

Automatic void content assessment of composite laminates using a machine-learning approach

João M. Machado^{a*}, João Manuel R. S. Tavares^{a,b}, Pedro P. Camanho^{a,b}, Nuno Correia^a

^a Instituto de Ciência e Inovação em Engenharia Mecânica e Engenharia Industrial, Porto, Portugal

^b Departamento de Engenharia Mecânica, Faculdade de Engenharia, Universidade do Porto, Porto, Portugal

* Corresponding author:

E-mail: jmmachado@inegi.up.pt

Tel: (+351-229578710)

Abstract

Voids have a substantial impact on the mechanical properties of composite laminates and can lead to premature failure of composite parts. Optical microscopy is a commonly employed imaging technique to assess the void content of composite parts, as it is reliable and less expensive than alternative options. Usually, image thresholding techniques are used to parse the void content of the acquired microscopy images automatically; however, these techniques are very sensitive to the imaging acquisition conditions and type of composite material used. Additionally, these algorithms have to be calibrated before each analysis, in order to provide accurate results.

This work proposes a machine-learning approach, based on a convolutional neural network architecture, with the objective of providing a robust tool capable of automatically parsing the void content of optical microscopy images, without the need of parameter tuning.

Results from training and testing datasets composed of microscopy images extracted from three distinct types of laminates confirm that the proposed approach parses void content from microscopy images more accurately than a traditional thresholding algorithm, without the need of a previous calibration step. This work shows that the proposed approach is promising, despite sometimes lower than expected precision in individual void statistics.

Keywords

Microscopy, Deep Learning, Convolutional neural network, Materials characterization

Introduction

Voids are created during manufacturing and can display different morphologies. Research suggests that these characteristics are dependent to the manufacturing process used, as well as the optimization degree of the process parameters [1–3]. Voids have a negative impact on the mechanical properties of composite laminates, especially those that are matrix dominated [4–7], as fatigue resistance and compression strength [7–11]. Therefore, void content assessment is an essential step to monitor the quality of manufactured parts, guaranteeing the reliability of the composite structure.

Optical microscopy is the most commonly employed imaging technique to evaluate void content [1,4,12–17], since it provides reasonable accuracy and detail, and is simple. In order to avoid the time consuming task of evaluating the relative void content of micrography images manually [4,17], a commonly employed technique is automatic image segmentation by pixel intensity thresholding [1,13,18,19]. This technique relies on the different pixel intensities of the composite and the void and is usually user calibrated. However, despite the good results that can be achieved, this methodology can be very sensitive to the illumination conditions during the acquisition of the images as well as the type of material being analysed. Therefore, the calibration of the algorithm parameters (including the threshold value), is a necessary step before the analysis of a given set of images.

Several methods exist to enable the assessment of void content in composite parts, namely density-based methods (acid digestion, matrix burn-off), optical or electron microscopy, ultrasonic testing, thermography, and X-Ray micro-CT.

Non-destructive methods such as ultrasonic testing and thermography have the added advantage of preserving the part, while allowing to estimate void content. On the other hand, although X-ray micro-CT is not a destructive technique by nature, microscopy, X-ray micro-CT and density techniques usually require the partial or total destruction of the composite part in order to assess void content in smaller samples [20].

Another relevant issue in void analysis is the extraction of void characteristics, such as dimensions, shape, and number count. Such analysis requires a high level of detail, which not all analysis techniques can provide, especially when the intended voids are small enough to be located inside or in between the fibre tows. It is known that density-based techniques are not able to provide such data, whereas despite the advancements in ultrasound testing techniques, still ultrasound and thermography usually do not provide the ideal level of detail for such analysis [21,22].

Usually, microscopy and X-ray micro-CT techniques are reported to provide a good level of detail, which enables the accurate measurement and parametrization of void characteristics on the smaller length scales [20,22,23]. Due to its simplicity, lower cost and reasonable accuracy and detail, optical microscopy is still a commonly employed imaging technique to conduct void content analyses [1,4,12–17].

In order to avoid the time-consuming task of evaluating the relative void content of micrography images manually [4,17], a commonly employed technique is automatic image segmentation by pixel intensity thresholding [1,13,18,19]. This technique relies on the different pixel intensities of the composite and the void and is usually user calibrated. However, despite the good results that can be achieved, this methodology can be very sensitive to the illumination conditions during the acquisition of the images as well as the type of material being analysed. Therefore, the calibration of the algorithm parameters (including the threshold value), is a necessary step before the analysis of a given set of images.

Another problem that can undermine thresholding approaches is the appearance of large voids in laminates. As void size increases, light coming from the microscope illumination can be reflected from the inside of the void cavities, which in turn originates lighter areas inside the dark ones. This translates to high pixel intensities, which should be classified as voids, that are mistakenly classified as matrix, due to its naturally higher pixel intensity. In turn, this renders the common thresholding approaches ineffective, as these techniques are not able to detect the void areas entirely (Figure 1).

The adoption of machine-learning algorithms to do automatic detection of voids has been reported in the literature for several void assessment techniques, such as X-ray micro-CT [24–26], thermography [27] and ultrasound testing [28]. Luo *et al.* used a deep learning framework based on DeepLabV3+, which achieved good void segmentation results in optical microscopy images [23]. However, their results show that the segmentation accuracy of a thresholding algorithm is very close to the one obtained by the deep learning algorithm. In turn, it is plausible to infer that the images present in their dataset might not have the complexity that is added when large pixel intensity scattering exists due to the presence of large voids and reflections. This increased complexity could produce larger differences between thresholding and machine-learning results than the ones Luo reported.

In this work, a machine vision algorithm based on machine-learning was developed, to analyse microscopy images for void detection, in order to overcome the shortcomings of common thresholding approaches, which reliability is greatly affected by the pixel intensity variability.

For that matter, a machine learning approach based on convolutional neural networks was used to analyse microscopy images and obtain the corresponding void contents.

This article is organized as follows: Section one presents a brief introduction to convolutional neural networks is given. Section two describes the methodology used to create the machine learning framework. Section three shows the results obtained with the proposed model. Finally, section four presents the conclusions taken from this study.

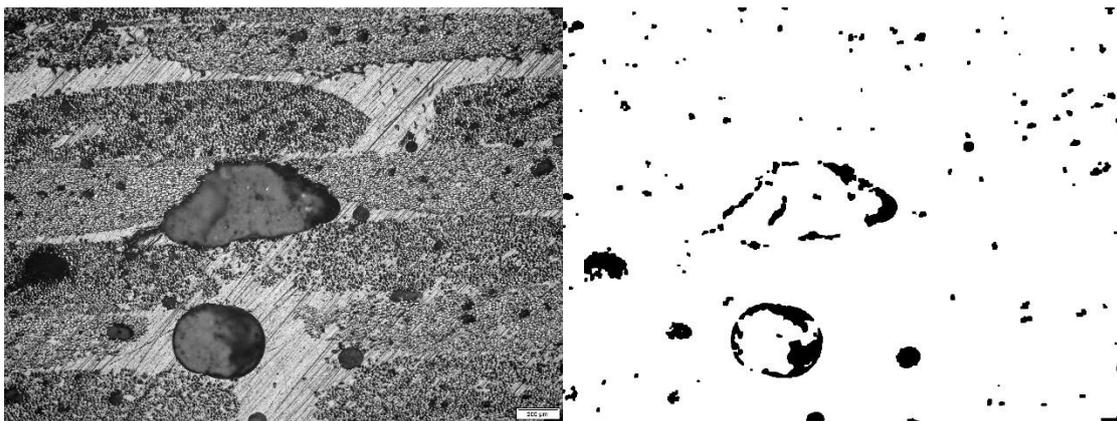


Figure 1 – Microscopy image with light reflecting voids (on the left) and poor performance of thresholding based segmentation method (on the right).

Convolutional neural networks

In an artificial neural network, a set of inputs is mapped to an output, by means of a mathematical function [29]. If the inputs are mapped directly to an output, it is denominated

as a single-layer neural network. On the contrary, if the inputs are mapped to an output through a succession of subsequent (hidden) layers, the neural network is denominated as of the multi-layer type [29,30]. The universal approximation theorem states that a neural network with at least one hidden layer can be used to approximate any function well, provided that the network has enough hidden units [29,30].

Similarly to traditional artificial neural networks, the architecture of convolutional neural networks is built upon layers, which are connected in a logical sequence. In analogy to neural networks and the universal approximation theorem, a convolutional neural network can be used to approximate any continuous function to a desired non-zero amount of error, provided that the depth of the convolutional neural network is large enough [31].

However, unlike traditional artificial neural networks, convolutional neural networks can possess different types of layers: fully connected layers, convolutional layers and pooling layers.

