. Profilometry sensors, such as those described by Sacco et al. (2020), employ various technologies, including laser-based or structured light-based profilometry, to capture detailed depth maps of the composite part's surface. These depth maps represent the elevation of the surface, providing valuable insights into the structural integrity of the manufactured component. The synergy of these imaging techniques enriches the dataset for training AI models, enhancing their accuracy and effectiveness in AFP quality inspection.

Sacco et al. (2020) review the applications of machine learning in composite manufacturing processes and present a case study of state-of-the-art inspection software for AFP processes. The presented inspection method uses a deep convolutional neural network for semantic segmentation to classify defects on a per-pixel basis. They use about 800 scans which is relatively a large dataset in this domain, yet the results show their method often misses some defects. Object detection is a well-developed subfield of computer vision in which models learn to recognize specific objects from a large dataset of labelled bounding boxes. Zhang et al. (2022) offer an alternate approach using object detection. This work implements a modified YOLOv5 network, which is a popular and commercially available object detection model. With a large dataset of 3000 images containing five different defect types, their proposed model demonstrated effective performance in detecting those five defect types.

To achieve better real-time inspection, Meister and Wermes (2023) evaluate the use of convolutional and recurrent neural network architecture for analyzing laser-scanned surfaces line by line as 1D signals. The different network structures are assessed on both real and synthetic datasets, demonstrating sufficient performance. Through experimentation, the authors evaluate the effects of training and testing on differing data types (real or synthetic), realizing that deviations between the training and testing domain have a greater potential to impact the results of their proposed 1D analysis methodology.

These supervised learning methods, however, can be impractical in industrial projects because they require large, unambiguously labelled training datasets which are not typically available. There are three key reasons for the lack of labelled training data. First, collecting real-world data from production machines is expensive and disruptive to existing production schedules. Second, defects and

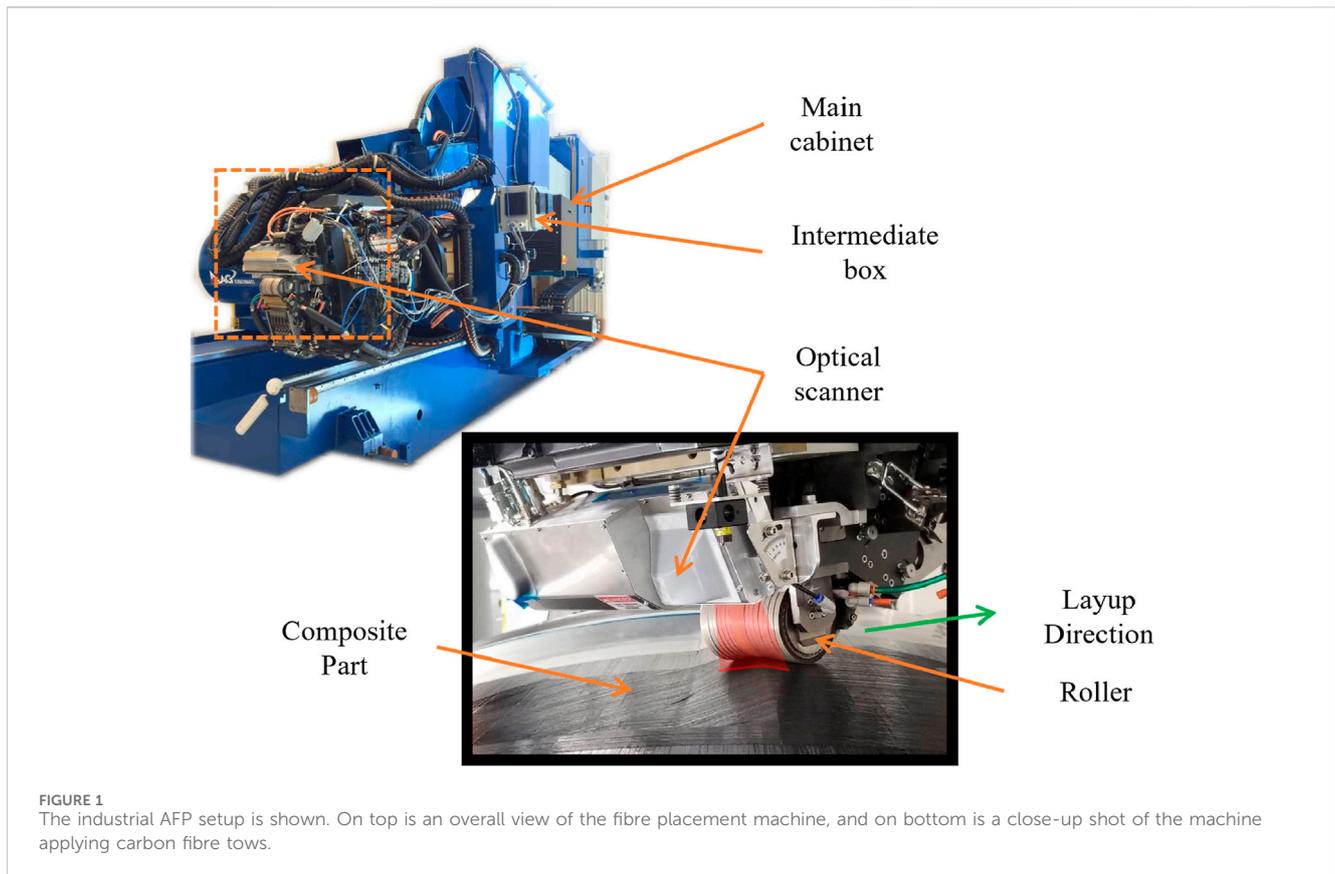
anomalies in real-world production are rare. To collect enough defect and anomaly samples for the models to learn from, one must collect a very large amount of data, adding to the training cost. Third, real-world defects can take many different forms, and there is no universally accepted standard of how a human inspector should delineate and record anomalies and defects, not to mention how to label them for machine learning (Heinecke and Willberg, 2019). Industrial practitioners of AFP manufacturing typically rely on organizational-specific standards and individual professional practices to identify and correct AFP anomalies. These separate standards and practices cannot be easily translated into a well-defined labelling strategy for labelling training sets.

Islam et al. present a novel approach to optimizing the AFP process through the integration of machine learning and a Virtual Sample Generation (VSG) method. The method generates virtual samples to address the challenge of incomplete information within the collected experimental data. The proposed approach incorporates physical experiments, finite element analysis, and VSG with the objective of refining the training process in scenarios with limited data, ultimately improving the accuracy of the machine learning model. However, the proposed approach lacks comprehensive validation, and although the integration of VSG holds promise for enhancing the training of machine learning models in AFP, the efficacy of this approach for defect detection tasks should be further explored.

One solution for addressing data limitations is to create synthetic datasets that can be utilized for training supervised models. Using AI models for synthetic dataset creation has been explored in many applications with varying success. To address the limited defect data in AFP, Meister et al. (2021a) compares different data synthesis techniques for generating defect data useful to AFP inspection applications. The paper compares synthetic image datasets generated by various GAN-based models, even implementing a CNN-based defect classifier for analysis. However, there is a lack of comparison of the generated datasets to real-world data regarding image diversity and realism (Meister et al., 2021a). In a similar problem, for the task of machine fault detection where faulty data is scarce and normal data is abundant, Ahang et al. (2022) implemented a Conditional-GAN for generating fault data in different conditions from normal data samples. This paper provides a more in-depth analysis comparing generated data to real ground truth fault data, showcasing that generated feature distributions are similar to those of real faults. Such data generation approaches have demonstrated effectiveness, though they still require sufficiently large and representative datasets, and cannot generalize to unseen defects or anomalies.

Circumventing the need for large datasets with labelled defects, unsupervised anomaly detection methods focus on learning the high-level representations of non-defective, normal data to identify outlying anomalies. In AFP manufacturing, normal data is typically well-defined thanks to the simple, invariant structure of the tows (narrow strips of composite material) and limited layup manners used. In this study, we utilize the non-defective samples, which constitute the majority of any real-world AFP datasets, to train a classifier capable of discerning normal and abnormal composite structures.

Autoencoders, known for their capability of reconstructing input data, have emerged as potent tools for detecting anomalies within images (Albuquerque Filho et al., 2022). An autoencoder



works by learning to encode normal input samples into a lower dimensional latent vector that can be decoded to reconstruct the original sample. The reconstruction is compared with the original, and a reconstruction error metric is calculated. When an abnormal sample is provided, the reconstruction errors will be high since the autoencoder was never trained with similar images. A threshold is then applied to the reconstruction errors to determine if the sample is normal, with a higher error indicating the sample is more likely to be abnormal or defective.

