DEFECTS DETECTION IN ROTOR COMPOSITE PARTS USING INSTANCE SEGMENTATION

Nicolas GRISELIN, Airbus Helicopters, Marignane, France
Denis DUBOIS, Alten, Paris, France
Pierre BARBIER, Airbus Helicopters, Marignane, France

Abstract

During the manufacturing process of rotor composite parts, each part is systematically controlled using a Radiographic Testing (RT) approach. This is a non-destructive testing (NDT) method, which uses x-rays to examine the internal structure of manufactured components identifying any flaws or defects within the material. In this paper, an instance segmentation approach is used to detect a specific defect in radiographic images of composites parts, based on a Mask R-CNN model. Instance segmentation represents a significant advancement in computer vision compared to more conventional approaches (such as classification or object detection). Application to an industrial case is presented here with a precision to predict "defects" of 85%, a recall of 96% and a F1 score of 90%. We also introduce an original labeling technique well suited for industrial purpose.

1. NOTATION

Acronyms:

AI	Artificial Intelligence
CNN	Convolutional Neural Network
CV	Computer Vision
DICOM	Digital Imaging and Communications in Medicine
ERF	European Rotorcraft Forum
IoU	Intersection over Union
NDT	Non-Destructive Testing
R-CNN	Region-based CNN
ROI	Region of Interest
RPN	Region Proposal Network
RT	Radiographic Testing

2. INTRODUCTION

With the rapid advancements in artificial intelligence (AI) over recent years, the field of Non-Destructive Testing (NDT) has undergone a significant transformation, particularly through the application of AI techniques for Human Augmentation in defect detection. AI-driven tools and systems now significantly enhance the capabilities of human inspectors by providing advanced defect detection algorithms. These technologies improve the accuracy and efficiency of NDT processes, enabling quicker and more reliable identification of defects and anomalies in materials and structures.

In previous work, classification approaches were used, where algorithms were trained to classify inspection data into categories such as "defective" or "non-defective." While effective to a certain extent, these methods often lacked the granularity needed for precise defect localization.

In contrast, current work leveraging instance segmentation offers a significant improvement. Instance segmentation not only classifies the data but also provides detailed information about the exact location and shape of each defect. This method uses advanced deep learning models to produce pixel-wise masks for each defect, allowing for much more precise and actionable insights. For example, instead of simply identifying an area as defective, instance segmentation can highlight the exact boundaries of a defect, enabling more accurate defect characterization and assessments. This not only augments the inspectors' skills but also reduces the time required for thorough inspections.

3. COMPOSITE PARTS MANUFACTURING

3.1. Overall Process

Composite materials are now widely used in industry and their uses are no longer limited to parts with a simple geometry. Many structures with complex geometries are made of composites and can be found on some of our helicopter's parts.

The manufacture of those parts involves the use of various manufacturing techniques, including ply preparation, molding, curing, de-molding, machining of parts, quality control and finishing phases such as painting.

It is important that the manufactured parts meet the quality and safety requirements necessary for reliable operation of the helicopter. Parts must be tested and inspected to ensure that they meet the appropriate specifications and safety standards.

3.2. Quality Control

Inspection of the parts is performed through the quality control phase. During this phase a number of radiographic pictures are taken to fully cover the whole surface of the part. Till today, each image is analyzed using a dedicated viewing software that allows visualizing the image as well as modifying the contrast and brightness of the image in order to identify any potential defect within the material. This task is performed by a specialist who needs specific training to be able to recognize any given defect within the radiographic image.

This can be a time-consuming process to analyse all the x-ray images taken on a single part. It is also a very human dependent process where you need skilled qualified inspector to be able to detect defects on images, to characterize the defects and to be able to interpret the predefined criteria for acceptance (or refusal) of the part.

3.3. Ripple Defect

The complex geometry of composite parts can lead to problems of ply ripple formation in some areas of the part. Understanding the mechanisms of ripple formation is a lever to optimize the forming process in order to limit the formation of these waves as much as possible.