Fully connected layers are a type of layer in which every neuron is connected each neuron of the previous layer by a distinct set of weights, which are the layer trainable parameters:

$$z^l = \sum_{j=1}^n w_{ij}^{l-1} x_i^{l-1} + b^{l-1} \quad (1)$$

where z is the vector containing the input node values to layer l , w_{ij} is the connection weight between neurons, x^{l-1} is the activated neuron value of the previous layer, and b is the bias vector (omitted in Figure 2 for conciseness).

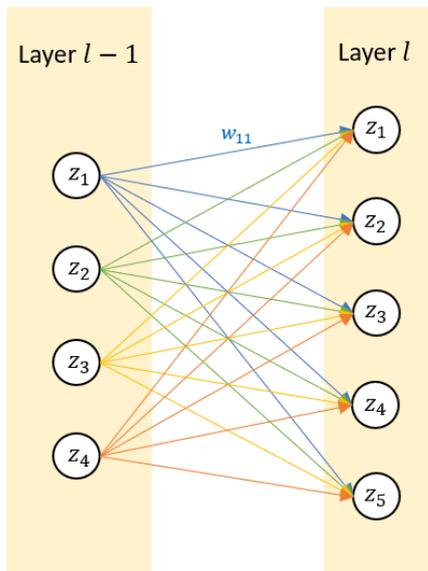


Figure 2 – Example of a fully connected layer

Fully connected layers are the staple of traditional artificial neural networks, which are only comprised by a succession of this type of layers. However, in convolutional neural networks, this type of layers can be commonly found in the ending layers of the network [32–34].

Convolutional layers implement the convolution operation, which for two-dimensional tensors can be written as:

$$S(i, j) = (I * K)(i, j) = \sum_m \sum_n I(m, n)K(i - m, j - n) \quad (2)$$

where I is the two-dimensional tensor being convolved, K is a two-dimensional kernel and S is the resulting tensor.

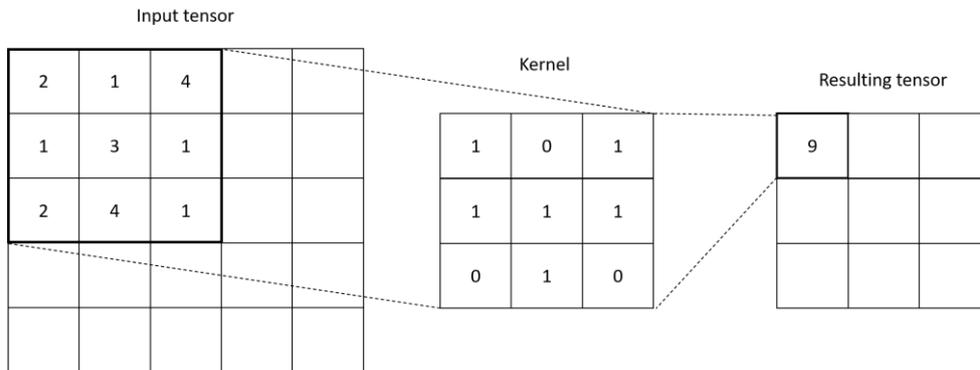


Figure 3 – Example of a convolution operation.

Unlike fully connected layers, which can only accept an input of a predetermined size, convolutional layers do not have this restriction, as the learnable parameters are embedded in the kernel, which size is independent of the input tensor (feature map).

One characteristic outcome of convolutions (observable in Figure 3) is the reduction of the size of the feature map, whose extent depends on the size of the kernel. In case this behaviour is not desirable, and if one intends to maintain the size of the feature map, padding can be added to the feature map before the convolution operation. This is done by adding an outer layer of values, usually in the form of zeros, an operation commonly designated as zero-padding (Figure 4). Moreover, if enough padding is added to the input feature map, one can obtain a bigger output than the original input, leading to a transposed convolution (also known as up-convolutions or deconvolutions). This increase of the size of the feature map, commonly designated as up-sampling, can be useful in certain network architectures, such as autoencoders [29].

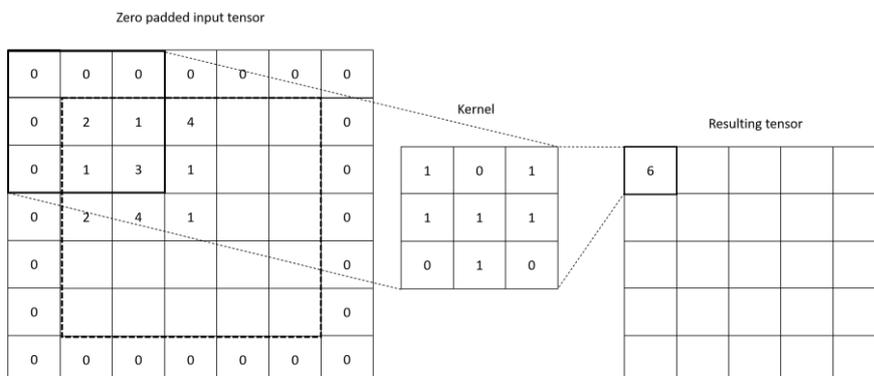


Figure 4 – Example of a padded convolution.

At the end of each convolutional layer or fully connected layer, a non-linear activation function can be commonly found. These functions have been found to allow the network to learn more complex features in data, compared to linear activation functions [29]. An activation non-linearity commonly used in convolution neural networks is the Rectified Linear Unit (ReLU), which is a piece-wise linear function that will output the received input, in case it is positive, otherwise the output is zero. This activation function is particularly relevant for deep learning (neural network architectures with several layers), as it better preserves the gradient information across several layers deep, compared to the logistic, or commonly designated sigmoid activation function, which can suffer from saturation for large activation values (Figure 5) [30]. The ReLU activation can be written as:

$$f(x) = \begin{cases} x & \forall x > 0 \\ 0 & \forall x \leq 0 \end{cases} \quad (3)$$

The logistic, or sigmoid, activation function (Figure 6) can be written as:

$$f(x) = \frac{1}{1 + e^{-x}} \quad (4)$$

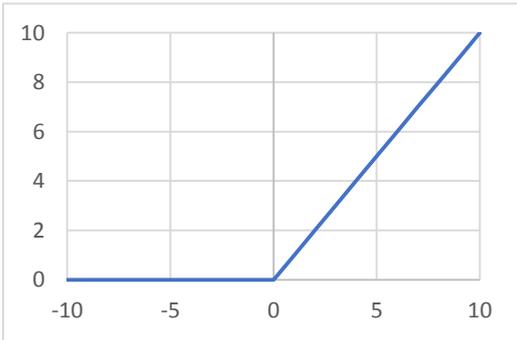


Figure 5 - ReLU activation function.

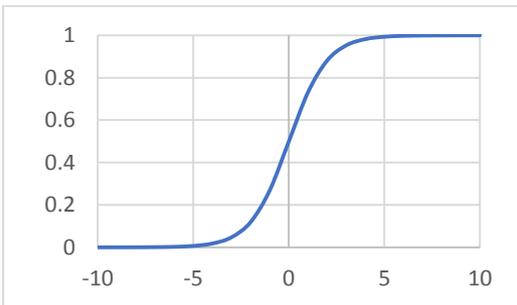


Figure 6 - Sigmoid activation function.

The last staple in a convolutional neural network structure is the pooling layer. Pooling layers can either extract the maximum, minimum or average value inside a sliding window with a predetermined size and stride (Figure 7). No trainable parameters exist in this type of layers, as

their objective mostly relies in reducing the size of the feature map, an operation commonly designated as down-sampling. However, unlike the convolutional layers, pooling layers are not constrained to an input of a predetermined size, as the pooling window slides throughout the entire tensor with a defined stride, independently of the tensor size.

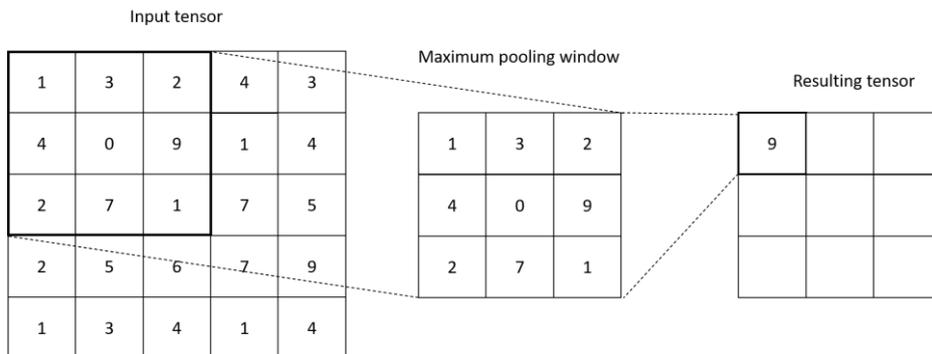


Figure 7 – Example of a maximum pooling operation.