There is a lack of research applying these methods in the AFP industry, however, the approach has demonstrated success in other similar defect detection tasks. [Ulger et al. \(2021\)](#) employ convolutional and variational autoencoders (CAE and VAE) for solder joint defect detection. Reconstruction errors guide classification, applying a threshold to differentiate between normal and abnormal inputs. [Tsai and Jen \(2021\)](#) also employ CAE and VAE for textured surface defect detection, favoring CAE in a Receiver Operator Characteristic (ROC) analysis. The proposed anomaly detector is tested on various textured and patterned surface types, including wood, liquid crystal displays, and fibreglass. Additionally, [Chow et al. \(2020\)](#) implement a CAE for concrete defect detection, introducing a window-based approach for high-resolution images. Their window-based implementation enables pixel-wise anomaly maps to provide localization and contextual understanding of anomalies.

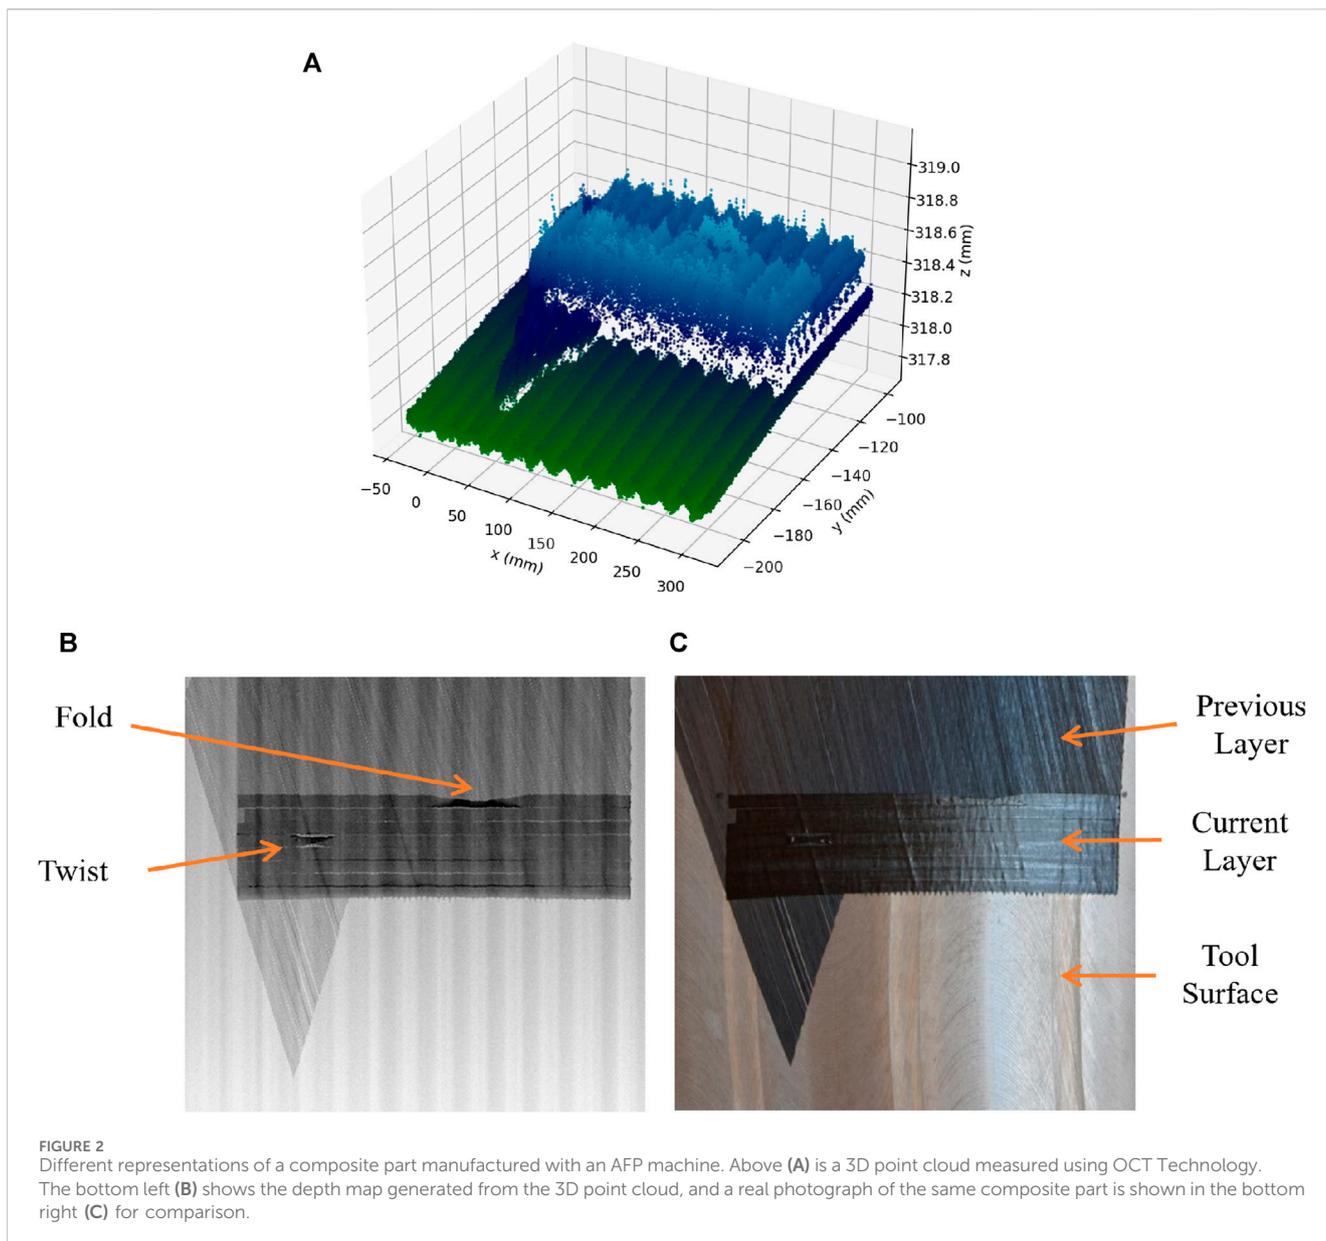
We propose a comprehensive framework for anomaly detection in AFP based on the autoencoder methodology. Compared to the existing methods that require a large labelled dataset including manufacturing defects, our approach is compatible with a small

training dataset of normal samples. Our autoencoder-based anomaly detector uses data collected from the AFP setup shown in [Figure 1](#). The autoencoder is trained on a collection of local samples taken from depth images of the composite carbon fibre surfaces. The depth images are obtained with an Optical Coherence Tomography (OCT) sensor installed on the AFP layup head, which captures high-resolution point clouds of the layup tow surfaces ([Palardy-Sim et al., 2019a](#); [Rivard et al., 2020](#)).

To simplify and enhance efficiency in processing information, the 3D point clouds are converted into 2D depth maps. Since the point cloud is a measurement of the surface elevation, it can be projected onto a 2D depth map without the loss of information. These depth maps are presented as grayscale images, where the brightness of each pixel corresponds to the surface elevation on the composite part. [Figure 2](#) provides various representations of a sample composite part, highlighting that defects are less discernible in the photograph due to reflections and low visual contrast. However, defects become more evident in the depth map, facilitating defect detection.

The work offers the following contributions.

1. We introduce a novel, end-to-end framework for anomaly detection and localization in Automated Fibre Placement, circumventing the data limitations in this industry. Our proposed framework has several advantages compared to the existing methods. It detects all types of anomalies, without the need for manual data labeling or defect samples. Also, it works with a small number of composite images.