However, this does not ensure that ripple formation will not occur and the aeronautical manufacturers have to deal with these non-conformities. All the interest is to succeed in identifying precisely these ripples in the composite structures in order to conclude about their viability and allow a better management of non-conformities.

In the literature, two types of ripples are distinguished: ripple in the plane of the ply, called "waviness", and ripple out-of-plane, called "wrinkle". In this study we will focus on "wrinkle" defects (see Figure 1).

Figure 1: Ripple ("wrinkle") defect

3.4. Defects Detection

Traditional visual inspection methods have long been utilized for quality control in industrial product manufacturing. The state of the art in defect detection based on machine vision and latest advancements in this field are described in Ref. 1, offering a comprehensive understanding of employing deep learning techniques for defect classification, localization, and segmentation.

3.4.1. X-Ray Images

Radiography is one of the common techniques used to detect defects in composite parts. It uses x-rays to penetrate the material and create an image that shows in-depth details through the different layers of the part. The resulting image is stored in a DICOM file.

DICOM is a standard format for storing, exchanging and displaying digital medical images, including radiographic images. It allows the storage of image information, such as the technical parameters used to produce the images, patient data (or product data in our case), serial data and acquisition data. DICOM images can be viewed on a computer screen using specialized viewing software, and can also be used for archiving. This rich set of features make it a versatile and widely-adopted format for radiographic imaging. It was thus selected for our application and has been used for many years now during the manufacture of our composite parts.

In our application, we use 16-bit grayscale image of 1024 x 1024 pixels DICOM files. This means that each file contains a digital radiographic image where each pixel in the image is encoded in 16 bits, which allows for 65,536 different shades of gray to be stored. This can improve image quality by allowing finer gray levels to be represented and reducing quantization artifacts.

For the purpose of this study, we have collected x-ray DICOM files from a specific area of a composite part.

3.4.2. Detection using Classification

Classification is widely used in various applications such as image tagging, content-based image retrieval, and large-scale image categorization. It serves as a foundational building block in many computer vision systems and was used in our previous work (see Ref. 2) where we described a three steps approach to detect defects in radiographic images of composite parts using simple Artificial Intelligence CNN models for image classification. However, the use of such an approach can bring limitation to count the number of defects as well as to localize and characterize precisely each defect.

3.4.3. Detection using Segmentation

Segmentation techniques, in image processing and computer vision, allows for partitioning images into meaningful regions or segments. These methods can have an important role in tasks requiring detailed object localization and segmentation, such as medical image analysis and object counting.

Instance segmentation and semantic segmentation are two prominent segmentation techniques used in computer vision for image analysis and understanding. While both semantic segmentation and instance segmentation operate at the pixel level to partition images into meaningful regions, they differ in their level of granularity.

Semantic segmentation focuses on assigning class labels to pixels based on semantic categories, while instance segmentation goes a step further by providing separate segmentation masks for each object instance within the same class. Both techniques are valuable tools in computer vision and have their respective applications depending on the specific requirements of the task at hand.

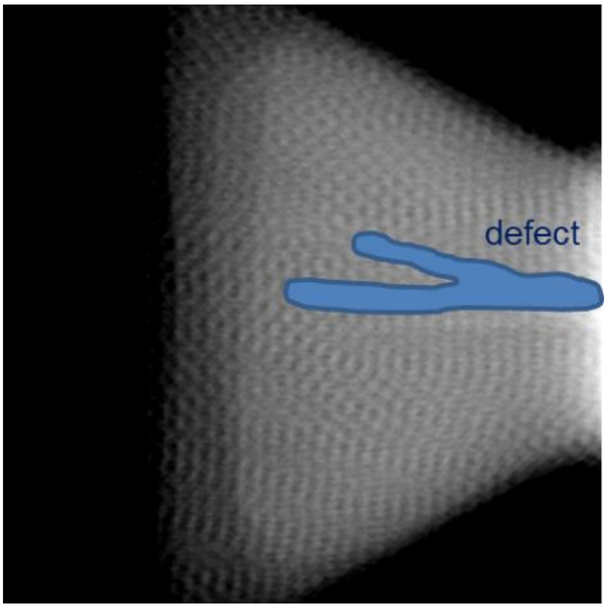

Figure 2: Semantic Segmentation for defect detection

An illustration of the difference between Semantic Segmentation and Instance Segmentation applied to our study is presented in Figure 2 and Figure 3.