In the field of machine vision, images are interpreted as a collection of pixel intensity values, which is represented as a tensor with varying depth, depending on the encoding of the image colours. A greyscale image can be therefore represented as a two-dimensional tensor, which dimensions match the resolution of the image, whereas an RGB encoded image can be represented as a three-dimensional tensor, for example, with a depth equal to three, representing the red, green and blue channels.

When the neural network processes an image, the number of inputs in the network will match the resolution of the image, multiplied by the number of channels it possesses. The number of inputs can be therefore substantially large. As convolutional neural networks consist mainly of convolutional and pooling layers, in which trainable parameters are not dependent on the size of the input tensor, this type of networks allow the design of deeper network architectures allied with a faster training, without recurring into memory and computational overloads, compared to traditional artificial neural networks [30]. Consequently, the field of machine learning applied to computational vision has seen important performance gains with the intensification of research around convolutional neural networks (CNNs) [29,35]. As a result, different research branches were created to solve problems such as image classification (assignment of a single class per image) [34], or problems which require a pixel-level type of inference, such as semantic segmentation (segmentation based on the classes existing) [36], as well as of instance segmentation (segmentation based on instances of each class present in the image) [37]. In the scope of this work, the automatic segmentation of microscopy images, for subsequent determination of relative void content, is a semantic segmentation problem.

U-net architecture

The U-net is a semantic segmentation model architecture proposed by Ronneberger *et al.* [38], which is built upon the previous fully convolutional network model for semantic segmentation [36], making the U-net architecture a fully convolutional network, itself (Figure 8). As mentioned in the former section, from the image segmentation point of view, as fully convolutional networks do not contain fully connected layers, they present the added advantage of being able to process images with a variable size, while reducing the computational overhead, due to connection sparsity and parameter sharing [30].

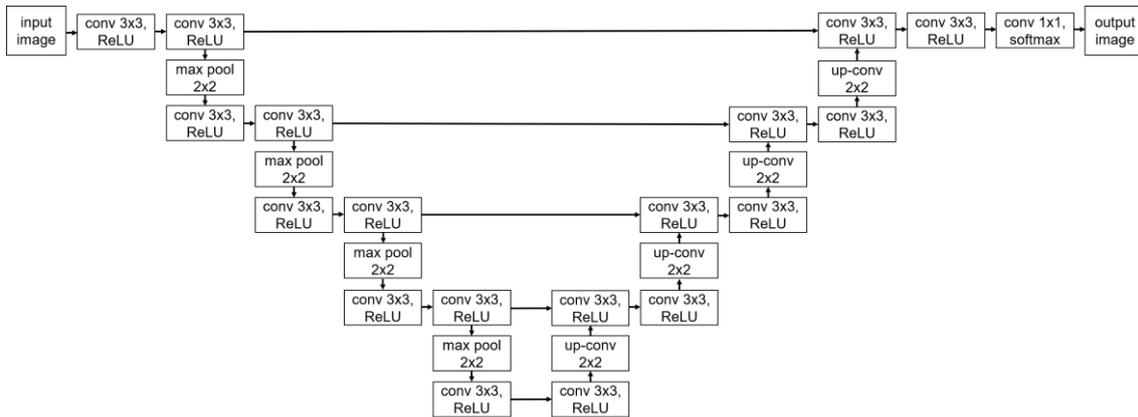


Figure 8 - Original U-Net architecture model [38].

The U-net architecture possesses a contracting path (encoder), where the initial feature-map (input image) is reduced in size, while the number of feature channels increases. This is done through a series of convolution layers followed by pooling layers, as in a regular CNN. The relevant features of the image (context) are intended to be captured in this portion of the network. The second portion of the network is an expansive path (decoder), where the size of feature-map is up-sampled and the number of feature channels is reduced. During the up-sampling procedure, the feature-map is concatenated with its corresponding pair of the contracting path. This strategy ensures that the network captures with more refinement the locations of the relevant features in the image. This forces the network to have an approximately symmetric architecture. The up-sampling of the feature map is achieved through transposed convolutions (or up-convolutions), followed by regular convolution layers. This portion of the network is thus intended to capture the location of the relevant features, to make a refined reconstruction of the image. Due to this network architecture strategy, the model can be trained with smaller datasets and higher learning rates than other convolutional neural network models [38]. Because of this, several follow-through model architectures have been proposed with the objective of enhancing the accuracy of the segmentation results mainly for specific biomedical imaging problems [39–42], as well as segmentation of aerial and satellite imagery [43,44].

Methodology

Network architecture

The implemented network architecture follows much of the original U-Net architecture proposed by Ronneberger *et al.* [38], with a contracting path composed by successive 3x3 convolution layers, each followed by a rectified linear unit (ReLU) (Figure 9). The down-sampling is achieved by 2x2 pooling layers after each pair of convolution layers. A feature that

can increase both the accuracy and reliability of the results, is adapting the pooling layers to the natural appearance of voids in micrographs, which translates to generally lower pixel intensity values. Hence, voids, which are the relevant features of the image under analysis, can be better captured during the down-sampling operation. This means that one can substitute the original maximum pooling layers by minimum pooling layers or, instead, invert the pixel intensities of the image (voids become lighter and matrix darker) and maintain the maximum pooling layers. In this work, the latter option was chosen, due to the lack of a minimum pooling layer implementation in the framework used. Batch normalization layers are added before each ReLU activation layer, as batch normalization is reported to increase both the stability and speed of the learning process [45].

As in the original U-net, the expansive path is constituted by four similar blocks, containing a set of different layers. These blocks start with 2x2 transposed convolution layers, which are responsible for up-sampling the feature map. In order to enhance the capacity of the network to capture more precisely the location of the relevant features (voids), after the transposed convolution layer, the resulting feature map is concatenated with its corresponding pair of the contracting path. The resulting concatenation is fed to a pair of 3x3 convolution layers, each followed by batch normalization and a ReLU activation.

The eight components of the remaining feature vector are mapped to the desired number of classes, adding a 1x1 convolution layer to the end of the network. As the intended pixel labelling is binary (void or matrix), the 1x1 convolution layer maps to a single tensor, where a final sigmoid activation translates the values obtained into values in the range]0, 1[. These values can be seen as probabilities of the pixel belonging to the void class. After applying a probability threshold to the obtained values, pixels with a value of 1 (one) are considered voids, whereas pixels with a value of 0 (zero) are considered matrix.

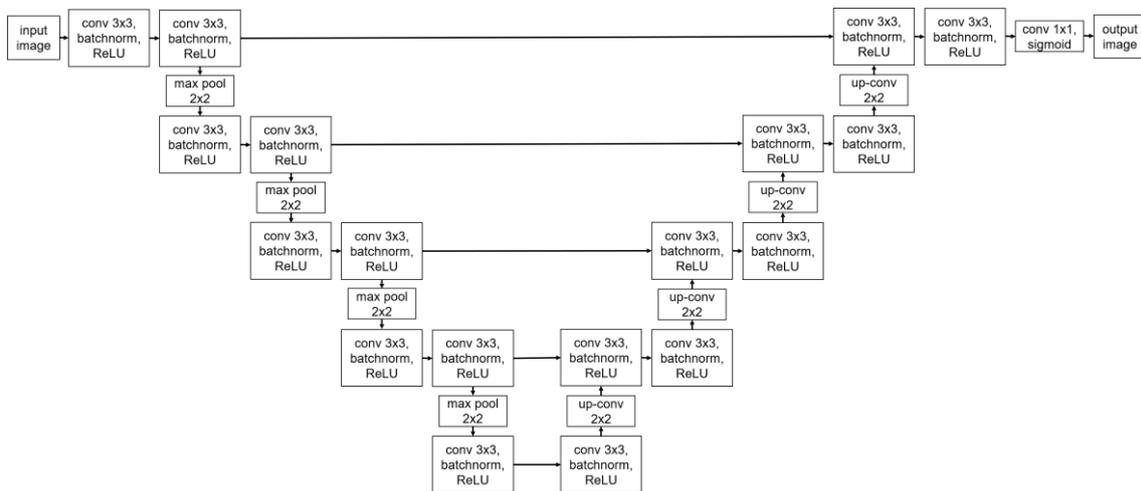


Figure 9 - Proposed modified U-Net architecture.

Dataset

The dataset used for this study is comprised by microscopy images captured at INEGI, under the polishing and image capturing conditions described in Table 1. The samples come from three different types of composite laminates (Figure 13): glass fibre and epoxy laminate processed by vacuum infusion (Type A laminates); carbon fibre and epoxy laminate processed by resin transfer moulding (Type B laminates); carbon fibre and epoxy laminate processed by vacuum infusion (Type C laminates).