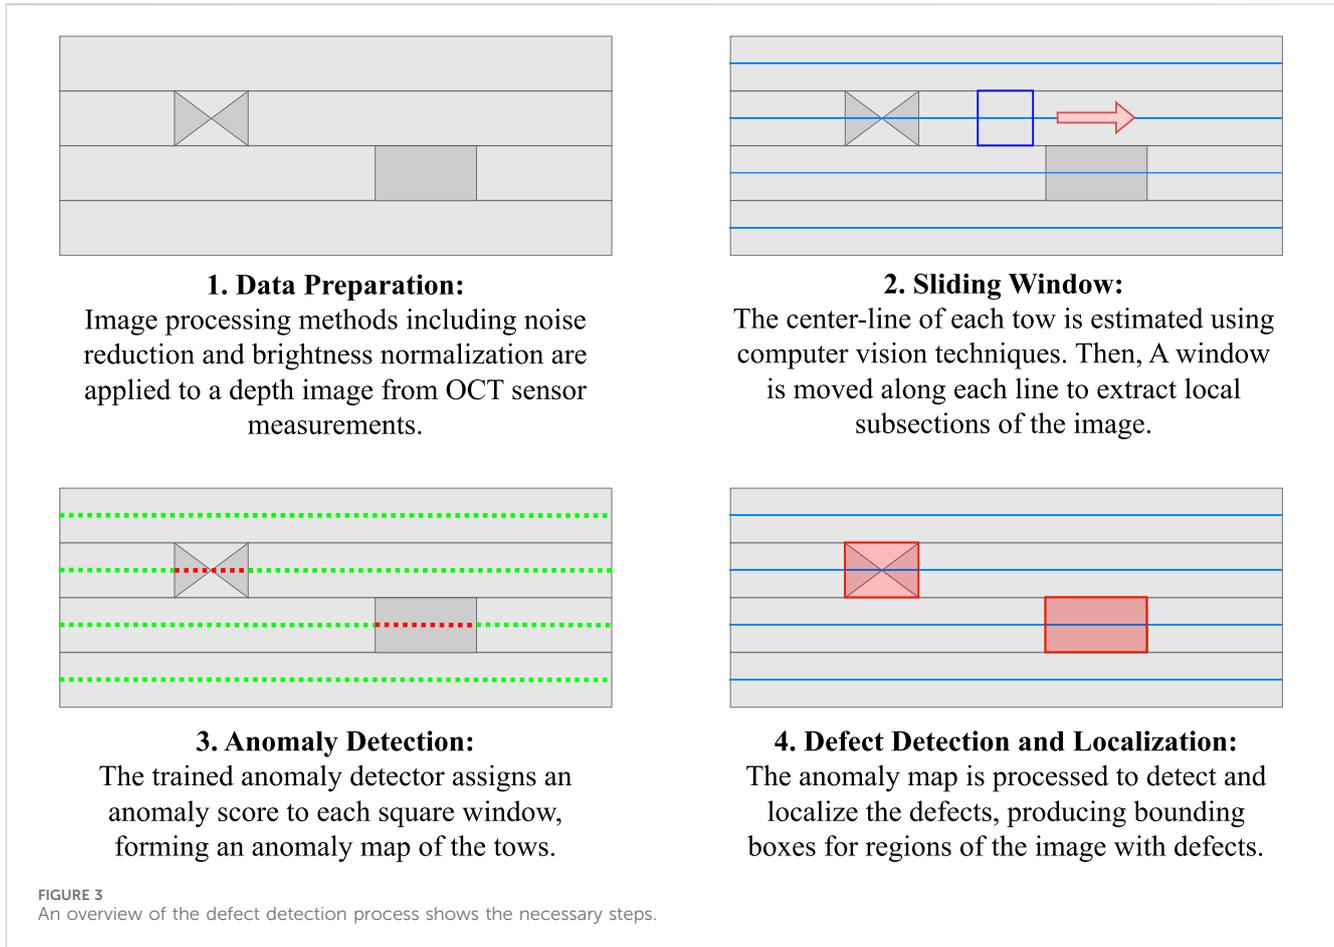


- Another major contribution of this work is an efficient data extraction methodology that can convert a limited number of composite images into a large dataset of local samples. Utilizing classical computer vision algorithms to detect the boundaries of composite tapes, this method generates a dataset that exploits the inherent symmetry of AFP composite materials.
- We design and validate an autoencoder with the optimal size of the latent domain that can identify the best distinctive features to differentiate between normal and defective samples.
- The proposed framework generates a map representing the local anomaly score of the AFP-manufacture parts and visualizes this map on the original composite scan. This visual representation serves as a valuable tool for AFP technicians, aiding them in the identification and resolution of anomalies within the composite structure.

The organization of this manuscript is as follows: **Section 2** describes, the whole procedure of anomaly detection, including data preprocessing, training of the AI model, and implementation details. In **Section 3**, the evaluation results of both the anomaly detector and the localization system are presented. Finally, concluding remarks of this work are provided in **Section 4**.

2 Methodology

Figure 3 summarizes our anomaly detection framework. First, a composite scan is processed and local samples are extracted from the images. Then, the trained autoencoder generates an anomaly map, used to detect and locate the defects in the image.



2.1 Data preparation

The raw depth maps contain impulse artifacts, also known as salt and pepper noise, which can be detrimental. To remove the noise, we used a median filter with the kernel size 3×3 applied to the whole depth image (Azzeh et al., 2018). Compared to a Gaussian filter with a small radius (a low-pass filter) (Deng et al., 2016), a median filter has less risk of losing high-frequency features.

Another data preparation step is needed because different raw depth-map images have inconsistent ranges of values depending on the distance of the laser origin to the composite surface. This effect is commonly caused when the OCT sensor is mounted to a fixed location behind the AFP head while scanning a contoured surface. These variations can cause undesired behaviour in the defect detection methodology. To address this, all the images undergo a min-max normalization so that the minimum depth value is mapped to zero and the maximum value is mapped to one. By applying this linear transformation, the visual contrast of the images is improved while keeping the original depth ratio. The normalization function is provided in Eq. 1 in which $z_{i,j}$ is the original depth value, $p_{i,j}$ is the normalized pixel value, and Z is the whole depth map matrix.

$$p_{i,j} = \frac{z_{i,j} - \min(Z)}{\max(Z) - \min(Z)} \quad (1)$$

2.2 Local sample extraction

The training dataset used in this work is composed of depth maps from 42 non-defective composite surfaces. Developing an effective end-to-end network for defect detection using such a limited number of scans presents a significant challenge. However, we address this issue by leveraging the consistent uniformity along the composite tows and extracting cropped windows of the scans to form a dataset with many more localized samples. This is possible under the assumption that each cropped section conforms to a similar distribution, given that defect-free tows should exhibit minimal to no disparity across their segments. Consequently, the analysis of smaller regions allows us to employ a more compact neural network to learn from a broader spectrum of local samples rather than relying on a larger network to process full scans. Moreover, by extracting localized samples along the tows, our network gains exposure to a greater variation of tow structure.

One of the most basic methods to detect local objects in an image is to move a window on the image and classify the smaller region inside the window (Dalal and Triggs, 2005). This method is known as “sliding window” in computer vision literature and has its own limitations. For example, the scale of the object may vary depending on how close the object is to the camera. Consequently, multiple sizes of sliding windows are required which can be computationally expensive. In our current dataset, on the other hand, most of the

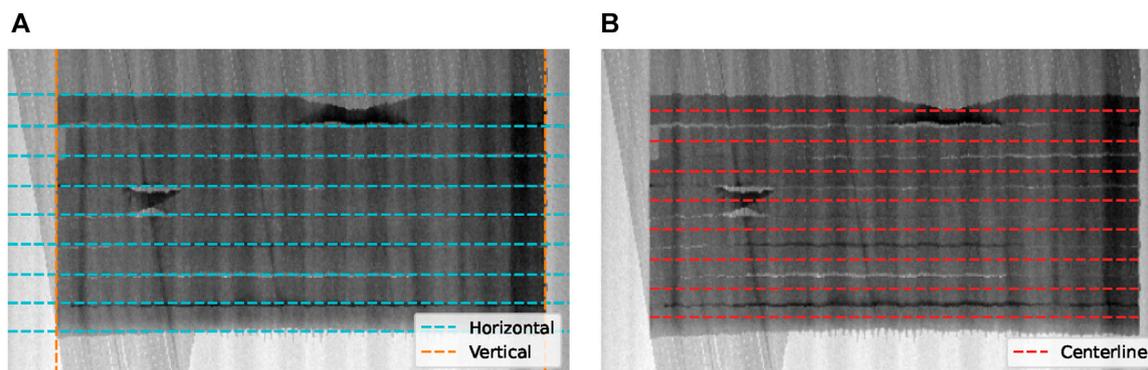


FIGURE 4
The centerline detection procedure contains two main steps: detecting horizontal and vertical lines (A) and estimating tow centerlines from the detected lines (B).

defects are localized to one tow, and therefore approximately the same relative size and there is no wide variation in perspective or orientation of the objects. Consequently, only one scale of sliding windows is sufficient for this use case. This also helps to keep computation complexity relatively low for this approach. Besides, there is preliminary knowledge of the composite part scans, like the number of tows and the general direction they follow. This enables a customized sliding window method that makes use of the known information. Moreover, the depth maps generated from our OCT scans have a specific structure. For example, all tows are placed straight and horizontal in the images, the number of tows is known, and their width is also known. To incorporate this predetermined knowledge, a line detector algorithm based on Hough Transform (Hough, 1959; Duda and Hart, 1972) detects the vertical and horizontal edges of the tows. After detecting the boundaries of the tows, the center of them (centerlines) are calculated by averaging each two consecutive horizontal lines, bounded within the detected vertical lines. This process of centerline detection is illustrated in Figure 4. Finding the center of the tows makes it possible to directly focus on the regions that are candidates for defect instead of scanning the whole image. In other words, it creates a skeleton that directs and constrains the region of interest. This can reduce additional effort on the classifier side.