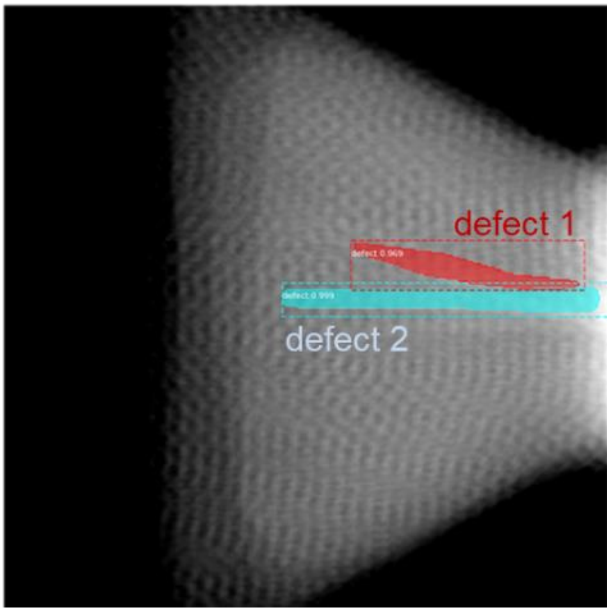

Figure 3: Instance Segmentation for defect detection

While Semantic Segmentation can improve localization of defects in the radiography without discriminating between different instances of objects belonging to the same class (see Figure 2), it does not allow counting the number of defects and characterizing each defect.

On the other hand, Instance Segmentation (see Figure 3) allows classifying each pixel into semantic categories and distinguish between individual object instances of the same class. This gives us the possibility to identify and differentiate each object instance presents in the image and thus characterize each defect according to their length and width.

Recent advances in deep learning, particularly with the advent of convolutional neural networks (CNNs) and architectures like R-CNN, have greatly improved the accuracy and efficiency of instance segmentation models (see Ref. 3).

In this work, we have implemented an Instance Segmentation approach using a Mask R-CNN model to detect defects but also to be able to localize and characterize each defect in the radiography.

4. SEGMENTATION APPROACH

4.1. Image Pre-treatments

4.1.1.Image Extraction

We have used the exact same images database as in our previous study (see Ref 2). This database consists of 16 bits full size grayscale images (with size of 1024 x 1024 pixels) containing one type of defect with different size and location.

These images were extracted from the original DICOM file collected after Radiographic Testing of the composite part and converted to .png format to reduce the size of the image file without any quality loss (lossless image).

4.1.2.Image Filtering

After we remove unnecessary borders, normalize grey level, invert intensity scale, an image filtering is performed to filter pixel intensity within defect density range. From a raw image where it is difficult (if not impossible for human eyes) to detect any defect (see Figure 4) we produce an image where defects can clearly be seen (see Figure 5). A similar approach is usually performed by the inspector when looking for defects in x-ray images in order to visually detect any defect. It has been automatized here for the purpose of our study.