Laminate type	Sandpaper grit (last polishing)	Optical microscope
A	2000	Olympus PMG3 w/ CCD camera
B	1000	Olympus PMG3 w/ CCD camera
C	1000	Olympus PMG3 w/ CCD camera

Table 1 - Polishing and image acquisition conditions

For each image in the dataset, a corresponding ground-truth mask was generated. The ground truth masks consist of binary 8-bit gray-scale images, where the pixels representing voids have a value of 255, whereas pixels representing matrix or fibers have a value of 0. Therefore, these ground-truth masks allow to determine unequivocally which pixels are voids (the object of interest), and which pixels are either matrix or fibres (no distinction is necessary in our study). The ground-truth masks were generated by running a thresholding-based in-house software, while further segmentation corrections were made manually, using GIMP open-source image processing software. From a total of sixty images in the dataset, thirty were used for training and the other half for validation. The selection of the images was random.

In order to obtain a characterization of the dataset, for each type of laminate, void size and frequency were measured automatically, using the contours of the voids present in the ground truth masks. Contouring is a well-established method for the representation of the geometry of voids in binarized images, allowing sub-pixel accuracy [20,46]. In turn, the accuracy of the binarization process dictates the accuracy of the extracted void related data. Python methods available in the OpenCV library were used for automatic contour extraction and respective area calculation. The computation of the frequency measures was also carried by a Python script written for this purpose, while further statistical analysis was done in R statistical language.

In turn the frequency measures were grouped into five bins with equal width. Mean pixel intensity and standard deviation were also determined, encoding the original image in greyscale format. The different backgrounds of the microscopy samples were expected to foment an increase in the reliability of the network inference results, due to a higher capacity for generalization.

As it can be observed in Tables 2, 3 and 4, the size and frequency of voids varies significantly depending on the manufacturing process. Laminates manufactured by vacuum infusion (type A and type C) seem to possess a higher number of voids, as well as a higher variance in void sizes, compared to the laminates manufactured by resin transfer moulding (type B). This can be due to the lack of a mould packing procedure in vacuum infusion, whereas in resin transfer moulding, it is possible to do so by increasing resin injection pressure after mould filling, which in turn compacts the existent voids, therefore minimizing void content in the part [3]. However, conducting such type of quantitative analysis is out of the scope of this study.

Nevertheless, the void sizes obtained for all laminate types in our dataset are in agreement with the range of void sizes found in the literature [21,23]. Moreover, plotting the frequency measures of void sizes into an histogram, for each laminate type, it can be observed that the distribution of void sizes follows a Weibull distribution (Figures 10, 11, 12), which is also consistent with the reported literature [47]. For the sake of easiness of visualization, Figures 10 and 12 only plot the frequency measures on the first bin of laminate types A and C, respectively, as the remaining bins have only residual frequency values.

Void area bin [μm^2]	Number of voids	Frequency	Mean area	Area standard deviation	Coefficient of variation	Mean pixel intensity	Pixel intensity standard deviation
22.72 - 104962.73	5932	99.48%	1303.43	5255.83	403%	62.34	35.93
104962.73 - 209902.75	18	0.30%	138140.45	26567.77	19%	75.55	46.38
209902.75 - 314842.76	9	0.15%	258181.69	22393.60	9%	98.34	52.08
314842.76 - 419782.77	3	0.05%	359158.30	26297.97	7%	128.60	43.88
419782.77 - 524722.79	1	0.02%	440786.31	0.00	0%	57.19	21.30

Table 2 - Properties of type A laminates.

Void area bin [μm^2]	Number of voids	Frequency	Mean area	Area standard deviation	Coefficient of variation	Mean pixel intensity	Pixel intensity standard deviation
74.62 - 24913.59	42	73.68%	10430.46	8569.67	82%	45.86	6.56
24913.59 - 49752.57	11	19.30%	32487.34	4695.19	14%	44.40	5.93
49752.56 - 74591.54	1	1.75%	62479.84	0.00	0%	40.88	1.98
74591.54 - 99430.52	0	0.00%	--	--	--	--	--
99430.52 - 124269.50	3	5.26%	113349.23	7794.01	7%	47.57	10.07

Table 3 - Properties of type B laminates.

Void area bin [μm^2]	Number of voids	Frequency	Mean area	Area standard deviation	Coefficient of variation	Mean pixel intensity	Pixel intensity standard deviation
29.21 - 93240.56	774	97.60%	3862.02	10122.17	262%	85.94	26.74
93240.56 - 186451.90	5	0.63%	124782.00	28115.30	23%	104.63	12.34
186451.90 - 279663.25	5	0.63%	262437.56	8067.09	3%	93.58	10.41
279663.25 - 372874.60	5	0.63%	320011.96	25303.36	8%	83.79	10.31
372874.60 - 466085.94	4	0.50%	449319.15	14296.00	3%	81.16	7.92

Table 4 - Properties of type C laminates.

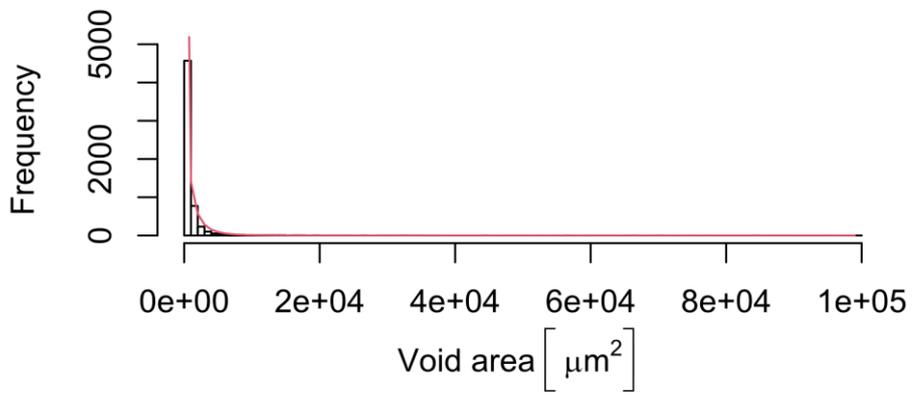


Figure 10 - Void size frequencies from the first bin in type A laminates and fitted weibull distribution (in red)

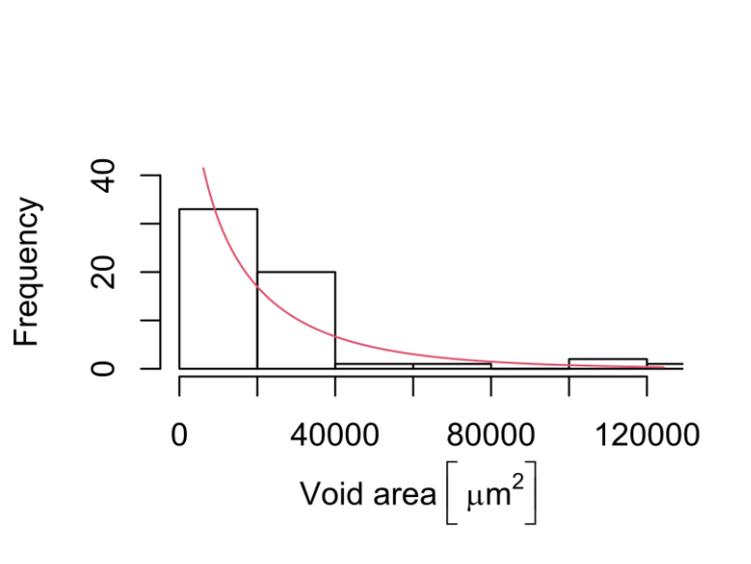


Figure 11 - Void size frequencies from all bins in type B laminates and fitted weibull distribution (in red)

