

From microstructural images to properties - an interpretable deep learning approach to predict elastic-plastic properties of fiber composites

Emmanouil Efthymiou¹, Yazhini C Pradeep², Mehtab Pathan¹, Jayaprakash Krishnasamy⁴, Vito L. Tagarielli³, Nik Petrinic⁵ and Sathiskumar Anusuya Ponnusami^{1,*}

¹Department of Mechanical Engineering and Aeronautics, City, University of London, Northampton Square, London, EC1V 0HB, UK

²Humanising Autonomy Ltd, Somerset House, Strand, London WC2R 1LA, UK

³Department of Aeronautics, Imperial College London, London, SW7 2AZ, UK

⁴Faculty of Aerospace Engineering, TU Delft, Kluyverweg 1, 2629 HS, Delft, Netherlands

⁵Department of Engineering Science, University of Oxford, Oxford, OX1 3PJ, UK

Abstract

The application of machine learning in the field of materials engineering can facilitate materials design and enable faster discovery of novel materials. This paper presents a deep learning approach for the prediction of the transverse elastic and plastic properties of unidirectional fibre reinforced composites directly from images of their microstructures. The training dataset consists of finite element predictions of the elastic-plastic properties of a set of 2D representative volume elements of unidirectional composites with different volume fraction and fiber diameter. Single-regression, Fully-connected Neural Networks (FcNNs) and a multi-regression Convolutional Neural Network (CNN) are designed and trained to predict 5 mechanical properties, using as input the images of the material's microstructure. The performance of the FcNNs and of the CNN are compared; the CNN is shown to perform the most effective regression, predicting the properties of the composites with an accuracy of 99%. A key value addition in this research lies in the explainability of the otherwise 'blackbox' CNN model, guiding the reader through the model's predictive process. We test the CNN on real, relatively low-resolution micrographs of composite microstructures and we find that this provides accurate predictions.

Keywords: composites; interpretable AI; deep learning; material properties; microstructure.

1. Introduction

Understanding the mechanical behaviour of composite materials and determining their properties is a key step to advance the applications of such materials to the aerospace, automotive and renewable energy sectors. Over the past decades, numerous researchers have proposed either analytical models or computational techniques to predict the mechanical properties of such heterogeneous materials. The process of determining the properties of a heterogeneous material at a macroscopic scale from the properties of its constituents at a lower scale is referred to as homogenisation. The underlying objective in either an analytical or a numerical homogenisation approach is to establish the microstructure-property relationships, which in turn yield the properties of an equivalent homogeneous material at the macroscale. Such an approach enables the design of composite materials and structures circumventing time-consuming material characterization experiments.

* Corresponding author. E-mail: sathiskumar.ponnusami@city.ac.uk

On one hand, analytical approaches of homogenisation are limited in their scope due to their inherent simplifying assumptions regarding the type of microstructure, modelled by unit cells or Representative Volume Elements (RVEs). Computational Homogenisation, on the other hand, provides accurate predictions albeit at very high computational cost and time, which limits their potential for practical deployment. There is a need for efficient tools to obtain quick and accurate predictions of the properties for a range of material combinations. In this context, deployment of reliable Artificial Intelligence (AI) methodologies can facilitate achieving such objective. In this current research, deep neural networks are adopted to develop a surrogate model to predict accurate material properties with negligible computational efforts.

Recent advances in machine learning have led to promising applications in the field of materials science and mechanics, ranging from understanding material behaviour and properties prediction, to design of new molecular structures and discovery of new materials [1-10]. Deep learning is a subset of machine learning based on neural networks. Recent research has demonstrated its potential in a variety of materials engineering applications. For instance, deep learning models trained on the elemental composition of materials have shown great results in the properties prediction, such as reactivity and toxicity [11, 12]. A general-purpose machine learning framework successfully used to predict properties of crystalline and amorphous materials, such as glass forming ability and band gap energy, was presented in [13].

The design of novel composite materials with tailored properties is a challenging engineering problem, due to the large search space that needs to be explored, with immediate translation and need by industry. In this regard, machine learning approaches can establish the structure-property relationships and thereby determine the properties in a quick and accurate manner. Several researchers have successfully applied machine learning techniques to predict the properties of composite materials and to optimize composite properties [14,15,16]. For the prediction of the modulus, strength, and toughness of composite materials, Abueidda et al [15] trained a 6-layer Convolutional Neural network on 4.3 million computer-generated composite configurations, achieving excellent results with more than 99% predictive accuracy in each property. In [16], a 5- layer CNN was proposed and achieved very good predictive results of the effective elastic properties of composites with arbitrary shapes and distribution of inclusions. Our research collaboration also presented a successful use of a gradient-boosted tree regression model, to accurately predict the macroscopic stiffness and yield strength of unidirectional fibre reinforced composite materials trained on the results of FE simulations [17]. In this work, the microstructural images are encoded via a two-point correlation, followed by dimensionality reduction based on principal component analysis. This yielded a set of principal components for each microstructural image, which, together with the mechanical properties predicted by FE simulations, formed the dataset for the training of the machine learning model.