Figure 4: Image extracted from DICOM file

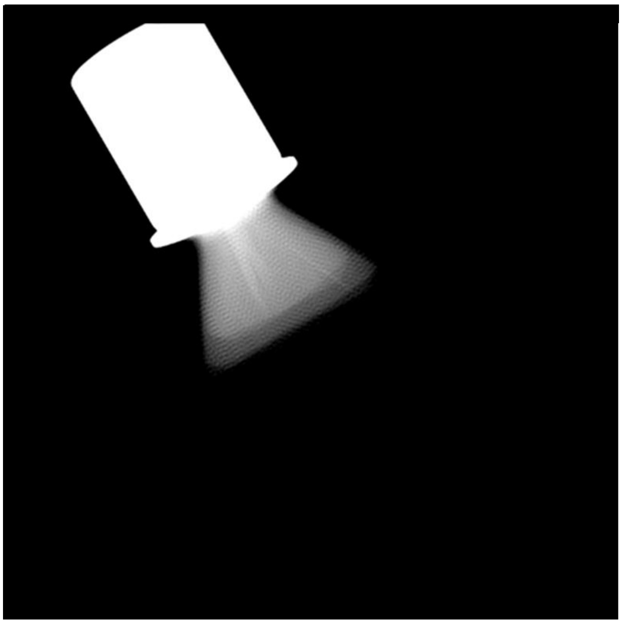

Figure 5: Filtered image after pre-treatment

4.1.3.Image Rotation

Image rotation is performed to align defects in the same direction: one direction is chosen and all images are rotated in the same direction (see Figure 6).

This preprocessing step creates a standardized dataset, which simplifies the learning process for machine learning models, particularly CNNs. Uniform orientation reduces the variability that the models need to account for, focusing their capacity on recognizing and classifying defects accurately. This standardization enhances feature extraction, making edge detection, texture analysis, and shape descriptors more reliable and meaningful.

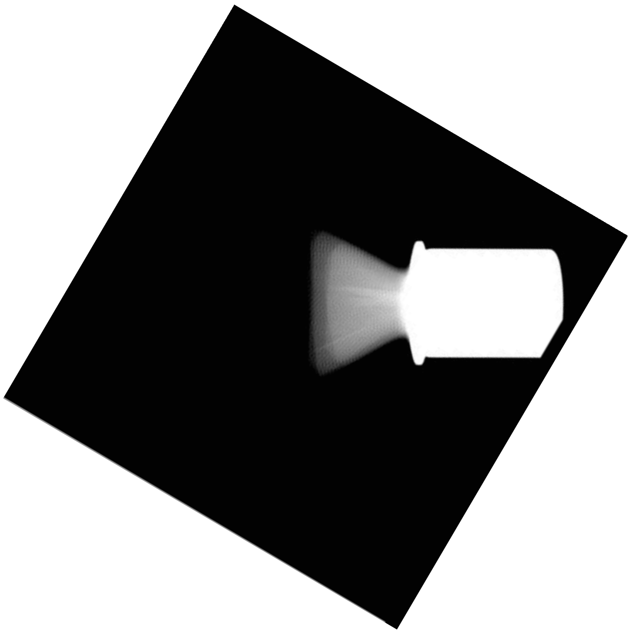

Figure 6: Image rotation

Additionally, image rotation facilitates more effective data augmentation by allowing techniques such as flipping, scaling, and cropping to be applied without introducing unnecessary variability. This results in a more robust training dataset.

4.1.4.Image Cropping

Images are cropped to focus on the region of interest (ROI) where defects are located. A 300 x 300 pixels crop is applied to extract the area where defect can be located (see Figure 7).

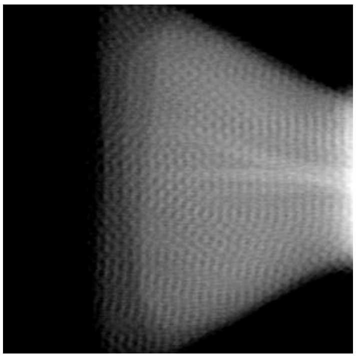

Figure 7: Cropped image (300 x 300 pixels)

Cropping offers significant advantages for computer vision: by eliminating irrelevant parts of the image, cropping reduces the amount of data that needs to be processed, which enhances computational efficiency. This targeted approach allows CV algorithms to concentrate solely on the critical areas, improving the accuracy and precision of defect detection.

4.2. Labelling

Data labeling for instance segmentation involves creating detailed annotations that identify and delineate each object instance within an image, specifying the precise pixel boundaries. Unlike semantic segmentation, which assigns a class label to each pixel without distinguishing between different instances of the same class, instance segmentation requires each object instance to be uniquely labeled.