Based on the detected centerlines, a square window slides across the tows to extract cropped regions. To select the size of the window we use the width of the tow plus a 50% of extra margin to consider the boundary of the tows. This 50% of extra margin is sufficiently large to cover the upper and lower boundary of the tows. A larger margin will cover a big portion of neighboring tows which is not desirable for the tow-by-tow approach proposed in this work. The window size calculated in this way will be 32×32 pixels. In this implementation, a step size of 8 pixels is used to move the window and sample the information cropped inside. This size of stride allows enough overlap between the nearby samples while keeping the samples sufficiently distinctive. Using a smaller stride increases the number of generated samples which leads to increased computational requirements for training. On the other hand, a larger stride allows less overlap between the windows, resulting in chance of overlooking or undervaluing key features of the AFP composite structure. Some of the extracted samples are presented in

Figure 5. At the time of inference, each window that is detected as an anomaly is considered to contain a defect, while the windows that are not anomalies are assumed to have normal tow structures. The next sections explain the approach to distinguishing between normal and abnormal samples.

2.3 Anomaly detection

As mentioned in Section 1, autoencoders have shown great success at identifying anomalies in images. An autoencoder is an unsupervised learning model that reconstructs the given input by learning to minimize the error between the input and reconstructed output. They do this by encoding the input to a vector of latent features, also known as the bottleneck, and then decoding those latent features to reconstruct the input. Convolutional Autoencoders (CAEs) are a group of autoencoders that use convolution layers in their network structure. Convolutional neural networks (CNN) are more popular for image-based autoencoders than basic fully-connected networks. This is because CNNs incorporate receptive fields using kernels that maintain the spatial relationships of the data. CNNs are also computationally efficient with sparse connectivity of neurons.

If only normal samples are used to train the autoencoder, it will be able to reconstruct similar normal samples accurately, and the reconstruction results for abnormal samples will be poor. Therefore, reconstruction error can be used as an indicator of how anomalous each input is. For inference, each cropped window of a composite material depth map is fed into the trained autoencoder. The reconstruction error of each window is then used as an anomaly score to create an anomaly map for the entire image. Reconstruction error of a window centred at (x,y) is calculated using Eq. 2 in which $p_{i,j}$ and $\hat{p}_{i,j}$ are the pixel value of the input and reconstructed output, respectively, and b is half of the size of each window.

$$m_{x,y} = \frac{1}{(2b)^2} \sum_{i=x-b}^{x+b-1} \sum_{j=y-b}^{y+b-1} [p_{i,j} - \hat{p}_{i,j}]^2 \quad (2)$$

In this work, a CAE is designed and used as the anomaly detector. The design incorporates symmetric encoder and

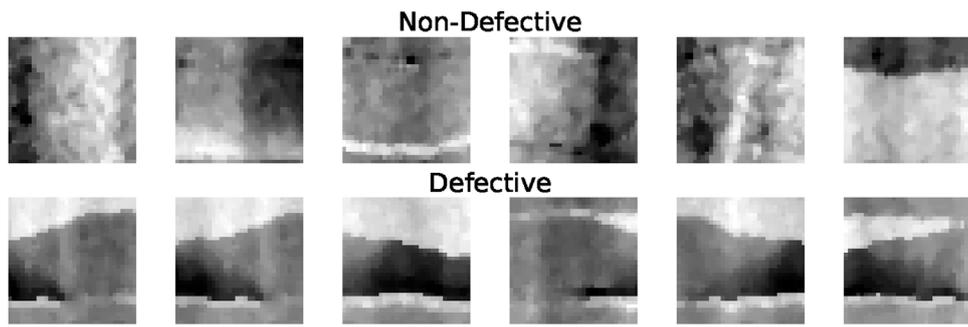


FIGURE 5
A dataset is created from cropped sections of the depth maps, using the sliding window method. Normal samples are shown on top and abnormal samples are shown below.

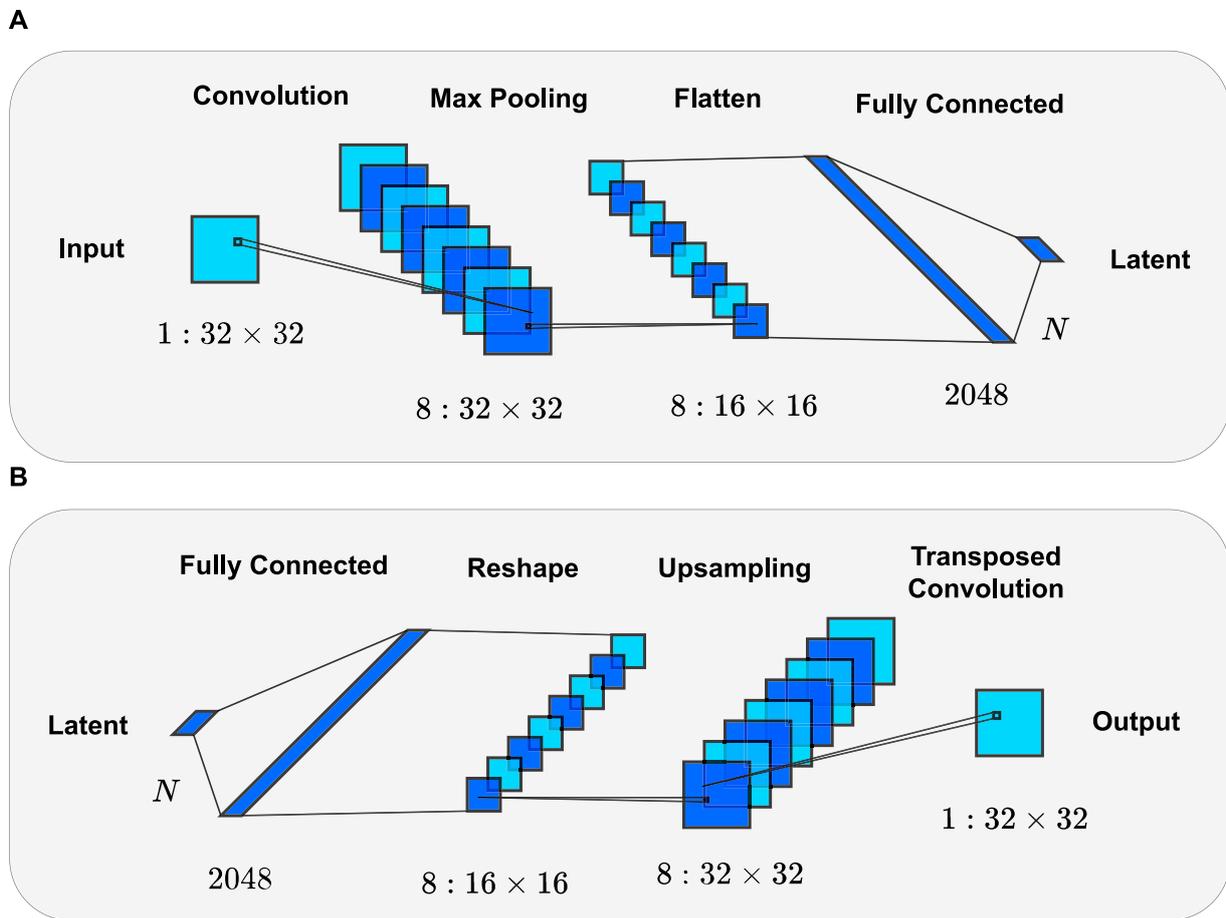


FIGURE 6
A graphic depicting the network structure of the proposed autoencoder. Above is the encoder structure (A), and below is the decoder structure (B).

decoder structures shown in Figures 6A, B respectively. For training the model, mean squared error is employed as the loss function. Although the full method uses continuous valued anomaly maps to identify the defects rather than a binary prediction, binary classification can still be useful in validating the model performance. For this, a threshold parameter is introduced to

classify samples based on their reconstruction error. To select the threshold value, a Receiver Operating Characteristic (ROC) curve is applied. The ROC curve plots the true positive rate against the false positive rate while varying the threshold value. In an ideal case, the selected threshold would give a true positive rate of 1 and a false positive rate of 0. In the ROC plot, this corresponds to the upper left

corner and hence the best threshold value is selected from the curve at the point closest to that corner.

2.4 Defect localization

The anomaly detection generates an array of anomaly scores for each tow, which can be considered as a 1D digital signal. Any area of this signal with a concentration of high values indicates the presence of a defect. In Computer Vision literature, these areas are called blobs (Danker and Rosenfeld, 1981; Kong et al., 2013). For detecting the blobs, we use the Difference of Gaussian (DoG) method (Lowe, 2004). In this approach, the signal ($f(x)$) is filtered using Gaussian kernels with increasing values for standard deviations (σ) as described in Eq. 3. Then, the subtractions of each two successively filtered signals are calculated. The local maxima of $g(x, \sigma)$ represents the blobs. In such maxima points, x and σ correspond to the location and characteristic scale (size) of the blob, respectively.

$$g(\sigma, x) = \sigma^2 \frac{\partial m_\sigma^2}{\partial x^2} * f(x) \quad (3)$$

For each defect, two parameters are detected, radius and center. With this information the detected blobs can be transferred from anomaly map to image space.