The present paper focuses on fibre-reinforced composite materials and their elastoplastic response in the transversely isotropic plane. Fully connected Neural Networks (FcNNs) and a multi-regression Convolutional Neural Network (CNN), which are not well-explored in the published literature, are proposed to solve the above problem. We proceed to uncover the physical correlations learnt by the CNN model, while establishing the structure-property relationships. We also aim at demonstrating the applicability of the computational framework to real composites microstructures, while using training datasets which are computer-generated.

The article is organized as follows. Section 2 describes the production of the training datasets and the architecture of the neural networks. Section 3 examines the fidelity of the models and their interpretation. Section 4 presents the application of the model to real micrographs of

unidirectional composite materials. In the conclusions we highlight the key results from the conducted research and the scope for further work.

2. Methods

2.1. Procedure

We focus on predicting the elastoplastic properties of unidirectional (UD) fiber-reinforced composites in the transversely isotropic plane (the 2-3 plane, perpendicular to the fiber's direction). The computational framework adopted in this research starts with dataset generation, design of the deep neural networks to training and validation, as schematically shown in Fig. 1. The dataset corresponds to 1800 computer-generated images of composite microstructures and their associated mechanical properties, produced by computational homogenisation and presented in [17,18]. Five properties, namely the Young's moduli in the two transverse directions E_{22} and E_{33} , shear modulus G_{23} and two transverse yield strengths σ_{22} and σ_{33} are obtained through FE simulations and computational homogenisation of RVEs of UD composite microstructure. On the deep learning architecture, 5 single-output fully connected networks and a convolutional deep neural network were designed to perform the regression.

We then apply recently developed inference techniques to uncover the decision-making process of the neural networks [19]. This is particularly important to assist the development of tailored designs of next-generation fibre composite materials. In the following sections we provide details of each step.




Figure 1: Schematics of the modelling framework.

2.2. Dataset creation: microstructure and material properties

In this study, the dataset utilised to develop the deep learning models includes 1800 microstructural images of UD fiber composites and the corresponding mechanical properties [17]. Each microstructural image corresponds to a statistical volume element (SVE) containing random arrays of circular fibres of equal radii, embedded in a uniform matrix. Six different fiber volume fractions ϕ_f (in the range of 10-60%) and three different fibre radii, $R_f = 4, 5$ and $6 \mu\text{m}$, were considered to generate square SVEs of side length $50\mu\text{m}$. For each pair of volume fraction and radii, 100 different realizations were generated, leading to a total of 1800 microstructures. Each microstructural image forms a raw input X_i , where $i = 1, 2, \dots, 1800$ for the deep learning models. The corresponding homogenized, macroscopic mechanical properties Y_i of each microstructure image X_i , were determined using finite element simulations of the SVEs under periodic boundary conditions using the commercial FE software Abaqus/Standard. The fibres were modelled as an isotropic, linear elastic solid with properties

similar to those of glass fibres (Young's modulus $E_f = 30$ GPa, Poisson's ratio $\nu_f = 0.2$). The matrix was modelled as an isotropic, elastic-perfectly plastic solid with Young's modulus $E_m = 3$ GPa, Poisson's ratio $\nu_m = 0.3$ and yield strength $\sigma_{yield} = 120$ MPa, representative of an epoxy resin.




Figure 2: Schematics of the applied loading on the microstructural volume element: (a) uniaxial tension in direction 2, (b) uniaxial tension in direction 3, and (c) pure shear in the 2-3 plane.

The 2-D geometry of the SVEs was discretized using 4-noded plane strain quadrilateral CPE4 elements of Abaqus, following a mesh convergence study. Three quasi-static loading cases were considered as shown in Fig. 2: uniaxial tension in directions 2 and 3, as well as pure shear. To ensure an elastic-plastic transition in all cases, macroscopic normal or shear strains up to 4% were imposed. Following the analysis, the homogenized transverse Young's moduli E_{22} and E_{33} , the transverse shear modulus G_{23} and the transverse normal yields σ_{22} , σ_{33} were obtained. The yield strengths were evaluated as the flow stress at plastic normal or shear strains of 0.2%. The yield strength in shear σ_{23} was not included in the deep learning exercise, due to difficulties in convergence of some of the FE simulations for this property.