The usual labeling process for instance segmentation primarily involves drawing polygonal masks around each object instance in the images (see Figure 8).

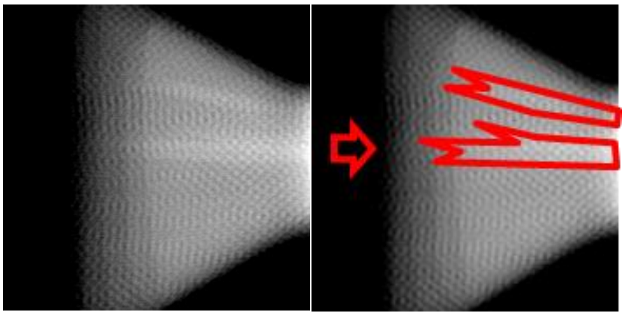

Figure 8: Polygonal Masks

Specialized annotation tools can be used like LabelMe or VGG Image Annotator (VIA) to achieve this. The process involves to manually click around the boundary of each object to create a closed polygon, ensuring that the mask precisely delineates the object's shape. Each polygon is then assigned a class label and a unique instance identifier to differentiate between multiple instances of the same class. This meticulous process of drawing polygons ensures that the model can accurately learn to identify and segment objects based on their exact boundaries.

However, this process can be quite time consuming and is not very well suited for an industrial application where the detailed contour of the defect is not the most important concern of the inspector but rather the start and end of each defect (thus allowing to measure the length of the defect and its position in the composite parts).

In the future, labeling of such type of defect will probably need to be performed during the Quality Control process by qualified inspectors who have been trained to perform such task in order to increase the confidence level of our algorithm in regards to future EASA requirements for safety-related machine learning applications (see Ref. 4).

For this reason, and to make sure that labeling will not add too much extra time for the inspector during the Quality Control process, we propose a new simplified way to perform the labeling of such a defect using existing task performed by the inspector. During current inspection process, each defect is systematically measured by identifying the start and end of each defect using existing software that provides the corresponding length of the defect. By using those 2 points (already spotted by the inspector) we define a line with a given width (that has been optimized during the study) to create a mask that covers most of the defect area (see Figure 9).

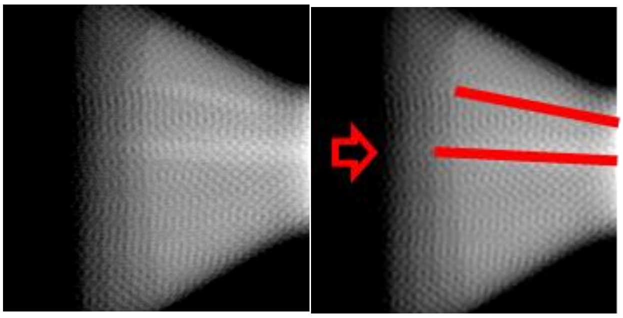

Figure 9: Simplified Labeling

4.3. Mask Creation

From the labeling of each defect, we created a pixel-level segmentation mask. An example is presented in Figure 10.

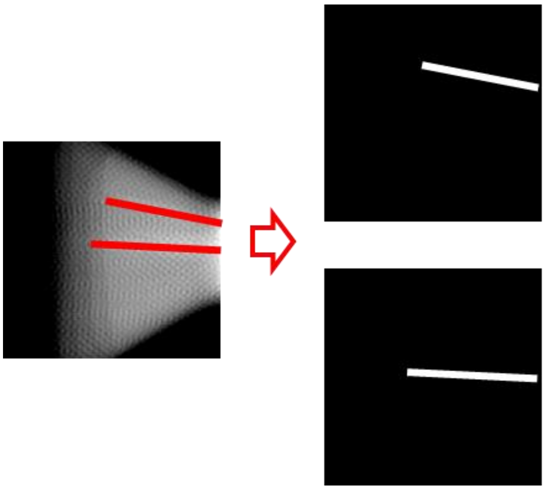

Figure 10: Example of Masks Created

This mask indicates which pixels in the image belong to that particular instance. A unique mask is created for each defect, meaning that for an image with 2 defects, we will have 2 distinct mask images (1 for each defect). By assigning unique labels and generating segmentation masks, instance segmentation algorithms can accurately delineate and separate individual objects within the same class