3 Results and discussion

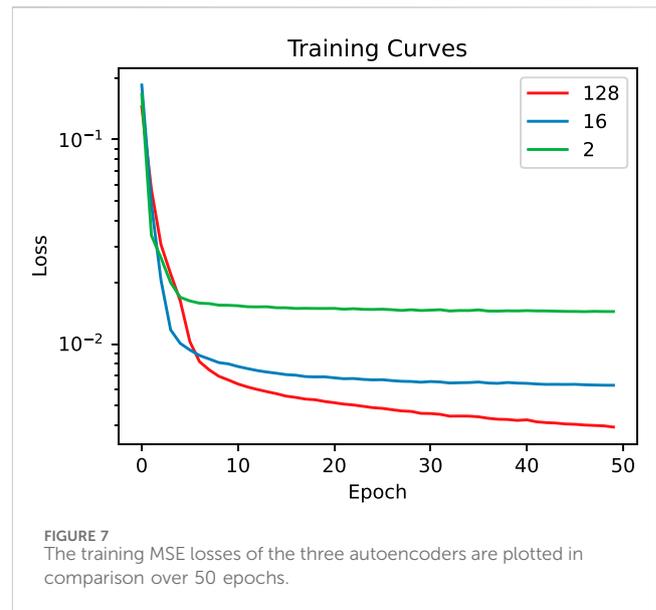
The performance of the anomaly detection and localization system depends on two factors. First is the number of samples the anomaly detector is correctly classifying as normal or abnormal. Second is the size and location accuracy of the predicted defects. This section evaluates these two aspects using a test dataset with an additional two composite surfaces containing defects.

3.1 Anomaly detection

The network structure proposed in Figure 6 is implemented using three different latent dimensions 2, 16, and 128 for comparison. Each network is trained with a dataset consisting of 27406 only normal samples. An Adam optimizer is employed with an MSE loss function to train the network. The batch size is set to 128. Each autoencoder undergoes training for 50 epochs, completing in under 5 min on a computer with the following specifications:

- Processor (CPU): Intel(R) Xeon(R) E5-1607 v4 @ 3.10 GHz
- Graphics (GPU): NVIDIA GeForce GTX 1080
- Memory (RAM): 32.0 GB

The curves in Figure 7 demonstrate the training losses of each autoencoder. Comparing the training loss curves shows that the models' reconstruction ability improves with higher dimensional latent space. The curves also show that the models are learning relatively quickly with tiny improvement in the later epochs.



Reconstruction results for the autoencoders are demonstrated in Figure 8. The original samples are randomly selected from normal and abnormal classes in test set. These results clearly show the improved reconstruction performance with a higher dimensional latent space. It shows that the autoencoder with a 128-dimensional latent vector is able to produce good reconstructions for both normal and abnormal samples. The 16-dimensional autoencoder, on the other hand, produces relatively good reconstructions for normal samples and poorer reconstructions of defect samples. This is ideal for the classification method to distinguish anomalies. Finally, the autoencoder with only a 2-dimensional latent space is unable to make good reconstructions for any of the input samples.

Figure 9 shows comparisons of the reconstruction error. In Figure 9A the distributions of mean square error are shown for the 16-dimensional autoencoder on the training set and test set. Note that the training set only includes normal samples whereas the test set contains both normal and abnormal samples, separated accordingly. As the figure suggests, the normal samples have a similar distribution in both the training and test sets. On the other hand, the abnormal samples have a generally higher MSE with a slight overlap on normal sample distribution. In an ideal case, if there were no overlap between these two distributions we could find a perfect threshold as the decision boundary to classify the samples into normal and abnormal categories. With existing overlap, however, an ROC curve can help to select the decision boundary that makes the best trade-off between true and false positive rates. Taking a closer look at the difference between the three autoencoders, Figure 9B shows boxplots of the MSE for normal and abnormal samples in the test set, using different numbers of latent features. Here it shows how separable the two classes are based on reconstruction error alone. For the 2-dimensional autoencoder, the interquartile ranges are separable, but there is a significant overlap when considering the whiskers. For the 16-dimensional autoencoder, the separation is greatly improved with minimal overlap between the whiskers. The 128-dimensional autoencoder, however, does not show significant separation and would be impossible to accurately classify the two classes based on MSE alone.

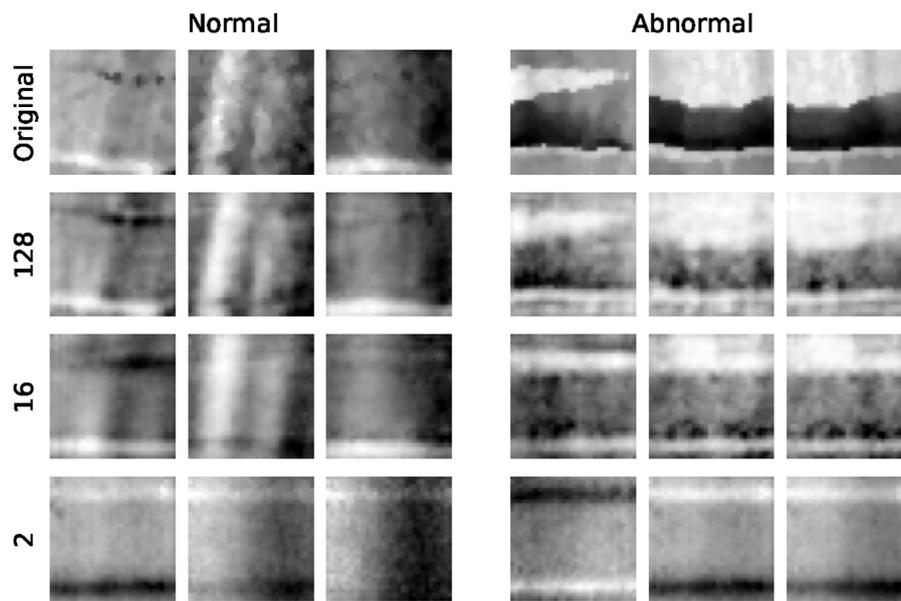


FIGURE 8 The resulting reconstructions from the autoencoders with various latent sizes are compared for both normal and abnormal test samples.

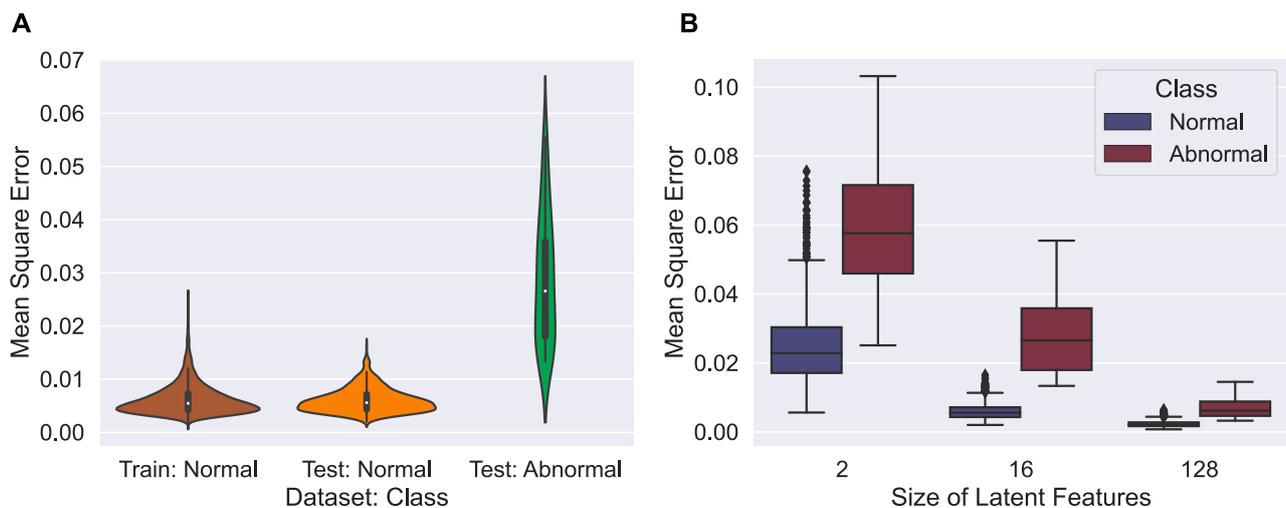
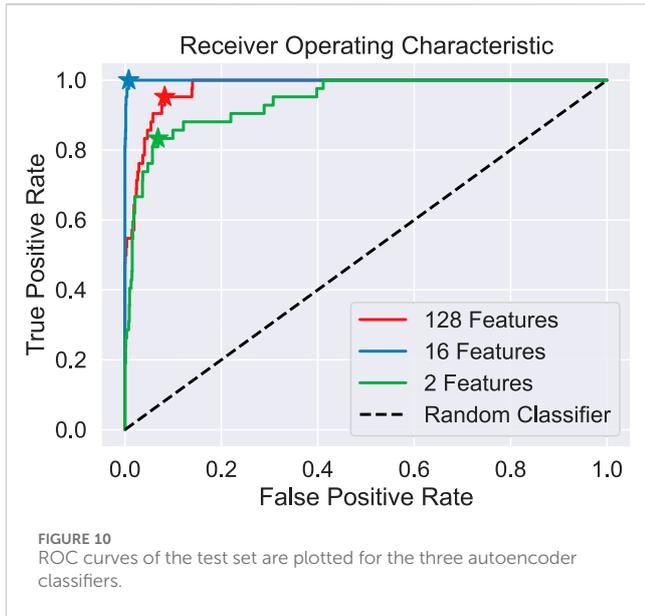


FIGURE 9 Distributions of MSE for different input types and latent sizes are presented for comparison. (A) Comparing test and training. (B) Comparing different latent sizes.