Only 1485 of the 1800 computer-generated microstructures were collected, i.e. those for which all the 5 elastoplastic properties were obtained from successful FE simulations. The resolution of the images used influences the accuracy of the machine learning application, but also the computational time associated with the training process. In this study we do not perform a dimensionality reduction step as in [17]; therefore, a resolution of 200 by 200 pixels is chosen, downsizing the original high-resolution images, to reduce the computational cost of the training process while ensuring a sufficiently detailed description of the geometry. The resulting input images were converted to square matrices of the form,

$$X_i = \begin{bmatrix} P_1 & \cdots & P_{200} \\ \vdots & \ddots & \vdots \\ P_{39801} & \cdots & P_{40000} \end{bmatrix}, \quad i = 1, 2, \dots, 1485 \quad (1)$$

where P_j is a binary pixel variable and $j = 1, 2, \dots, 40000$. The completed training dataset D is the combined list of the input arrays X_i and their corresponding output vectors Y_i

$$D = \begin{bmatrix} X_1, Y_1 \\ \vdots \\ X_{1485}, Y_{1485} \end{bmatrix}. \quad (2)$$

2.3. Fully connected Neural Network

A Neural Network or an Artificial Neural Network (ANN) is a network of neurons. An artificial neuron is a unit performing a sequence of mathematical operations that mimics the function of a biological neuron. Specifically, as it can be seen from Fig. 3, dendrites are represented by the product of inputs x_i by constants w_i (weights) and are responsible for the connection between the input and the cell body. Subsequently, all products $x_i w_i$ enter the 'cell body' of the artificial neuron, where they are summed, and a constant value b (bias) is added to the result of this summation. Finally, in order to determine whether the neuron will activate and propagate an output 'signal' y , a non-linear activation function f is applied to the result of this calculation. The artificial neurons are assembled to form a layer and a series of such layers are stacked, such that the output of one neuron y on a particular layer becomes the input x to one or more neurons in the following layer. With this structure, one can create an ANN model. An ANN model is called an FcNN when each layer in the ANN is a fully connected layer. A fully connected layer is one where each neuron in one layer is connected to all neurons in the succeeding next layer [20].




Figure 3: A visualization of an artificial neuron

In this article five identical fully connected neural networks are designed. Each of the networks receives as input a microstructural image and predicts independently one of the five mechanical properties of interest ($E_{22}, E_{33}, G_{23}, \sigma_{22}, \sigma_{33}$). In Fig. 4, the chosen architecture of a single FcNN is shown. The single FcNNs each consist of five layers: an input layer, three hidden layers and an output layer. The input layer receives the microstructural image and therefore possess 40000 binary neurons. The three hidden layers have respectively 256, 128 and 64 neurons and a sigmoid is used as the activation function for all. Finally, the output layer has one neuron, whose output is the final prediction of the model. To increase the generalization of the model and avoid over-fitting, two dropout layers of 50% probability each were included, after the first and the third hidden layers. A dropout layer randomly drops out certain nodes from the model with a given probability, to avoid over-fitting the outputs to the given data [21]. The architecture of the FcNN was the result of a systematic iterative process aimed at maximising accuracy while minimising computational efforts.




Figure 4: Design of a single fully connected neural network (FcNN).




Figure 5: Example of convolutional operation of one kernel on a 2D binary array. w is the weight of the kernel and f is the activation function. The numbers (2), (3), (4) describe the direction of the operation.

2.4. Convolutional Neural Network

A convolutional neural network (CNN) is designed with the capability to predict all the five properties of interest using the microstructural image as input. A CNN is a commonly used deep learning technique to perform predictions involving imagery data. It consists of convolutional layers followed by fully connected layers. A convolutional layer performs a convolutional operation to the input image, which is an element-wise matrix multiplication (see Fig. 5). The four main parameters of a convolutional layer are the number and size of kernels (or filters), the stride and the padding. The kernel is the matrix with which the input image will be convolved and whose weights w will be tuned during the training process. The number of kernels can be described as the equivalent of the number of neurons in the case of a fully connected layer. Each kernel, when applied on the input image, produces an output image, thus the number of kernels define the depth of the convolutional layer. The stride is the step with which the kernel will scan the input image, while the padding is the addition of extra pixels to the image that will be convolved, preventing the image from shrinking. [22]




Figure 6: Design of the convolutional neural network (CNN).

The design of the convolutional neural network developed in this paper is the result of an iterative procedure and it is summarised in Fig. 6. It consists of one convolutional layer and two fully connected layers. In particular, the convolutional layer has 32 kernels of size 3x3, stride 1 and zero padding (output image retains its size). Succeeding to the convolutional layer is a dropout layer of 50% probability, to avoid overfitting and to increase the generalization of the model. Thereafter, the 32 output images (i.e., 2D arrays of size 200 by 200 outputted by the convolutional layer) are converted to a single vector of 1,280,000 elements and move into the following fully connected layer. The last fully connected layer usually referred as the output layer consists of 5 neurons containing the predictions of the five mechanical properties ($E_{22}, E_{33}, G_{23}, \sigma_{22}, \sigma_{33}$).