4.4. Dataset Split

For the training of the model, we only considered images with defects (corresponding to 345 images collected, having one or more defects). This dataset is divided into 3 distinct groups: TRAIN, VALIDATION and TEST images datasets. We used a 60/20/20 ratio to split the dataset. This means that 60% of the images are used for training the model, 20% for validation of the model and the remaining 20% are used for evaluating the model's performance on unseen data.

4.5. Image Augmentation

Image augmentation allows enhancing the diversity of training data without the need for additional image collection. This is a critical technique in deep learning when the amount of data available is quite small.

This process helps models generalize better and avoid overfitting, particularly in tasks like instance segmentation. Several common augmentation techniques used here include flipping, adjusting brightness and contrast, and applying blur and sharpen filters.

Flipping involves creating mirror images by flipping the original image either horizontally or vertically. Horizontal flipping was used here. It is particularly useful for ensuring the model can recognize objects regardless of their orientation.

Brightness and contrast adjustments simulate varying lighting conditions. Increasing or decreasing brightness helps the model handle different lighting scenarios, while contrast adjustments make features either more prominent or subtle, enhancing the model’s ability to focus on significant details under diverse conditions.

Blurring and sharpening are techniques used to alter the clarity of images. Blurring, achieved using Gaussian blur, smoothens the image and helps the model become invariant to minor, irrelevant details. Sharpening, on the other hand, enhances edges and makes features more defined, helping the model to better recognize and delineate objects.

4.6. Instance Segmentation using Mask R-CNN

Mask R-CNN (Mask Region-based Convolutional Neural Network) is a state-of-the-art framework designed for instance segmentation in computer vision. It extends the Faster R-CNN model, which performs object detection by predicting bounding boxes around objects, by adding a branch for predicting segmentation masks.

Mask R-CNN is widely used in applications requiring precise object segmentation, such as autonomous driving, medical image analysis, and video surveillance. It excels in scenarios where it is crucial to distinguish between overlapping objects and segment each one accurately. The ability to predict both bounding boxes and masks makes Mask R-CNN particularly valuable in environments where pixel-level accuracy is necessary. Its architecture efficiently balances accuracy and computational complexity, making it a powerful tool for instance segmentation tasks applied to our industrial application.

The architecture of a Mask R-CNN model comprises three main components: a backbone network, a region proposal network (RPN), and task-specific heads. The backbone network, typically a ResNet, acts as a feature extractor. It processes the input image to produce a feature map. The RPN then scans these feature maps to propose candidate object regions, generating region of interest (RoI) proposals. For each RoI, the model has three separate branches (heads): one for classification (assigning a class label), one for bounding box regression (refining the coordinates of the bounding box), and one for mask prediction (producing a binary mask that highlights the object within the bounding box).

Training a Mask R-CNN model involves preparing a labeled dataset where each object instance in the training images is annotated with a class label, a bounding box, and a segmentation mask. The model is trained using a multi-task loss function that combines classification loss, bounding box regression loss, and mask prediction loss. This allows the network to learn not only to detect objects and classify them but also to generate precise segmentation masks. Transfer learning is commonly employed, starting with a pre-trained backbone network (here we have used a backbone trained on the COCO dataset) to accelerate convergence and improve performance.

5. RESULTS

The trained model was applied on the 69 images from the test dataset (corresponding to 132 defects as one image can have multiple defects included).

To evaluate the performance of the model we have calculated the Intersection over Union (IoU) metric to measures the overlap between the predicted segmentation mask and the ground truth mask, providing a clear indication of how well the model has delineated each object.