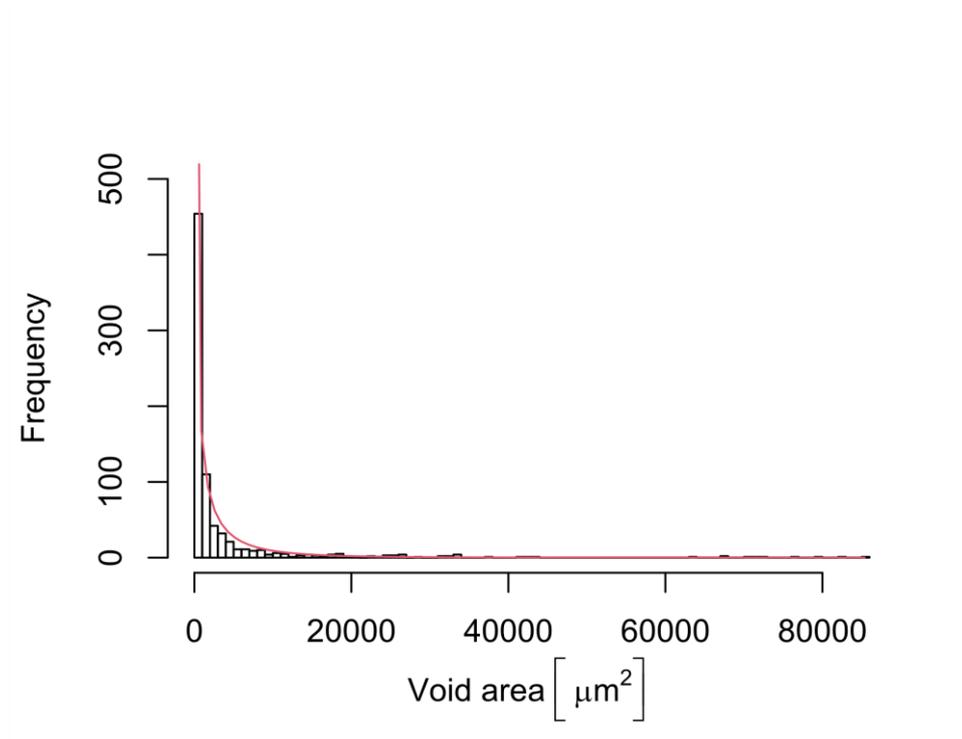


Figure 12 - Void size frequencies from the first bin in type C laminates and fitted weibull distribution (in red)

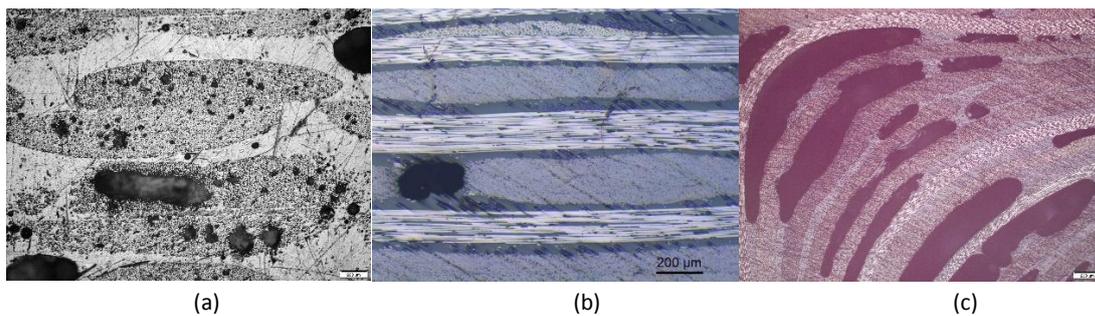


Figure 13 – Microscopy samples of the used dataset: glass fibre laminate processed by vacuum infusion – Type A laminate (a); carbon fibre laminate processed by resin transfer moulding – Type B laminate (b); carbon fibre laminate processed by vacuum infusion – Type C laminate (c)

Training

To conduct the U-net training, each of the high-resolution training images was partitioned into a set of twenty grayscale 256x256 pixel smaller images. The benefits of this strategy were twofold: Firstly, this strategy allowed to increase the number of filters of the network without incurring in GPU memory overloads. Moreover, this strategy allowed each training batch to contain images of all types of laminates. As the network training is based on gradient optimization with an update of the network weights on a per-batch basis, this strategy allows a better estimation of the gradient, and therefore, a more efficient training.

The network was trained using the Adam optimization algorithm [48] with an initial learning rate of 0.001, binary cross-entropy loss function and a batch size of 40 images, for a total of 400 epochs. The number of batches per epoch was estimated to assure that theoretically all 256x256 dataset images would be processed during a training epoch. The model was implemented in Keras, using Tensorflow and an Nvidia Quadro RTX6000 with 24GB of memory.

Results

Four different metrics were calculated for both the training dataset, as well as the validation dataset, using a probability threshold of 0.35: accuracy, precision, recall and intersection on union (IoU), to evaluate the performance of the proposed deep learning network. These metrics can be calculated from the confusion matrix, which stores the frequency measures for each true positive (TP), false positive (FP), false negative (FN) and true negative (TN) pixel classifications, as demonstrated in Figure 14.

		Actual	
		Void	Matrix
Predicted	Void	True Positive (TP)	False Positive (FP)
	Matrix	False Negative (FN)	True Negative (TN)

Figure 14 - Confusion matrix.

Subsequently, accuracy, precision, recall and IoU can be calculated as:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (5)$$

$$Precision = \frac{TP}{TP + FP} \quad (6)$$

$$Recall = \frac{TP}{TP + FN} \quad (7)$$

$$IoU = \frac{TP}{TP + FP + FN} \quad (8)$$

The use of different metrics allows one to obtain answers to different questions. Accuracy reflects the number of correctly classified pixels by the total number of pixels analysed. Precision allows one to assess out of all pixels classified by the network as voids, how many are voids. Recall allows one to assess out of all pixels which are voids, how many were classified by the network as voids. Lastly, IoU evaluates out of the group composed by all the pixels classified as voids, as well as the pixels which are actually voids (union), how many pixels are actually correctly classified as voids by the network (intersection). One relevant matter for assessing metrics, is the fact that accuracy can be sensitive to unbalanced datasets (datasets in which one class is more representative than the others), possibly giving biased results. In the case of an unbalanced dataset, preference should be given to metrics such as IoU, as these are less sensitive to unbalances between the dataset classes.

A physical interpretation of the segmentation results achieved by the network, was produced by frequency measures, which were computed for the different void sizes present in the validation dataset. In turn, these measures were compared to the ones obtained by the segmentation results of the network. Using the same confusion matrix analogy for void instance statistics, intersection on union was computed for each computed bin of void sizes. Tables 6, 7 and 8 contain the obtained results for each type of laminate under analysis.

Metric	Training Dataset	Validation Dataset
Accuracy (binary)	0.9970	0.9936
Precision	0.9491	0.9299
Recall	0.9907	0.9114
Intersection on Union (IoU)	0.9650	0.9241

Table 5 - Network performance evaluation.

Void area bin [μm^2]	Void n ^o	Voids detected	IoU
22.72 - 104962.73	3188	2349	73.68%
104962.73 - 209902.75	12	11	91.67%
209902.75 - 314842.76	4	4	100%
314842.76 - 419782.77	0	0	100%
419782.77 - 524722.79	2	2	100%

Table 6 - Segmentation results for type A laminate samples.

Void area bin [μm^2]	Void n ^o	Voids detected	IoU
381.92 - 20089.10	23	30	76.67%
20089.10 - 39796.28	8	7	87.5%
39796.28 - 59503.46	1	0	0%
59503.46 - 79210.65	0	1	0%
79210.65 - 98917.83	1	1	100%

Table 7 - Segmentation results for type B laminate samples.

Void area bin [μm^2]	Void n ^o	Voids detected	IoU
29.21 - 93240.56	83	100	83%
93240.56 - 186451.90	2	2	100%
186451.90 - 279663.25	2	3	66.66%
279663.25 - 372874.60	3	2	66.66%
372874.60 - 466085.94	2	2	100%

Table 8 - Segmentation results for type C laminate samples.

	Mean void content error	Error standard deviation
Type A laminate dataset	0.66%	0.29%
Type B laminate dataset	0.34%	0.27%
Type C laminate dataset	0.50%	0.20%

Table 9 - Error analysis of global void content detected in dataset images.