2.5. Training process

The training process of the neural networks is an optimization problem shown schematically in Fig. 7. The input data X_i enters the neural network, which makes a prediction based on its initialized weights and biases. Thereafter, the prediction is compared with the target output Y_i using a loss function which quantifies the error between the predicted output and the known

ground truth value. This loss is then passed on to the training algorithm, which updates the values of weights and biases, such to minimize this loss function. The search space of the loss function is determined by the number of variables (weights and biases) of the model.




Figure 7: Representation of neural network training process.

Parameters	CNN	Fully connected
Epochs	100	200
Batch size	128	450
Cross validation folds	10	10
Training algorithm	Adam	Adam
Loss function	Mean absolute percentage error	Mean absolute percentage error

Table 1: Training parameters for the FcNN and CNN models

The training algorithm also depends on the parameters listed in Table 1 for the CNN and FCNNs used in this study. An epoch is one cycle through the full dataset and consists of certain number of training iterations. The batch size is the number of training examples (subsets of the full training dataset) in each iteration, based on which the variables of the model will update in that iteration. Thus, the number of iterations is equal to total number of training examples in the full dataset divided by batch size. In case of a batch size equal to 1, an epoch will have as many iterations/variable updates, as the total number of training examples. The selected training algorithm for the 2 models was the Adam optimizer [23], which is a specifically designed one for neural networks. In addition, to increase the models' generalization and validate their performance, they were trained with a 10-fold cross validation. Specifically, the dataset of 1485 microstructures was divided to 10 randomly populated subsets (10-fold cross validation), 9-folds of which containing data of 149 microstructures each and 1-fold of 144 microstructures. 90% of data in the subsets were used for training purposes, while the remaining 10% was used for validation. The loss function used for both regression models was the mean absolute percentage error (MAPE) which can be defined as

$$MAPE = \frac{1}{n} \sum_{i=0}^{n_{samples}} \left| \frac{y_i - \hat{y}_i}{y_i} \right| \quad (3)$$

The accuracy of the model is defined as $A = 1 - MAPE$. In eq. (3), $\hat{y}$ are the predicted values and y are the outputs (ground truth values) forming part from the training dataset. A further evaluation of the models' performance is made via the coefficient of determination r^2 defined as

$$r^2(y, \hat{y}) = 1 - \frac{\sum_{i=0}^{n_{samples}-1} (y_i - \hat{y}_i)^2}{\sum_{i=0}^{n_{samples}-1} (y_i - \bar{y})^2} \quad (4)$$

where $\bar{y} = (\sum_{i=0}^{n_{samples}} y_i) / n_{samples}$.

The parameters in Table 1 were chosen after an optimisation procedure evaluating multiple combinations of layers (0-4), nodes (8-512), kernels (8-128), activation functions (ReLU, sigmoid, tanh), as well as epochs (10-500), batch size (10-450), cross validation folds (2-20) and loss functions (mean squared error, mean absolute error and mean absolute percentage error). The training of all models was performed on a GeForce RTX2060 GPU.

2.6. Interpretation of CNN model

Understanding how an AI-based model decodes the physical correlations between the microstructural features and the mechanical properties can be of interest and can lead to constructing new constitutive models with the help of explainable artificial intelligence [24-28]. Deep neural networks are often treated as black boxes, however, recent advances in the novel field of explainable artificial intelligence have greatly improved their interpretability. Specifically, for the interpretation of the proposed CNN, we extracted 32 output images, corresponding to 32 kernels (or filters) of the network's single convolutional layer, providing an insight of what such convolutional layer detects. An output image from a convolutional layer is often called a feature map because it represents the deep features that the convolutional layer has learnt. Both, analysing these features and determining what contribution to the final predictions each of them carries aids the understanding of the model's logic.

Here we apply the Layer-wise Relevance Propagation method (LRP) [19], a backward propagation technique specifically designed for explanation of deep learning models with proved effectiveness in image applications; it helps quantifying the influence of different feature maps for a particular input/output pair and gives us a quantitative insight into the model's decision making. The general propagation rule of LRP is implemented as given by the following equation.

$$R_j = \sum_k \frac{a_j w_{jk}}{\sum_{0,j} a_j w_{jk}} R_k \quad (5)$$

Here, k and j indicate the neurons of two successive layers and $R_j: R_{j \leftarrow k}$ the share of the relevance R_k that is redistributed to neuron j . The parameter a_j represent the output value (signal) of the neuron j and w_{jk} is the weight that connects the two neurons j and k . In Fig. 8, a high-level graphical representation of the LRP procedure in a deep neural network is provided. We redirect the reader to [19] for additional details.




Figure 8: Schematic of the Layer-wise Relevance Propagation procedure. Red arrows indicate the relevance propagation flow.

3. Results and Discussion

We proceed to present the results obtained using both FcNNs and CNN and the results of the interpretation exercise conducted on the CNN model.