Figure 10 shows the ROC curves for each of the three autoencoders and the selected best threshold points shown as stars. These results further demonstrate that classification performance does not correspond with reconstruction performance, as the 2D and 128D autoencoders' ROC curves have worse classification performance than the 16D model. The best-performing model is the autoencoder with a 16-dimensional latent space, achieving a high true positive rate with a low false positive rate.

Classification results of each autoencoder with the selected threshold values found by the ROC curves are summarised in Table 1. The table reports precision, recall, F1 score, area under the ROC curve (AUC), and selected threshold for each classifier.

Figure 11 shows the effect of the latent vector size on the performance of the model. Figure 11A suggests that larger latent dimensions will produce lower reconstruction errors. However, an accurate classification model does not require the lowest reconstruction error, but a moderate reconstruction that leads to

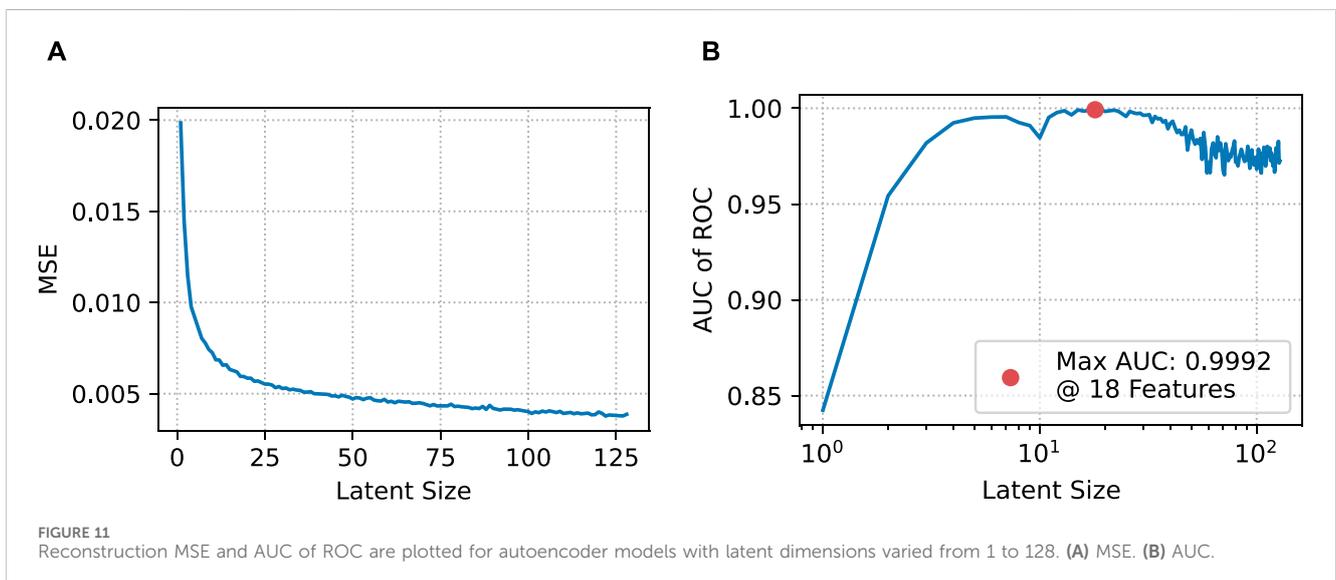


better classification performance. In **Figure 11B** the best latent dim is found by calculating the minimum AUC of the ROC curve while varying the latent dim.

Figure 12 shows the classification confusion matrix while using the optimal latent size. The off-diagonal values in this matrix are low which shows that most of the samples from both normal and abnormal classes are classified correctly.

TABLE 1 Test metrics of convolutional autoencoder.

Latent Dim.	AUC	Best Threshold	Precision	Recall	F1-Score
128	0.978	3.86	0.918	0.929	0.923
16	0.999	13.69	0.992	0.976	0.984
2	0.942	42.08	0.922	0.810	0.862



3.2 Defect localization

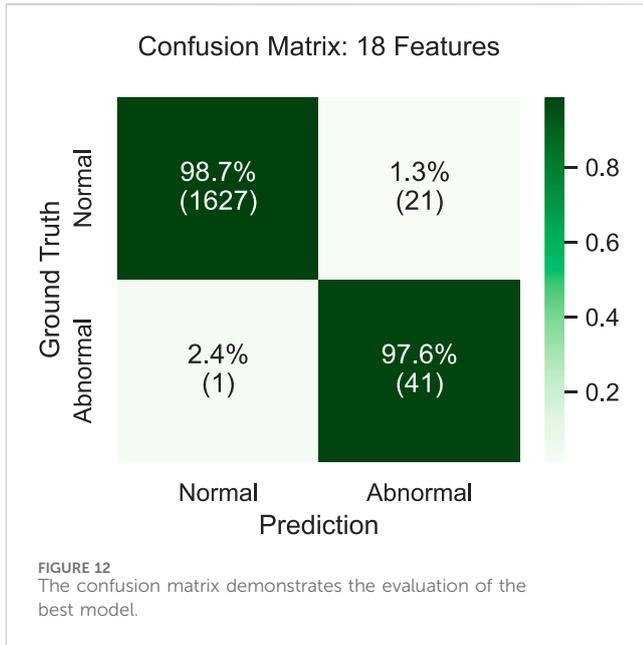
Figure 13 illustrates the results of the anomaly detector on a 2D depth map. The colour of each point indicates the normalized MSE for reconstructing a small window around the point with the anomaly detector. As can be seen, the defective areas have a large density of points with a higher MSE. This information can be used to detect these areas while ignoring individual outlier values.

In **Figure 14** the process of detecting the defects from the anomaly map is illustrated. The elevation in each curve represents the MSE value for one tow (represented by colour in the previous figure). The arrows show the detected blobs after applying the Derivative of Gaussians method. It is observed that only the areas with an extended length of high MSE values are detected as blobs.

Figure 15 shows the final output of the computer vision pipeline, comparing the annotated defect bounding boxes (ground truth) with the predicted bounding boxes.

3.3 Qualitative comparison

The proposed framework is unique and to the best of our knowledge, no other studies have implemented an end-to-end unsupervised defect detection method for AFP inspection. Unfortunately, there are no publicly available datasets in this domain to serve as a benchmark for AFP inspection tasks. Additionally, the dataset used in this work is insufficient for training supervised learning models, which constitute the



majority of current studies in this field. For these reasons, an explicit quantitative comparison of our method with other state-of-the-art approaches is not possible. However, a qualitative comparison of the most relevant studies is presented in Table 2, showcasing the advantages of our framework.

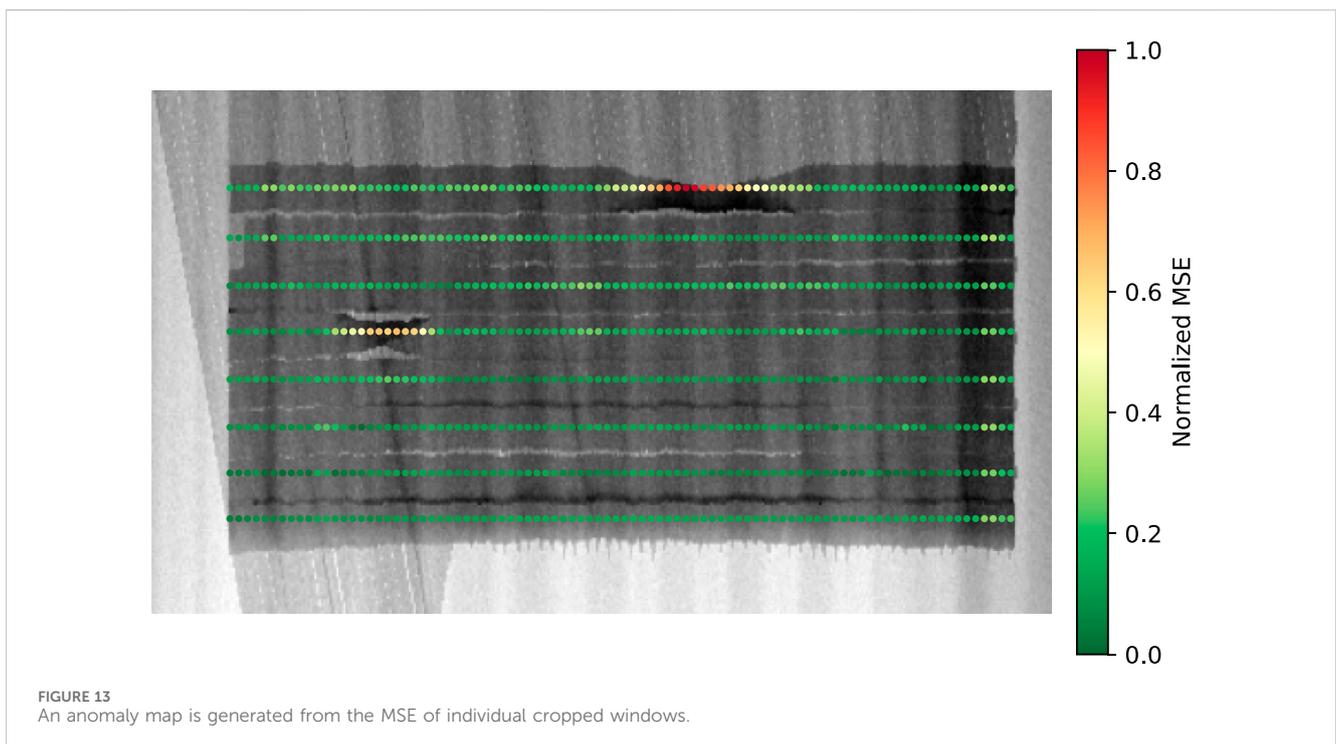
In regard to other defect detection methods, the main advantages of our proposed approach stem from the unsupervised learning process which enables learning with data limitations. Foremost, our method detects all types of