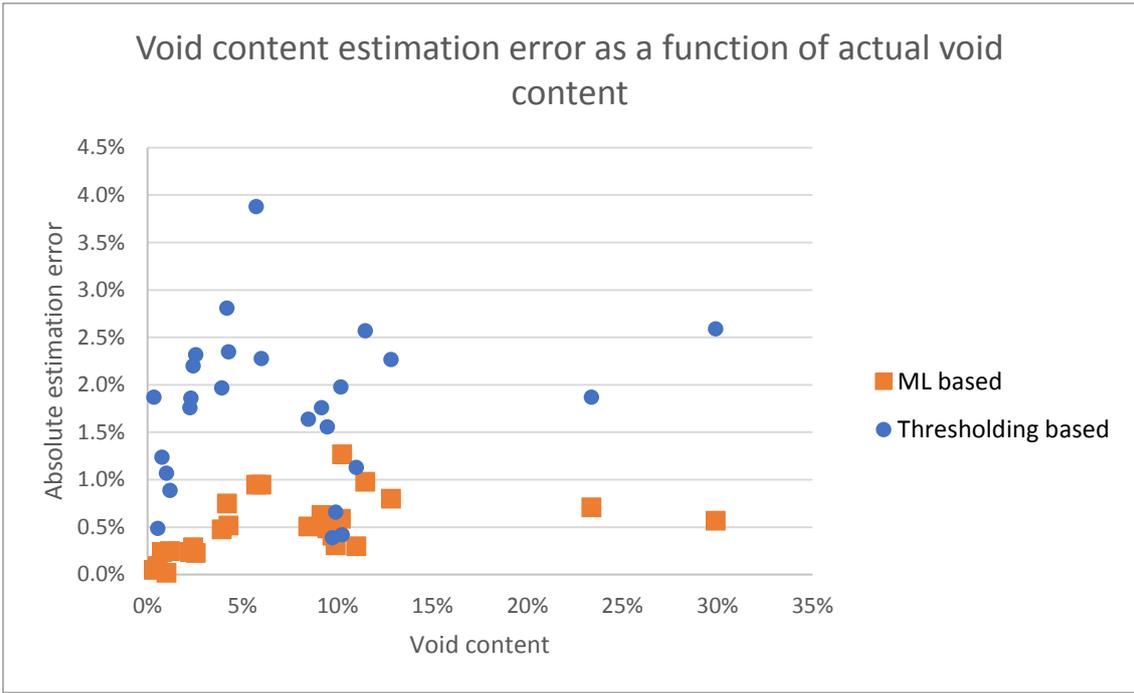


Figure 15 - Comparison of void content absolute estimation error between the proposed machine learning algorithm and a thresholding based algorithm

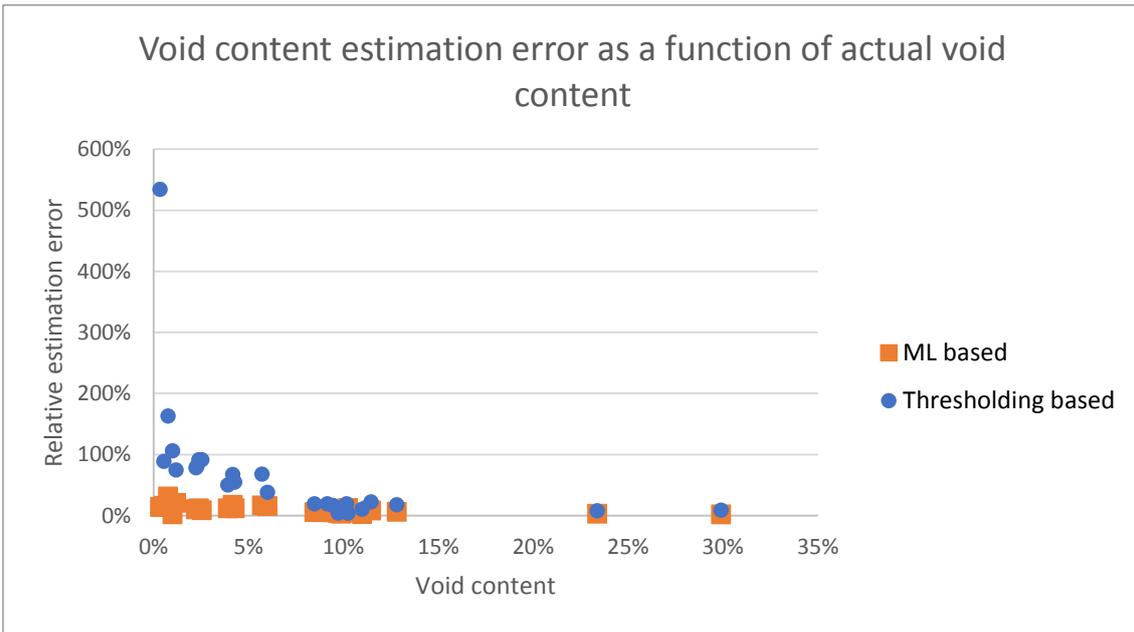


Figure 16 - Comparison of void content relative estimation error between the proposed machine learning algorithm and a thresholding based algorithm

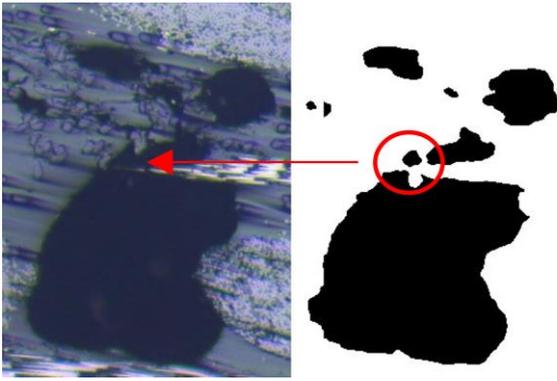


Figure 17 - Decreased capacity in void edge delineation, for voids containing fuzzy edges (laminate type B)

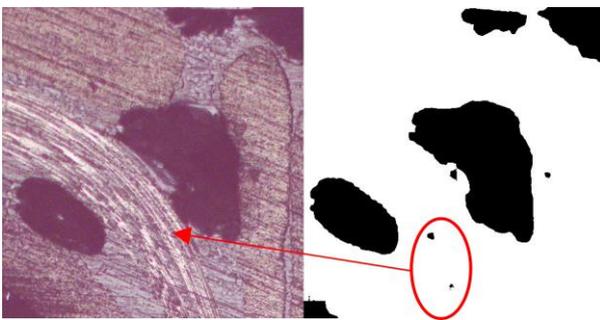


Figure 18 – Overdetection of small voids (laminate type C)

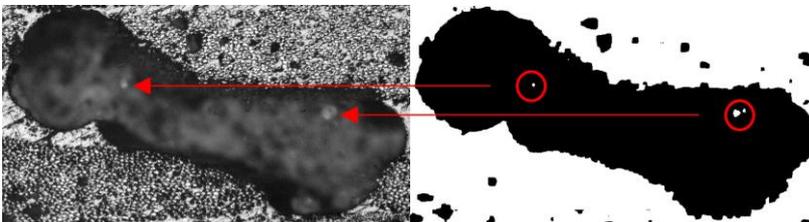


Figure 19 - Underdetection of void area, for big voids (laminate type A)

From the results presented in Tables 6, 7 and 8, it can be seen that the network correctly identified the majority of voids present in the micrography images, whereas for the type B laminate dataset, the network had its worst performance. This lack of performance may be due to the slightly decreased capacity of the neural network in delineating edges of the voids, when these have fuzzy edges (as exemplified in Figure 17). In turn, the overall detected area of the void is smaller than in reality, leading to a biased detection of voids in the presented statistics.

Additionally, the network performed worse in detecting smaller voids (first bin of void sizes for all laminate types), probably due to noise present in the images. This noise is composed by abrupt color changes, which may be due to small scratches in the matrix, or darker matrix features. In turn, as the convolutional neural network may have not learned entirely which set

of features is characteristic to smaller voids, it may be producing a slight difference in the predicted quantity of voids, as obtained in the current analysis where laminates type B and C have an overprediction of voids, whereas laminte A suffers from an undeprediction. An example can be seen in Figure 18.

Finally, in certain large bubbles, the network still fails to capture the entire area of the bubble. This happens probably due to the high scattering of pixel intensities that can be found inside certain bubbles, which intensity values can reach close to the maximum, 255. However, this effect is localized in very small regions of the voids, compared to their global area, which in turn do not affect greatly their measured areas. This effect can be seen in Figure 19.

The inference metrics appear to be better than the reported void detection results. A reasonable explanation is that inference metrics are measured on a by-pixel basis, which has no relation to void instance statistics. This means that although the pixel classification done by the network is good enough globally, the segmentation may not be entirely precise, due to the reasons delineated above. Nevertheless, this relative segmentation error is low enough as the mean absolute void content error is below 1%, independently of the laminate type (Table 9). This error is the mean of the estimated errors for each validation dataset (Equation 9). As expected, no error was obtained between void contents derived from the real images and the inference results of the training dataset.

The global void content estimation error of the proposed algorithm was compared to the results obtained with a thresholding based algorithm developed prior to this study by the same authors. The thresholding parameters were optimized on a laminate type basis. Absolute and relative errors were derived for each algorithm type:

$$Error_{absolute} = |void_{real} - void_{algorithm}| \quad (9)$$

$$Error_{relative} = \frac{|void_{real} - void_{algorithm}|}{void_{real}} \quad (10)$$

From Figure 15, it can be seen that the absolute void content estimation errors associated with the proposed machine learning algorithm were lower than the ones obtained using the manually thresholding-based algorithm. No clear dependence between real void content and estimation error was detected; however, it can be observed that with the void content increased the standard deviation of the estimation errors also increases. Nevertheless, a cause-effect study is out of the scope of this study. Regarding relative estimation errors, plotted in Figure 16, it can be seen that the errors were proportionally higher for lower void contents. This was expected, independently of the segmentation algorithm, as the segmentation error does not reach an absolute null value. Therefore, as the void content approaches zero, since the estimation error does not drop accordingly, the relative error tends to rise. Nevertheless, the errors related to the proposed machine learning algorithm were lower than the errors related to thresholding alternative.