3.1. Accuracy of the predictions of FE results

In Figs. 9 and 10, the predicted values of the transverse elastic and plastic properties of the material, respectively, are plotted against the finite element predictions for both models. It is evident that the predictive accuracy of the CNN is much higher than the FcNNs. This is quantified in Table 2, providing the error metrics for both the FcNN and CNN predictions. From the table we can see that the CNN model achieved approximately 99% accuracy in both the elastic and yield properties, whereas the FcNN achieved approximately 98% accuracy. However, it can be noted that the FcNN model has maximum errors in the range 10.9-17.3%, while this range is only 2.6-6.6% for the CNN. The considerably better performance of the CNN emerges also by comparing the coefficients of determination of the FcNNs and CNN.




Figure 9: Comparison of the model predictions against the FE simulations (ground truth) for the elastic properties. The top row corresponds to FcNN and the bottom row corresponds to CNN results.




Figure 10: Comparison of the model predictions against the FE simulations (ground truth) for the yield strengths. The top row corresponds to FcNN and the bottom row corresponds to CNN results.

We also note that the accuracy of the CNN model is noticeably higher than the regression model proposed in [17] using the same dataset, particularly in the yield strengths, despite the CNN does not encode the images by a 2-point correlation; the addition of this preliminary step might increase the accuracy of the CNN even further.

The CNN achieved better performance than the FcNNs despite being a single multi-regression model with a very simple structure of 1 layer and 4,000,000 variables, compared to 5 independent fully connected networks of 3 layers and 10,000,000 variables each. It is clear that the self-learnt kernels in the convolutional layers are able to learn high-order relationships between neighbouring pixels, and are thus able to provide more accurate predictions of the elastoplastic properties of UD composites.

Property	MAPE		Maximum error		r^2	
	CNN	FCNNs	CNN	FCNNs	CNN	FCNNs
E_{22}	0.9%	1.9%	6.6%	16.6%	0.84	0.76
E_{33}	0.9%	1.7%	5.9%	10.9%	0.84	0.75
G_{23}	1.3%	1.8%	6.7%	12.4%	0.82	0.76
σ_{22}	0.2%	1.9%	2.6%	10.9%	0.82	0.47
σ_{33}	0.3%	2.5%	6.1%	17.3%	0.81	0.39

Table 2: Predictive performance of fully connected and convolutional networks.

3.2. Interpretation of the CNN model

In this section we present details of the predictions by the CNN, highlighting correlations between the input geometric characteristics of RVEs and the homogenised mechanical properties. For brevity we focus here on a composite with fiber volume fraction of 50%. Three selected images (out of the 32 feature images corresponding to 32 filters of the CNN model) are shown in Fig. 11. In these images the darker and lighter shades represent the undetected and detected parts of the microstructural image by three selected filters. This shows that filter 23 detects the complete interface of the circular fibres, filter 16 detects the material inside the fibres and filter 13 detects the remaining portion (matrix phase) of the microstructure. The remaining filters detected different portions of the fiber/matrix interfaces; it is anticipated that these additional filters would help detecting the presence of fibres of different cross-sectional shapes.

With the use of the LRP method, the contributions of all the filters to the final predictions are quantified. Results of this exercise are shown in Fig. 12. The bar plots show the relevance of the features detected by the CNN filters on two properties, namely E_{22} and σ_{22} , for two different volume fractions used in this study (10% and 60%). It can be observed that, regardless of the volume fraction of the microstructure, the resulted feature maps of filters 13 and 16 (matrix phase and fibre phase respectively) influence the model prediction the most. As the volume fraction increases from 10% to 60%, clearly the dominance of filter 13 (detecting fibers) decreases while the relevance of filter 16 (detecting the matrix) increases.

Taking the case of 10% fiber volume fraction, 57% of the predicted value of the Young's modulus E_{22} is ascribed to filter 13 (matrix) and 29% to the filter 16 (fiber). The next highest contributing filter, though not easily readable from the figure, is the filter 23 which detects the interface between the fiber and the matrix. This indicates that the fiber and matrix phases clearly dominate the predictions of the mechanical properties. We note that this is expected, as the mechanical response of the interface is not explicitly modelled in the FE simulations (i.e. the interfaces play no physical role in our simulations, apart from the geometric role, which is however accounted for by filters 13 and 16).

For the yield strength σ_{22} , 83% of the relevance is attributed to filter 13 and 14% to filter 16, for fiber volume fraction of 10%. These two filters combined account for 86% of the predictions of E_{22} and 96% of the predictions of σ_{22} . In contrast, for a volume fraction of 60%, 8% and 64% of the predictions of E_{22} are attributed to filters 16 and 28, respectively, while these figures are 23% and 60% for the yield strength σ_{22} . It is also interesting to note from Fig. 11 that, as the volume fraction increases from 10% to 60%, the contribution from the other filters increases both for modulus and strength.

