Deep Learning–based Mechanical Characterization of Synthetically Generated Plant Fibers under Transverse Compression

Nikolaos Kegkeroglou[1*], Bilal Sabillah[2], Myrto Inglezou[3], Vincent Placet[2], Cédric Clevy[2], Antonios Porichis[3]

[1*] Tech Hive Labs, Iliados 7, Halandri, 152 32, Greece.
[2]Université Marie et Louis Pasteur, SUPMICROTECH, CNRS, Institut FEMTO-ST, Besançon, F-25000, France.
[3]AI Innovation Centre, University of Essex, Little Abington, Cambridge, CB21 6AL, UK.

*Corresponding author(s). E-mail(s): nikolaos.kegkeroglou@thlabs.eu;
Contributing authors: bilal.sabillah@femto-st.fr; mi23878@essex.ac.uk; vincent.placet@femto-st.fr; cedric.clevy@femto-st.fr; ap21186@essex.ac.uk;

Abstract

Plant fibers are emerging as promising reinforcements in composites due to their attractive mechanical performance and sustainability benefits. However, the intrinsic variability of their microstructure leads to significant scatter in their mechanical behavior, making accurate prediction of their response under transverse compression critical for the design and optimization of bio-based composites. In this study, we propose a deep learning framework to enable fast prediction of both the sequential evolution of stress fields and key mechanical properties in individual fibers under transverse loading. A dataset of synthetically generated, pseudo-realistic fiber geometries is created, and finite element simulations of transverse compression are performed to provide ground truth data. We employ a unified architecture in which two coupled recursive U-Nets predict the compression sequence: a deformation U-Net first predicts the evolving fiber geometry step by step, and a stress U-Net predicts the corresponding physical stress field at each step, conditioned on this predicted deformation. The final predicted stress field is then fed into a secondary regression convolutional neural network to estimate the transverse Young's modulus. The results demonstrate that the proposed framework accurately reproduces the evolving stress field over multiple

compression steps while providing reliable estimates of the modulus. By significantly reducing computational cost compared to finite element simulations, this framework offers a fast and reliable tool for the mechanical characterization of individual pseudo-realistic plant fibers, with promising applications in the design of natural fiber-reinforced composites.

Keywords: Plant fiber composites, Deep learning surrogate model, Stress field prediction, Transverse compression

1 Introduction

Plant fibers are increasingly being employed as renewable reinforcement constituents in composite materials owing to their favorable specific mechanical properties and sustainability advantages [1]. Over the past decade, they have gained growing importance as reinforcing agents in a wide range of applications, including automotive lightweighting [2] and soft robotics [3]. Extensive research efforts have focused on developing hybrid or fully biobased composites that combine plant fibers with biopolymer matrices, aiming to reduce reliance on non-renewable resources and mitigate environmental impact [4].

Plant Fiber Reinforced Composites (PFRCs) have emerged as strong candidates to replace conventional synthetic fiber composites such as glass fiber–reinforced composites. It was projected that by 2025, bio-based fibers could account for up to 28% of the total reinforcement materials market [5]. Compared to their synthetic counterparts, PFRCs exhibit competitive mechanical properties, while remaining environmentally friendly. Bast fibers such as flax, hemp, and nettle are considered high-performance fibers owing to their tensile strengths and favorable aspect ratios [6]. These characteristics make them attractive candidates for both structural and functional bio-composite applications.

Nevertheless, a central obstacle to the wider adoption of PFRCs lies in the transverse mechanical behavior of the fibers themselves. In unidirectional composites, the transverse direction is typically the weakest, and local stress concentrations in this plane often govern the onset of damage. Characterizing the transverse stiffness and stress response of individual fibers is therefore essential, yet remains difficult. The Single Fiber Transverse Compression Test (SFTCT) has only recently been established as a direct means of probing this behavior in plant fibers [7], and reliable measurements demand micro-mechatronic testing platforms with sub-micron precision combined with advanced microscopy [8]. Compounding this difficulty, plant fibers exhibit pronounced variability in microstructure and composition arising from differences in species, cultivation, harvesting, and processing, which translates into substantial scatter in mechanical performance and undermines predictable composite design [9].

Extracting transverse properties from such tests further relies on inverse identification through finite element simulations. While these finite element analyses yield valuable mechanistic insight, a single transverse compression analysis requires high-performance computing resources and can take several minutes per fiber, rendering

high-throughput characterization across statistically representative fiber populations impractical. This computational bottleneck motivates the search for fast, accurate surrogate models capable of bypassing repeated finite element evaluations.