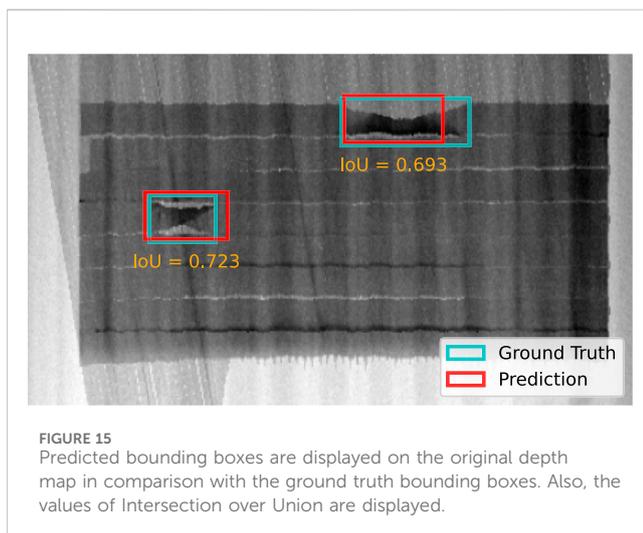
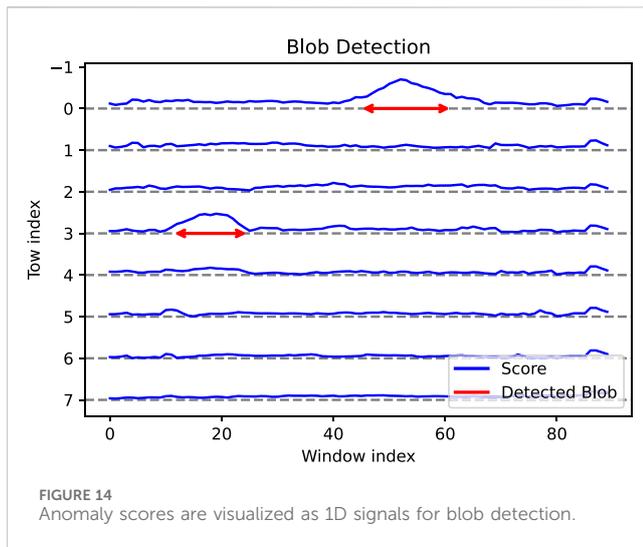
surface anomalies in AFP, whereas existing methods are limited to specific defect types. Additionally, unlike other methods, ours does not require labelling which is time-consuming and prone to errors. Moreover, our proposed framework works with fewer composite scans and it does not need any samples of defects.

Some methods implement semantic segmentation which requires explicit pixel-wise labelling. This is not necessarily needed, as our method provides sufficient localization with bounding boxes and anomaly maps. Unlike other methods, ours does not classify the defects, however, a separate classification module could easily be integrated using the detected bounding boxes. Besides, in an industry where the majority of the inspected parts are non-defective, directly detecting the defective parts can reduce most of the effort.

The proposed defect detection approach relies on layer-by-layer inspection during the AFP process. Severe defective structure should be detected and corrected before applying the next layer. If additional AFP layers are added on top of defects, their accumulation should leave a footprint that will affect the latest layer. Depending of the type of defects present underneath and how they combine with each other's, the surface will appear differently from normal AFP structure and the proposed detection method should be trained to identify it as an anomaly. Note that, for the purposes of this initial study, the selected laminate is thin, resulting in little to no accumulation of defects in underlying layers and thus, the impact is negligible. Datasets containing many layers are required to conduct a study on this question.

AFP has planned and unplanned gaps and overlaps. This study purposefully does not include planned gaps and overlaps. When draping prepreg tows across a surface with compound curvature, planned gaps or overlaps are expected based on the





surface coverage methodology. These are features of the AFP manufacturing process and typically occur along course edges or at individual tow ends. Deviations of individual tows from the expected fibre paths can produce unplanned gaps and overlaps. Planned gap and overlap size, location, and engineering tolerance are all known. This study used a simplified lamination sequence on a flat tool that does not include planned gaps and overlaps. As such, they are not considered in the evaluation. Only unplanned gaps and overlaps caused by individual tow movement during the lamination process are considered in the evaluation. In the next phase, training the suggested autoencoder with a dataset of curved composites that incorporates planned gaps and overlaps will lead the autoencoder to perceive them as normal structure. Any deviation from this predefined structure, such as planned gaps and overlaps, is categorized as abnormal. Consequently, the anomaly detector will specifically pinpoint and identify unplanned overlaps and gaps. If these unplanned gaps and overlaps exceed the manufacturing threshold, they are defects.

4 Conclusion

This paper introduces a practical and novel method for the inspection of composite materials manufactured by Automated Fibre Placement (AFP). The AFP process is susceptible to various types of defects which can significantly impact the final product's quality, necessitating thorough inspection of the composite parts. Manual human inspection has traditionally been employed for this purpose, but it is time-consuming, labour-intensive, and prone to human errors. To enhance the efficiency, accuracy and reliability of AFP, the development of an automated inspection system is crucial. Current inspection procedures mostly utilize profilometry technologies like laser scanning, thermal imaging, and optical sensors to generate visual measurements of the part's surface. The data used in this work is obtained from a laser scanner that operates based on OCT technology, though the framework presented is general and can be adapted to other types of profilometry data.

In AFP inspection, robust and generalized supervised learning methods are infeasible due to limitations in available labelled data. Anomaly detection methods, on the other hand, can circumvent this challenge by focusing on learning the structure of normal samples to identify any abnormalities. The proposed computer vision framework detects individual tows in AFP composites and creates a dataset of sub-images by sliding a window along the center of each tow. The extracted data is then used to train an autoencoder designed to detect anomalies. Using the same sliding window procedure, the autoencoder produces anomaly scores for local regions of the composite part. These scores are aggregated to form an anomaly map of the full image. This anomaly map can then be used as an explicit indication tool by an operator. We further process this anomaly map using a 1D blob detection algorithm to generate bounding boxes around defects.

Compared to other state-of-the-art automated inspection methods in AFP, our approach offers several advantages. Since the autoencoder learns the inherent structure of normal tows, it is capable of detecting all anomalies, unlike other methods which can only detect defects specific to their training data. Furthermore, our suggested framework operates with fewer composite scans and eliminates the requirement for defect samples. For verification purposes the autoencoder is evaluated in a classification task, achieving over 98% classification accuracy. Additionally, the overall framework is implemented on a set of test samples where bounding boxes generated by the method achieve an Intersection over Union (IoU) of 0.708. This demonstrates sufficient accuracy in the localization of detected defects.

This paper outlines a novel defect detection approach for AFP and emphasizes its practicality, particularly in addressing data limitations. There are several potential directions for future research to enhance the inspection system's capabilities and extend its relevance to broader domains. To improve dataset quality and quantity, we suggest investigating data engineering techniques like data augmentation and synthetic data generation. Additionally, while the current system identifies anomalies, it lacks the ability to classify specific defect types. To address this, we recommend utilizing the generated bounding boxes to collect training data for developing a classification model. We also

TABLE 2 A comparison of existing learning-based AFP defect detection methods with the proposed framework.