The content of Fig. 13 can be thought of as an 'AI-based rule-of-mixtures' derived automatically by the CNN using filters 13 and 16, which determine the volume fractions of matrix and fibers. The figure summarises the contributions of the two most significant filters (13 and 16) towards the modulus and strength for the various volume fractions considered in the study. It is obvious that, as the fiber volume fraction increases, the contribution from filter 16 increases for both modulus and the strength. An interesting observation is that the contributions from both matrix and fiber filters varies nonlinearly for modulus and close to linearly for the strength. Such trends show how the CNN model has learnt and established the 'physical' correlations between the features of the microstructure and the properties.




Figure 11: Feature maps of filters 13, 16 and 23 for a selected microstructural image corresponding to 50% volume fraction with 8 μ m diameter circular fibres. The lighter (brighter) pixels are the activated ones and the darker are inactive ones.




Figure 12: Percentage relevance of the feature map of 32 filters for the prediction of E_{22} and σ_{22} for RVEs of volume fractions of 10% or 60%, with fibre diameter of 8 μ m. The results for all the volume fractions are summarised in the plots shown in Figure 13.




Figure 13: Percentage contribution of the two filters as a function of volume fraction for the prediction of E_{22} and σ_{22} . The fibre diameter corresponds to $8 \mu\text{m}$.

4. Application to real micrographs

The models described above were obtained for a single choice of mechanical properties for the fibers and matrix. Within the assumptions made in the FE simulations (isotropic linear elasticity for the fibres and matrix, incompressible J2 plasticity with isotropic hardening and perfectly plastic strain-hardening for the matrix), the five homogenized elastic-plastic properties predicted by the simulations depend on five physical parameters, in addition to the geometry of the microstructures: the elastic constants for the fibers ($E_f = 30 \text{ GPa}$, $\nu_f = 0.2$) and the matrix ($E_m = 3 \text{ GPa}$, $\nu_m = 0.3$) and the yield strength of the matrix ($\sigma_{yield} = 120 \text{ MPa}$). Dimensional analysis dictates that the non-dimensional versions of the five outputs considered in this study (obtained normalising the output dividing by E_f , E_m or σ_{yield}) depend on only four independent non-dimensional groups, for example, a non-dimensional modulus $\bar{E} = E_f / E_m$, a non-dimensional yield stress $\bar{\sigma}_y = \sigma_{yield} / E_f$ and the Poisson's ratios ν_f and ν_m . In this study it is $\bar{E} = 10$, $\bar{\sigma} = 4$, $\nu_f = 0.2$, $\nu_m = 0.3$, and the predictions of the AI models can be accurate only for this choice of non-dimensional materials properties. While we do not pursue this here, the study could be repeated considering appropriate practical ranges of the four non-dimensional groups above, to construct a more versatile predictive tool. This is left as a topic for future studies.

An important question to address is whether or not approaches like those presented in this paper can be successful when using real microstructure images, in addition to computer-generated, binary ones. To answer this question we used our predictive CNN on a micrograph of a carbon-fibre-epoxy UD composite published in [29]. This paper characterised experimentally carbon fibres of $7 \mu\text{m}$ diameter and a neat epoxy matrix by DMTA tests; these two components were then used to manufacture a UD CFRP by vacuum assisted infusion, and the viscoelastic properties of the CFRP were measured by DMTA at different temperatures and frequencies, and also compared to FE predictions (with which they were found in good agreement). The measurements on fibres and matrix in [29] gave $E_f = 25 \text{ GPa}$, $\nu_f = 0.2$ and, at 1 Hz and 35°C , $E_m = 2.5 \text{ GPa}$, $\nu_m = 0.38$. Therefore, for this CFRP it is $\bar{E} = 10$, as for the material considered in this study. The Poisson's ratio of the matrix in [29], of 0.38 is,

however, different from the value of 0.3 used in this study. Further differences are that carbon fibres are strongly anisotropic, while the FE simulations in this study assume an isotropic response of the fibers; in [29] the (real) materials are viscoelastic and tested subject to strains oscillating harmonically, while in our simulations the RVEs are subject to quasi-static, monotonic loading. Despite these differences, the material in [29] is reasonably similar to the material considered here, especially in consideration of the fact that we are focusing on the elastic properties in the (2-3) plane. In this section we attempt predictions of the elastic properties in [29] using the CNN developed in this study.