It is important to emphasize that the results shown in this study should only be interpreted in the context of the dataset used. Convolutional neural networks, such as the U-Net, are

designed to make inferences based on the interpolation of several features present on the dataset provided. This means that the reliability of inference results outside the training and testing datasets may be greatly affected since these algorithms are not designed to make extrapolations outside the training data. Therefore, the generalization capability of such algorithms is linked to how general the dataset is itself.

Figure 20 depicts the segmentation results for different types of microscopy images.

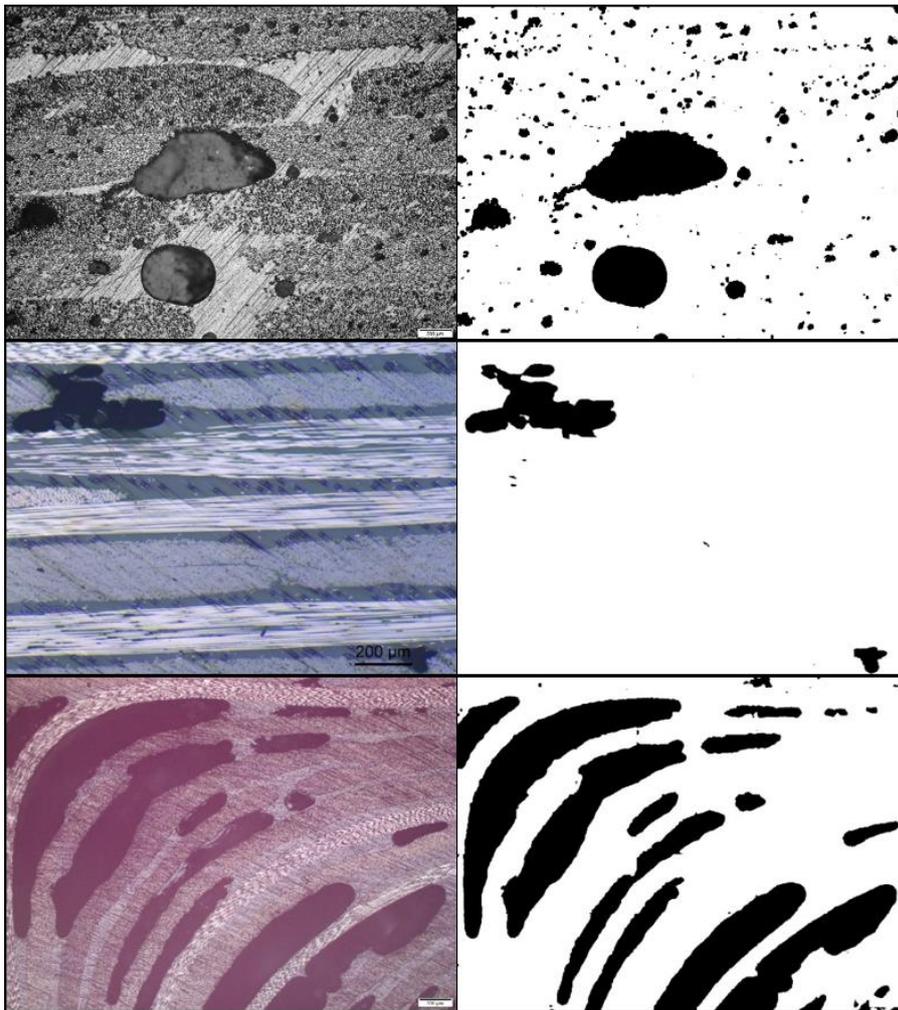


Figure 20 - Segmentation results (on the right) for different microscopy images (on the left).

Conclusion

It was successfully demonstrated that using machine learning techniques applied to computational vision, common micrography samples can be automatically segmented, in order to calculate their relative void content.

The u-net architecture is a rather convenient machine learning approach for semantic segmentation, as it needed very few annotated images and training time. Using a microscopy image dataset built for this study, the segmentation results suggest that the network performs worse in detecting smaller voids, while the appearance of fuzzy void edges may also affect the accuracy of the segmentation. At last, pixel intensity variability can also be a factor for incomplete segmentation of the void area. Nevertheless, the achieved inference results are very promising as the obtained average void content error was below 1%, regardless of the laminate type. These results have surpassed a thresholding based algorithm manually calibrated for each laminate type dataset, thus proving the applicability of this methodology.

Acknowledgements

The authors would like to acknowledge the support from the Associated Laboratory for Energy, Transports and Aeronautics (LAETA) under the Research Grant UIDB/50022/2020 and CompRTow (Composites from Recycled TOWs) NORTE-01-0247-FEDER-038349.

Data Availability

The raw/processed data required to reproduce these findings cannot be shared at this time as the data also forms part of an ongoing study.

References

- [1] Bodaghi M, Cristóvão C, Gomes R, Correia NC. Experimental characterization of voids in high fibre volume fraction composites processed by high injection pressure RTM. *Compos Part A Appl Sci Manuf* 2016;82:88–99. <https://doi.org/10.1016/j.compositesa.2015.11.042>.
- [2] van Oosterom S, Allen T, Battley M, Bickerton S. An objective comparison of common vacuum assisted resin infusion processes. *Compos Part A Appl Sci Manuf* 2019;125:105528. <https://doi.org/10.1016/j.compositesa.2019.105528>.
- [3] Bodaghi M, Costa R, Gomes R, Silva J, Correia N, Silva F. Experimental comparative study of the variants of high-temperature vacuum-assisted resin transfer moulding. *Compos Part A Appl Sci Manuf* 2019:105708. <https://doi.org/10.1016/j.compositesa.2019.105708>.

- [4] Olivier P, Cottu JP, Ferret B. Effects of cure cycle pressure and voids on some mechanical properties of carbon/epoxy laminates. *Composites* 1995;26:509–15. [https://doi.org/10.1016/0010-4361\(95\)96808-J](https://doi.org/10.1016/0010-4361(95)96808-J).
- [5] Harper BD, Staab GH, Chen RS. A Note on the Effects of Voids Upon the Hygral and Mechanical Properties of AS4/3502 Graphite/Epoxy. *J Compos Mater* 1987;21:280–9. <https://doi.org/10.1177/002199838702100306>.
- [6] Liu L, Zhang BM, Wang DF, Wu ZJ. Effects of cure cycles on void content and mechanical properties of composite laminates. *Compos Struct* 2006;73:303–9. <https://doi.org/10.1016/j.compstruct.2005.02.001>.
- [7] de Almeida SFM, Neto Z dos SN. Effect of void content on the strength of composite laminates. *Compos Struct* 1994;28:139–48. [https://doi.org/10.1016/0263-8223\(94\)90044-2](https://doi.org/10.1016/0263-8223(94)90044-2).
- [8] Hapke J, Gehrig F, Huber N, Schulte K, Lilleodden ET. Compressive failure of UD-CFRP containing void defects: In situ SEM microanalysis. *Compos Sci Technol* 2011;71:1242–9. <https://doi.org/10.1016/j.compscitech.2011.04.009>.
- [9] Maragoni L, Carraro PA, Peron M, Quaresimin M. Fatigue behaviour of glass/epoxy laminates in the presence of voids. *Int J Fatigue* 2017;95:18–28. <https://doi.org/10.1016/j.ijfatigue.2016.10.004>.
- [10] Sisodia S, Gamstedt EK, Edgren F, Varna J. Effects of voids on quasi-static and tension fatigue behaviour of carbon-fibre composite laminates. *J Compos Mater* 2015;49:2137–48. <https://doi.org/10.1177/0021998314541993>.
- [11] Talreja R. Studies on the failure analysis of composite materials with manufacturing defects. *Mech Compos Mater* 2013;49:35–44. <https://doi.org/10.1007/s11029-013-9318-6>.
- [12] Jeong H. Effects of voids on the mechanical strength and ultrasonic attenuation of laminated composites. *J Compos Mater* 1997;31:276–92. <https://doi.org/10.1177/002199839703100303>.
- [13] Guerdal Z, Tomasino AP, Biggers SB. Effects of processing induced defects on laminate response. Interlaminar tensile strength. *SAMPE J* 1991;27:39–49.
- [14] Hamidi YK, Aktas L, Altan MC. Three-dimensional features of void morphology in resin transfer molded composites. *Compos Sci Technol* 2005;65:1306–20. <https://doi.org/10.1016/j.compscitech.2005.01.001>.
- [15] Naganuma T, Naito K, Kyono J, Kagawa Y. Influence of prepreg conditions on the void occurrence and tensile properties of woven glass fiber-reinforced polyimide composites. *Compos Sci Technol* 2009;69:2428–33. <https://doi.org/10.1016/j.compscitech.2009.06.012>.
- [16] Grunenfelder LK, Nutt SR. Void formation in composite prepregs - Effect of dissolved moisture. *Compos Sci Technol* 2010;70:2304–9. <https://doi.org/10.1016/j.compscitech.2010.09.009>.
- [17] Purslow D. On the optical assessment of the void content in composite materials. *Composites* 1984;15:207–10. [https://doi.org/10.1016/0010-4361\(84\)90276-3](https://doi.org/10.1016/0010-4361(84)90276-3).
- [18] Little JE, Yuan X, Jones MI. Characterisation of voids in fibre reinforced composite materials. *NDT E Int* 2012;46:122–7. <https://doi.org/10.1016/j.ndteint.2011.11.011>.