Deep learning has emerged as a powerful tool for solving conventional mechanics problems, including stress field analysis [10] and mechanical property prediction [11], and is increasingly applied to fiber-reinforced composites specifically [12]. These models learn patterns directly from large datasets, capturing the nonlinear relationships between geometry and mechanical response that are otherwise difficult to express in closed form. As a result, they achieve high accuracy and robustness across diverse microstructural configurations, while operating directly on image-based representations of the material that encode its full morphological information. Once trained, such models act as fast surrogates for the underlying physics, offering near-instantaneous evaluation across many samples or loading scenarios that would be prohibitively time-consuming with traditional finite element (FE) simulations [13]. This combination of accuracy and computational efficiency makes deep learning especially attractive for high-throughput tasks, such as screening large populations of fibers or exploring the design space of bio-based composites.

Building on these developments, this study introduces a deep learning framework designed to predict the sequential evolution of stress fields and subsequent mechanical properties of individual plant fibers under transverse compression. The framework leverages a unified architecture built around two coupled recursive U-Nets, a deformation U-Net that predicts the evolving fiber geometry step by step, and a stress U-Net that predicts the physical stress field at each step, using the predicted deformation as a geometric prior, followed by a regression network that utilizes the final predicted stress field to estimate the transverse Young's modulus. Trained on a dataset of synthetically generated fiber geometries and FE-derived ground truth data, this approach provides a scalable and computationally efficient route for the mechanical characterization of natural fibers, thereby supporting the broader design and optimization of bio-based composite materials.

2 Related Work

2.1 Mechanical Properties

Recent advancements in materials science have seen a significant turn towards deep learning models for predicting the mechanical properties of composites [14]. The typical approach involves training models on large datasets, often generated through FE methods, to capture the complex, non-linear relationships between a composite's microstructure and its mechanical performance. A common approach is to regress mechanical properties directly from key microstructural features. For example, Huang et al. [15] compared several regression techniques, including deep learning models to predict the effective elastic modulus based on the features of the constituent materials. Similarly, in [16] various regression models were evaluated to predict the tensile, flexural, and impact strength of fiber composites using importance-selected features.

An alternative approach is to regress these properties directly from an image, which implicitly contains all the morphological information of the composite's phases.

For instance, Gholami et al. [17] has successfully used Convolutional Neural Networks (CNNs) to predict the Young's modulus and Poisson's ratio directly from microstructural images with a high degree of accuracy. A more recent work [18] developed customized CNNs to predict the Young's modulus and yield strength of composites from 2D images. That work also employed explainability methods to verify that the models were learning scientifically relevant geometrical features, thereby increasing their trustworthiness for engineering applications. In a similar vein, Su et al. [19] utilized a deep learning method based on 3D CNN to predict the elastic coefficients of composite materials with inclusions of arbitrary sizes, shapes and material parameters.

2.2 Stress Field

Beyond estimating material properties, deep learning is increasingly being applied to predict stress distributions. These models act as effective surrogate models that can bypass computationally expensive finite element simulations. Several studies have explored this approach for fiber-reinforced composites. In [10], a CNN based on the U-Net architecture [20] is used to predict stress field maps for composite material systems, consisting of circular fibers with varying spatial configurations. However, the use of idealized fiber shapes, such as perfect circles, may not capture the complexities of real-world materials. To address this limitation, [13] trained U-Net models on composite microstructures spanning uniform, non-uniform, and irregularly shaped fiber geometries to predict the corresponding von Mises stress maps. In [21], they explore the application of StressNet [22], a CNN originally proposed for linear elastic materials, on the prediction of stress field in 2D slices of segmented tomography images of a fiber-reinforced polymer specimen. In [23], a conditional generative adversarial network (cGAN) called StressGAN was proposed to predict stress distributions under varied geometry, load, and boundary conditions. Similarly, [24] utilized a cGAN architecture to map hierarchical composite microstructures directly to physical stress and strain fields across non-square tessellations like triangles and hexagons. In [25], a pix2pix cGAN model was developed to efficiently predict full-field principal stress maps based directly on microstructural images.

A smaller but growing body of work targets the temporal dimension of mechanical response, predicting how fields evolve under progressive loading rather than at a single instant. Wang et al. [26] combined a Temporal Independent CNN with a Bidirectional Long Short-Term Memory (Bi-LSTM) network [27] to predict the sequence of maximum internal stress during fracture propagation in brittle materials, although it captures only a scalar quantity over time rather than the full spatial field. More recently, Chen et al. [28] introduced an auto-regressive U-Net that recursively predicts the coupled evolution of stress and damage fields throughout the deformation history of carbon fiber-reinforced composites, achieving order-of-magnitude speed-ups over interface-enriched finite element analysis.