This is done by analysing an optical micrograph presented in [29] (Fig. 1a in [29]) and reproduced in Fig. 14. Fifty square portions were extracted at random locations from this image, each measuring $47\text{ }\mu\text{m}$ (corresponding to 200 pixels). We note that the average fibre radius of the carbon fibres shown in Fig. 14 was of 7 mm , therefore the ratio of side of the square to fibre diameter was of 6.7, which falls within the training range for the CNN presented in this study. Each of the 50 square portions was converted into grayscale image with the pixel brightness ranging from 0 to 255. After normalising the grayscale range to 0-1, we obtain a binary image using a threshold value of 0.7. The analysis of these binary figures provided an average fiber volume fraction of 58.7%, ranging from 52.5% to 63.2%. Such volume fractions are at the upper range of the interval (10%-60%) used to train the CNN. Examples of 2 images analysed are shown in figure 15, before and after their conversion to the binary format.

The images were provided as inputs to the CNN and the values of the mechanical properties were predicted. Here we focus on E_{22} and E_{33} , since in [29] only the transverse elastic modulus was measured. The CNN predicted an average transverse modulus of 11.6 GPa; the distribution of on E_{22} and E_{33} , normalised by the modulus of the matrix used in [29], is plotted in Fig. 16. The average non-dimensional modulus predicted by the CNN was of 3.88; the authors of [29] measured a corresponding value of 3.63, also confirmed by their FE prediction. The difference between the predictions of the CNN, which uses as input only the optical micrograph of the material shown in Fig. 14, and the measurements was therefore of only 6.5%, in line with the maximum error made by the CNN in predicting the properties of computer-generated composite microstructures; in consideration of the slight differences in properties between the real material in [29] and the virtual material used to produce our training dataset, this can be considered an excellent agreement. We note that the CNN can also be employed to estimate the spatial variability of the mechanical properties within a real material (as in Fig. 16), which is an important input when performing multiscale numerical simulations of the response of fibre composites [30].




Figure 14. Real microstructure of a CFRP composite, obtained using optical microscopy in [29].




Figure 15. Examples of the inputs provided to the CNN, consisting of randomly selected square portions of Fig. 14 measuring 200 x 200 pixels. The original figures are shown at the left, the processed binary inputs are given at the right.




Figure 16: Probability density function of the predicted normalised Young's modulus for the material in Fig. 14.

5. Conclusions

Deep neural networks were used to predict the elastic and plastic properties of RVEs of UD fiber composites of different volume fraction and fiber radius in their transversely isotropic plane, using as inputs virtual or real images of the microstructures of the composites. We proposed two models, one based on fully connected neural networks and one based on a convolutional neural network. Both models had accuracy greater than 98%, with the CNN clearly outperforming the FcNNs.

Interpretation of the CNN model, usually considered as a 'black box', was presented, to illustrate the ability of the CNN to detect relevant microstructural features and correlate these

with the mechanical properties analysed. The Layer-Wise Relevance Propagation method was applied to quantify the relevance of the different geometric features detected by the CNN on the predicted properties. The study reveals the potential of AI in establishing unidentified structure-property linkages which could be highly relevant in complex material microstructures.

The model was trained on computer-generated microstructures and corresponding properties as predicted by FE simulations, but it proved accurate also in predicting the measured properties of real composites, starting from optical micrographs of modest resolution.

As the problem at hand depends on a limited number of non-dimensional groups, it would be possible, with limited additional training, to implement a flexible predictive tool applicable to composites with arbitrary constituents. Finally, the CNN can be used to develop a generative multi-network model to designing fibre reinforced composites of tailored properties. These are left as topics for future studies.

Author Contributions

EE planned and executed the research and wrote a draft paper; MVP contributed to the dataset creation and analysis; YCP contributed to data processing, model design and interpretation study; JK & NP contributed to the discussions of the results and revision of the manuscript. VLT proposed the experimental validation of the model and contributed to discussions, writing and revision. SAP conceived and planned the research, supervised EE, coordinated the research with the co-authors and wrote the final version of the paper.

Competing interests

The authors declare no competing interests.

References

- [1] Schmidt, J., Marques, M. R., Botti, S., & Marques, M. A. (2019). Recent advances and applications of machine learning in solid-state materials science. *npj Computational Materials*, 5(1), 1-36.
- [2] Butler, K. T., Davies, D. W., Cartwright, H., Isayev, O., & Walsh, A. (2018). Machine learning for molecular and materials science. *Nature*, 559 (7715), 547-555.
- [3] Mueller, T., Kusne, A. G., & Ramprasad, R. (2016). Machine learning in materials science: Recent progress and emerging applications. *Reviews in Computational Chemistry*, 29, 186-273.
- [4] Liu, Y., Zhao, T., Ju, W., & Shi, S. (2017). Materials discovery and design using machine learning. *Journal of Materiomics*, 3(3), 159-177.
- [5] Elton, D. C., Boukouvalas, Z., Fuge, M. D., & Chung, P. W. (2019). Deep learning for molecular design—a review of the state of the art. *Molecular Systems Design & Engineering*, 4(4), 828-849.
- [6] Meredig, B. (2019). Five high-impact research areas in machine learning for materials science. *Chemistry of Materials*, 31(23), 9579–9581.
- [7] Morgan, D., & Jacobs, R. (2020). Opportunities and challenges for machine learning in materials science. *Annual Review of Materials Research*, 50, 71-103.
- [8] Agrawal, A., & Choudhary, A. (2019). Deep materials informatics: Applications of deep learning in materials science. *MRS Communications*, 9(3), 779-792.
- [9] Reimann, D., Nidadavolu, K., Vajragupta, N., Glasmachers, T., Junker, P., & Hartmaier, A. (2019). Modeling macroscopic material behavior with machine learning algorithms trained by micromechanical simulations. *Frontiers in Materials*, 6, 181.
- [10] Pilania, G., Wang, C., Jiang, X., Rajasekaran, S., & Ramprasad, R. (2013). Accelerating materials property predictions using machine learning. *Scientific reports*, 3(1), 1-6.