References	Detection	Classification	Localization	Segmentation	Detects All Abnormalities	Labeling not Required	Defects not Required	Dataset Size	Results	
									Classification	Localization
Meister and Wermes (2023)	✓	✓	-	-	-	-	-	469	Acc: 95.12%	-
Zhang et al. (2022)	✓	✓	✓	-	-	-	-	3,000	-	mAP ⁴ : 93.1%
Tang et al. (2022)	✓	✓	✓	✓	-	-	-	43	Pr ² : 84.4%, Re ³ : 77.1%	mIoU ⁵ : 0.776
Sacco et al. (2020)	✓	✓	✓	✓	-	-	-	800	Average Acc: 49.1%	-
Schmidt et al. (2019)	✓	✓	-	-	-	-	-	12,000	Acc: > 92%	-
This work	✓	-	✓	-	✓	✓	✓	44	Acc: 98.7%	mIoU: 0.708

1: Accuracy, 2: Precision, 3: Recall, 4: mean Average Precision, 5: mean Intersection over Union.

recommend adapting our framework for use in industries that share similar tape-by-tape structures. This can lead to enhancements in defect identification and quality assurance across different sectors.

Data availability statement

The data analyzed in this study is subject to the following licenses/restrictions: The data generated and analyzed in this paper is proprietary and due to its commercialization potential, is not made publicly available. The information is protected to maintain its commercial value. We regret any inconvenience this may cause and appreciate your understanding of the importance of preserving its proprietary nature. Requests to access these datasets should be directed to HN, najjaraan@uvic.ca.

Author contributions

AG: Project administration, Writing–review and editing, Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Software, Validation, Visualization, Writing–original draft. TC: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Software, Validation, Visualization, Writing–original draft, Writing–review and editing. LJ: Methodology, Writing–review and editing, Supervision. MR: Writing–review and editing. GL: Writing–review and editing. HN: Writing–review and editing, Funding acquisition, Supervision.

Funding

The author(s) declare financial support was received for the research, authorship, and/or publication of this article. This work is done with the financial support of LlamaZOO Interactive Inc. and Natural Sciences and Engineering Research Council (NSERC) Canada under the Alliance Grant ALLRP 567583 - 21.

Acknowledgments

We would like to acknowledge the research collaboration of LlamaZOO Interactive Inc., the National Research Council of Canada (NRC), and Fives Lund LLC.

Conflict of interest

Author LJ was employed by LlamaZOO Interactive Inc. Author GL was employed by Fives Lund LCC.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The author(s) declared that they were an editorial board member of Frontiers, at the time of submission. This had no impact on the peer review process and the final decision.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated

organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Ahang, M., Jalayer, M., Shojaeinasab, A., Ogunfowora, O., Charter, T., and Najjaran, H. (2022). Synthesizing rolling bearing fault samples in new conditions: a framework based on a modified cgan. *Sensors* 22, 5413. doi:10.3390/s22145413
- Albuquerque Filho, J. E. D., Brandão, L. C. P., Fernandes, B. J. T., and Maciel, A. M. A. (2022). A review of neural networks for anomaly detection. *IEEE Access* 10, 112342–112367. doi:10.1109/ACCESS.2022.3216007
- Azzeh, J., Zahran, B., and Alqadi, Z. (2018). Salt and pepper noise: effects and removal. *JOIV Int. J. Inf. Vis.* 2, 252–256. doi:10.30630/joiv.2.4.151
- Böckl, B., Wedel, A., Misik, A., and Drechsler, K. (2023). Effects of defects in automated fiber placement laminates and its correlation to automated optical inspection results. *J. Reinf. Plastics Compos.* 42, 3–16. doi:10.1177/07316844221093273
- Brasington, A., Sacco, C., Halbritter, J., Wehbe, R., and Harik, R. (2021). Automated fiber placement: a review of history, current technologies, and future paths forward. *Compos. Part C. Open Access* 6, 100182. doi:10.1016/j.jcomc.2021.100182
- Chow, J. K., Su, Z., Wu, J., Tan, P. S., Mao, X., and Wang, Y. H. (2020). Anomaly detection of defects on concrete structures with the convolutional autoencoder. *Adv. Eng. Inf.* 45, 101105. doi:10.1016/j.aei.2020.101105
- Dalal, N., and Triggs, B. (2005). "Histograms of oriented gradients for human detection," in *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, 886–893. doi:10.1109/CVPR.2005.177
- Danker, A., and Rosenfeld, A. (1981). "Blob detection by relaxation," in *IEEE transactions on pattern analysis and machine intelligence PAMI-3*, 79–92.
- Deng, X., Ma, Y., and Dong, M. (2016). A new adaptive filtering method for removing salt and pepper noise based on multilayered PCNN. *Pattern Recognit. Lett.* 79, 8–17. doi:10.1016/j.patrec.2016.04.019
- Duda, R. O., and Hart, P. E. (1972). Use of the Hough transformation to detect lines and curves in pictures. *Commun. ACM* 15, 11–15. doi:10.1145/361237.361242
- Heinecke, F., and Willberg, C. (2019). Manufacturing-induced imperfections in composite parts manufactured via automated fiber placement. *J. Compos. Sci.* 3, 56. doi:10.3390/jcs3020056
- Hough, P. V. (1959). "Machine analysis of bubble chamber pictures," in *Proc. Of the international conference on high energy accelerators and instrumentation*, 1959, 554–556.
- Kong, H., Akakin, H. C., and Sarma, S. E. (2013). A generalized Laplacian of Gaussian filter for blob detection and its applications. *IEEE Trans. Cybern.* 43, 1719–1733. doi:10.1109/tsmcb.2012.2228639
- Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vis.* 60 (2), 91–110. doi:10.1023/B:VISI.0000029664.99615.94
- Meister, S., Möller, N., Stüve, J., and Groves, R. M. (2021a). Synthetic image data augmentation for fibre layup inspection processes: techniques to enhance the data set. *J. Intelligent Manuf.* 32, 1767–1789. doi:10.1007/s10845-021-01738-7
- Meister, S., and Wermes, M. (2023). Performance evaluation of CNN and R-CNN based line by line analysis algorithms for fibre placement defect classification. *Prod. Eng.* 17, 391–406. doi:10.1007/s11740-022-01162-7
- Meister, S., Wermes, M. A. M., Stüve, J., and Groves, R. M. (2021b). Review of image segmentation techniques for layup defect detection in the Automated Fiber Placement process. *J. Intelligent Manuf.* 32, 2099–2119. doi:10.1007/s10845-021-01774-3
- Palardy-Sim, M., Rivard, M., Lamouche, G., Roy, S., Padioleau, C., Beaudesne, A., et al. (2019a). "Advances in a next generation measurement & inspection system for automated fibre placement," in *Proceedings of the manufacturing & processing technologies conference in the composites and advanced materials expo* (Anaheim, CA, USA: CAMX), 23–26.
- Palardy-Sim, M., Rivard, M., Roy, S., Lamouche, G., Padioleau, C., Beaudesne, A., et al. (2019b). "Next generation inspection solution for automated fibre placement," in *The fourth international symposium on automated composites manufacturing*, 64.
- Rivard, M., Palardy-Sim, M., Lamouche, G., Roy, S., Padioleau, C., Beaudesne, A., et al. (2020). "Enabling responsive real-time inspection of the automated fiber placement process," in *International SAMPE technical conference 2020-june*.
- Sacco, C., Baz Radwan, A., Anderson, A., Harik, R., and Gregory, E. (2020). Machine learning in composites manufacturing: a case study of Automated Fiber Placement inspection. *Compos. Struct.* 250, 112514. doi:10.1016/j.compstruct.2020.112514
- Schmidt, C., Hocke, T., and Denkena, B. (2019). Deep learning-based classification of production defects in automated-fiber-placement processes. *Prod. Eng.* 13, 501–509. doi:10.1007/s11740-019-00893-4
- Tang, Y., Wang, Q., Cheng, L., Li, J., and Ke, Y. (2022). An in-process inspection method integrating deep learning and classical algorithm for automated fiber placement. *Compos. Struct.* 300, 116051. doi:10.1016/j.compstruct.2022.116051
- Tsai, D.-M., and Jen, P.-H. (2021). Autoencoder-based anomaly detection for surface defect inspection. *Adv. Eng. Inf.* 48, 101272. doi:10.1016/j.aei.2021.101272
- Ulger, F., Yuksel, S. E., and Yilmaz, A. (2021). Anomaly detection for solder joints using β -VAE. *IEEE Trans. Components, Packag. Manuf. Technol.* 11, 2214–2221. doi:10.1109/TCPMT.2021.3121265
- Zemzemoglu, M., and Unel, M. (2022). "Design and implementation of a vision based in-situ defect detection system of automated fiber placement process," in *2022 IEEE 20th international conference on industrial informatics (INDIN)* (Perth, Australia: IEEE), 393–398. doi:10.1109/INDIN51773.2022.9976182
- Zhang, Y., Wang, W., Liu, Q., Guo, Z., and Ji, Y. (2022). Research on defect detection in automated fiber placement processes based on a multi-scale detector. *Electronics* 11, 3757. doi:10.3390/electronics11223757