- [19] Yang P, El-Hajjar R. Porosity Defect Morphology Effects in Carbon Fiber - Epoxy Composites. *Polym - Plast Technol Eng* 2012;51:1141–8. <https://doi.org/10.1080/03602559.2012.689050>.
- [20] Mehdikhani M, Gorbatikh L, Verpoest I, Lomov S V. Voids in fiber-reinforced polymer composites: A review on their formation, characteristics, and effects on mechanical performance. *J Compos Mater* 2018. <https://doi.org/10.1177/0021998318772152>.
- [21] Mehdikhani M, Gorbatikh L, Verpoest I, Lomov S V. Voids in fiber-reinforced polymer composites: A review on their formation, characteristics, and effects on mechanical performance. *J Compos Mater* 2019;53:1579–669. <https://doi.org/10.1177/0021998318772152>.
- [22] Abdelal N, Donaldson SL. Comparison of methods for the characterization of voids in glass fiber composites. *J Compos Mater* 2018;52:487–501. <https://doi.org/10.1177/0021998317710083>.
- [23] Luo L, Zhang B, Lei Y, Zhang G, Zhang Z, Meng B, et al. Identification of voids and interlaminar shear strengths of polymer-matrix composites by optical microscopy experiment and deep learning methodology. *Polym Adv Technol* 2021;32:1853–65. <https://doi.org/10.1002/pat.5226>.
- [24] Stamopoulos AG, Tserpes KI, Dentsoras AJ. Quality assessment of porous CFRP specimens using X-ray Computed Tomography data and Artificial Neural Networks. *Compos Struct* 2018;192:327–35. <https://doi.org/10.1016/j.compstruct.2018.02.096>.
- [25] Madra A, Hajj N El, Benzeggagh M. X-ray microtomography applications for quantitative and qualitative analysis of porosity in woven glass fiber reinforced thermoplastic. *Compos Sci Technol* 2014;95:50–8. <https://doi.org/10.1016/j.compscitech.2014.02.009>.
- [26] Madra A, Van-Pham DT, Nguyen MT, Nguyen CN, Breitkopf P, Trochu F. Automated identification of defect morphology and spatial distribution in woven composites. *J Compos Sci* 2020;4:1–17. <https://doi.org/10.3390/jcs4040178>.
- [27] Manzano C de JG, Ngo ACY, Sivaraja VK. Intelligent infrared thermography inspection of subsurface defects. In: Oswald-Tranta B, Zalameda JN, editors. *Thermosense Therm. Infrared Appl. XLII, SPIE*; 2020, p. 33. <https://doi.org/10.1117/12.2558958>.
- [28] Meng M, Chua YJ, Wouterson E, Ong CPK. Ultrasonic signal classification and imaging system for composite materials via deep convolutional neural networks. *Neurocomputing* 2017;257:128–35. <https://doi.org/10.1016/j.neucom.2016.11.066>.
- [29] Aggarwal CC. *Neural Networks and Deep Learning*. Cham: Springer International Publishing; 2018. <https://doi.org/10.1007/978-3-319-94463-0>.
- [30] Goodfellow I, Bengio Y, Courville A. *Deep Learning* 2016:785.
- [31] Zhou D-X. Universality of Deep Convolutional Neural Networks. *Appl Comput Harmon Anal* 2018;48:787–94. <https://doi.org/10.1016/j.acha.2019.06.004>.
- [32] Krizhevsky A, Sutskever I, Hinton GE. ImageNet classification with deep convolutional neural networks. *NIPS* 2012:1106–14. <https://doi.org/10.1145/3065386>.
- [33] Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, et al. Going deeper with convolutions. *Proc IEEE Comput Soc Conf Comput Vis Pattern Recognit* 2015;07-12-June:1–9. <https://doi.org/10.1109/CVPR.2015.7298594>.
- [34] He K, Zhang X, Ren S, Sun J. Deep Residual Learning for Image Recognition. *Proc IEEE*

- Comput Soc Conf Comput Vis Pattern Recognit 2015;2016-Decem:770–8.
<https://doi.org/10.1109/CVPR.2016.90>.
- [35] Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, et al. ImageNet Large Scale Visual Recognition Challenge. *Int J Comput Vis* 2015;115:211–52.
<https://doi.org/10.1007/s11263-015-0816-y>.
- [36] Long J, Shelhamer E, Darrell T. Fully Convolutional Networks for Semantic Segmentation. *IEEE Trans Pattern Anal Mach Intell* 2014;39:640–51.
<https://doi.org/10.1109/TPAMI.2016.2572683>.
- [37] He K, Gkioxari G, Dollár P, Girshick R. Mask R-CNN. *IEEE Trans Pattern Anal Mach Intell* 2017;42:386–97. <https://doi.org/10.1109/TPAMI.2018.2844175>.
- [38] Ronneberger O, Fischer P, Brox H. U-Net: Convolutional Networks for Biomedical Image Segmentation. *Med. Image Comput. Comput. Interv. – MICCAI 2015*, vol. 9351, 2015, p. 209–17. https://doi.org/10.1007/978-3-319-24574-4_25.
- [39] Weng Y, Zhou T, Li Y, Qiu X. NAS-Unet: Neural architecture search for medical image segmentation. *IEEE Access* 2019;7:44247–57.
<https://doi.org/10.1109/ACCESS.2019.2908991>.
- [40] Jin Q, Meng Z, Sun C, Wei L, Su R. RA-UNet: A hybrid deep attention-aware network to extract liver and tumor in CT scans 2018:1–13.
- [41] Li X, Chen H, Qi X, Dou Q, Fu CW, Heng PA. H-DenseUNet: Hybrid Densely Connected UNet for Liver and Tumor Segmentation from CT Volumes. *IEEE Trans Med Imaging* 2018;37:2663–74. <https://doi.org/10.1109/TMI.2018.2845918>.
- [42] Zeng Z, Xie W, Zhang Y, Lu Y. RIC-Unet: An Improved Neural Network Based on Unet for Nuclei Segmentation in Histology Images. *IEEE Access* 2019;7:21420–8.
<https://doi.org/10.1109/ACCESS.2019.2896920>.
- [43] Zhang Z, Liu Q, Wang Y. Road Extraction by Deep Residual U-Net. *IEEE Geosci Remote Sens Lett* 2018;15:749–53. <https://doi.org/10.1109/LGRS.2018.2802944>.
- [44] Ji S, Wei S, Lu M. Fully Convolutional Networks for Multisource Building Extraction from an Open Aerial and Satellite Imagery Data Set. *IEEE Trans Geosci Remote Sens* 2019;57:574–86. <https://doi.org/10.1109/TGRS.2018.2858817>.
- [45] Ioffe S, Szegedy C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *32nd Int Conf Mach Learn ICML 2015* 2015;1:448–56.
- [46] Nikishkov Y, Airoidi L, Makeev A. Measurement of voids in composites by X-ray Computed Tomography. *Compos Sci Technol* 2013;89:89–97.
<https://doi.org/10.1016/j.compscitech.2013.09.019>.
- [47] Melenka GW, Lepp E, Cheung BKO, Carey JP. Micro-computed tomography analysis of tubular braided composites. *Compos Struct* 2015;131:384–96.
<https://doi.org/10.1016/j.compstruct.2015.05.057>.
- [48] Kingma DP, Ba JL. Adam: A method for stochastic optimization. *3rd Int Conf Learn Represent ICLR 2015 - Conf Track Proc* 2015:1–15.