- [11] Goh, G. B., Hodas, N. O., & Vishnu, A. (2017). Deep learning for computational chemistry. *Journal of Computational Chemistry*, 38(16), 1291-1307.
- [12] Jha, D., Ward, L., Paul, A., Liao, W. K., Choudhary, A., Wolverton, C., & Agrawal, A. (2018). Elemnet: Deep learning the chemistry of materials from only elemental composition. *Scientific reports*, 8(1), 1-13.
- [13] Ward, L., Agrawal, A., Choudhary, A., & Wolverton, C. (2016). A general-purpose machine learning framework for predicting properties of inorganic materials. *npj Computational Materials*, 2(1), 1-7.
- [14] Chen, C. T., & Gu, G. X. (2019). Effect of constituent materials on composite performance: exploring design strategies via machine learning. *Advanced Theory and Simulations*, 2(6), 1900056.
- [15] Abueidda, D. W., Almasri, M., Ammourah, R., Ravaioli, U., Jasiuk, I. M., & Sobh, N. A. (2019). Prediction and optimization of mechanical properties of composites using convolutional neural networks. *Composite Structures*, 227, 111264.
- [16] Ye, S., Li, B., Li, Q., Zhao, H. P., & Feng, X. Q. (2019). Deep neural network method for predicting the mechanical properties of composites. *Applied Physics Letters*, 115(16), 161901.
- [17] Pathan, M. V., Ponnusami, S. A., Pathan, J., Pitisongsawat, R., Erice, B., Petrinic, N., & Tagarielli, V. L. (2019). Predictions of the mechanical properties of unidirectional fibre composites by supervised machine learning. *Scientific reports*, 9(1), 1-10.
- [18] Pathan, M. V., Tagarielli, V. L., Patsias, S., & Baiz-Villafranca, P. M. (2017). A new algorithm to generate representative volume elements of composites with cylindrical or spherical fillers. *Composites Part B: Engineering*, 110, 267-278.
- [19] Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K. R., & Samek, W. (2015). On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS one*, 10(7), e0130140.
- [20] Anderson, D., & McNeill, G. (1992). Artificial neural networks technology. *Kaman Sciences Corporation*, 258(6), 1-83.
- [21] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1), 1929-1958.
- [22] Vincent, D., & Francesco, V. (2016). A guide to convolution arithmetic for deep learning. *arXiv preprint arXiv:1603.07285*, 4.
- [23] Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- [24] Hansen, K., Baehrens, D., Schroeter, T., Rupp, M., & Müller, K. R. (2011). Visual Interpretation of Kernel-based prediction models. *Molecular Informatics*, 30(9), 817-826.
- [25] Huang, S. L., Xu, X., Zheng, L., & Wornell, G. W. (2019, July). An information theoretic interpretation to deep neural networks. In *2019 IEEE International Symposium on Information Theory (ISIT)* (pp. 1984-1988). IEEE.
- [26] Palmieri, F. A., Baldi, M., Buonanno, A., & Di Gennaro, G. (2020). Probing a Deep Neural Network. In *Neural Approaches to Dynamics of Signal Exchanges* (pp. 201-211). Springer, Singapore.
- [27] Du, M., Liu, N., & Hu, X. (2019). Techniques for interpretable machine learning. *Communications of the ACM*, 63(1), 68-77.
- [28] Ancona, M., Ceolini, E., Öztireli, C., & Gross, M. (2017). Towards better understanding of gradient-based attribution methods for deep neural networks. *arXiv preprint arXiv:1711.06104*.
- [29] Pathan, M. V., Patsias, S., Rongong, J. A., & Tagarielli, V. L. (2017). Measurements and predictions of the viscoelastic properties of a composite lamina and their sensitivity to temperature and frequency. *Composites Science and Technology*, 149, 207-219.

- [30] Schiffer, A. and Tagarielli, V.L. (2015). Predictions of the interlaminar tensile failure of a carbon/epoxy composite laminate. *Composite Structures*, 133, 997-1008.