To calculate IoU, the intersection and union of the predicted and ground truth masks are computed. The intersection is the area where both masks overlap, while the union is the total area covered by both masks. IoU is then calculated as the ratio of the intersection area to the union area:

IoU = Area of Intersection / Area of Union

An IoU score ranges from 0 to 1, where 1 indicates a perfect match between the predicted and ground truth masks, and 0 indicates no overlap at all. Higher IoU scores reflect better segmentation performance, as they indicate a greater degree of overlap and therefore more accurate predictions.

In the context of instance segmentation using models like Mask R-CNN, IoU is used not only to evaluate individual object masks but also to determine whether a predicted mask is a true positive (IoU above a certain threshold, commonly 0.5) or a false positive. This makes IoU a fundamental metric for assessing and comparing the effectiveness of different segmentation models and their configurations.

The results on our test dataset are presented in the confusion matrix below (see Table 1) for an IoU threshold of 0.3.

Table 1: Confusion Matrix for test dataset

		Predicted	
		No Defect	Defect
Actual	No Defect		FP = 23
	Defect	FN = 6	TP = 126

Where TP is true positive, FP is false positive, and FN is false negative.

The following performance metrics have been derived from the confusion matrix:

Precision:
TP / (TP + FP)

Recall (Sensitivity or True Positive Rate):
TP / (TP + FN)

F1 Score:
2 * ((Precision * Recall) / (Precision + Recall))

In our application, we want to have the number of False Negative as low as possible because the cost of false negatives is much higher than false positive. For this reason, the Recall is a good metric to use. However, we don't want to have too many False Positive predicted by the model to avoid the Inspector losing confidence in the model. Thus, we also consider Precision (as it focusses on TP and FP) and F1-score as it considers both precision and recall.

The results are rather good with a *Precision* to predict "defects" of 85%, a *Recall* of 96% and a *F1-score* around 90%.

6. CONCLUSIONS

This paper presents a concrete application of defects detection in rotor composite parts using Mask R-CNN Instance Segmentation model.

1. We presented a multi-steps methodology to build a Mask R-CNN model, including image pre-treatment, image labeling, image augmentation and training of the model.
2. We proposed a new, simplified approach to perform the labeling of ripple defect using a 2 points-line creation from existing task performed by the inspector.
3. We described the metrics used to evaluate the performance of our model (based on calculating the ioU between the predicted segmentation mask and the ground truth mask) and presented the rather good results of our model applied on the Test set: *Precision* to predict "defects" of 85%, a *Recall* of 96% and a *F1-score* around 90%

Instance segmentation models (like Mask R-CNN) can significantly enhance the Radiographic Testing approach in industrial applications. These models offer several advantages:

- Precision: by providing pixel-level precision for defect localization and characterization, enabling accurate measurements and better-informed decisions.
- Fast Detection: by speeding up the inspection process, enhancing efficiency and consistency, and reducing human error.
- Enhanced Visualization: by generating detailed segmentation masks for better defect visualization, facilitating clear and comprehensive inspection results.

Further development will focus on "Explainable AI" and the application of such an Artificial Intelligence model to industrial and certification point of view. For certification matter, the EASA has published a specific guideline (see Ref. 4) that gives inspiring guidance for development of similar AI applications. Other defects will also be included in the detection algorithm in future work.

Author contact:

Nicolas GRISELIN nicolas.griselin@airbus.com

Denis DUBOIS denis.dubois.external@airbus.com

Pierre BARBIER pierre.p.barbier@airbus.com

7. REFERENCES

1. Ren, Z., Fang, F.Z., Yan, N. and Wu, Y., "State of the Art in Defect Detection Based on Machine Vision", Vertical International Journal of Precision Engineering and Manufacturing-Green Technology, 2022
2. Griselin, N., and Barbier, P., "Defects Detection in Rotor Composite Parts using AI", Vertical Flight Society 79th Annual Forum Proceedings, West Palm Beach, FL, 2023
3. Ren, Z., Fang, F.Z., Yan, N. and Wu, Y., "Deep CNN-based visual defect detection: Survey of current literature", Computers in Industry 148, 2023
4. EASA Concept Paper: First usable guidance for Level 1&2 machine learning applications, easa.europa.eu/ai, 2023