While these works demonstrate the potential of deep learning for stress field prediction in composites, they predominantly predict a complete stress field for a single, fixed configuration rather than the step-by-step evolution of the complete spatial stress field as loading progresses. Moreover, they target idealized synthetic-fiber arrangements and have not addressed individual plant fibers, whose complex, realistic geometries

and the potential presence of a lumen pose an additional challenge. Our framework addresses precisely this setting, recursively predicting the full physical stress field across successive compression steps for individual pseudo-realistic plant fibers.

3 Methodology

This section describes the complete methodology underlying the proposed framework, from data generation to model design. We first detail the construction of the synthetic dataset, in which pseudo-realistic plant fiber cross-sections are generated and subjected to finite element simulations of transverse compression to obtain ground-truth stress field sequences, together with the apparent transverse Young's modulus recovered through inverse identification. We then describe the pre-processing operations that convert the raw simulation output into a representation suitable for deep learning. Finally, we present the modeling framework: a recursive, U-Net-based architecture that predicts the evolving deformation and stress fields step by step, coupled with a regression network that estimates the transverse Young's modulus from the final predicted stress field.

3.1 Dataset Generation

The numerical workflow used to generate the pseudo-realistic fiber cross-sections is based on the method described in [29]. The generation process was based on the geometries of 20 real hemp fibers, whose cross-sectional contours were obtained from microscopic images of hemp fibers, comprising 9 fibers without lumen and 11 with lumen (Fig. 1). To ensure structural and physical realism, non-realistic geometries were filtered out during the generation loop. Approximately 40% of the initially generated cross-sections were discarded based on four geometric conformity criteria: convexity limitations, geodesic circularity thresholds, the presence of self-intersecting loops, and cases where the lumen boundary was located within less than 10% of the external fiber outer perimeter. Ultimately, 1,000 geometrically conforming pseudo-realistic fibers were retained (representing 60% of the total generated cases), evenly split into 500 fibers without a lumen and 500 fibers with a lumen.

Two-dimensional plane strain finite element simulations of transverse compression were conducted using COMSOL Multiphysics® (version 6.2) to resolve the internal stress distributions. The numerical framework, including the boundary conditions and contact mechanics, replicates the model validated in [31]. The simulation computes the stress fields at successive steps under the applied load, producing a sequence of images that shows the evolution of stress throughout the material (Fig. 2). The compression is driven by a rigid upper platen over 20 discrete displacement steps up to a maximum total displacement of 1 μm (yielding an incremental step size of 0.05 μm per step). The fiber material is modeled as linear elastic and transversely isotropic, defined by the following parameters:

$$E_{\mathrm{L}} = 50 \text{ GPa}, \quad E_{\mathrm{T}} = 1 \text{ GPa}, \quad \nu_{\mathrm{TT}} = 0.07, \quad \nu_{\mathrm{LT}} = 0, \quad G_{\mathrm{TT}} = \frac{E_{\mathrm{T}}}{2(1+\nu_{\mathrm{TT}})}$$

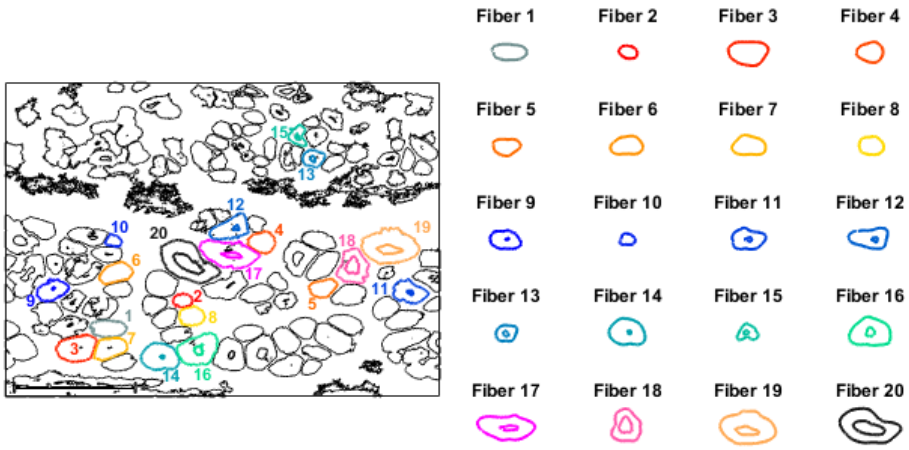


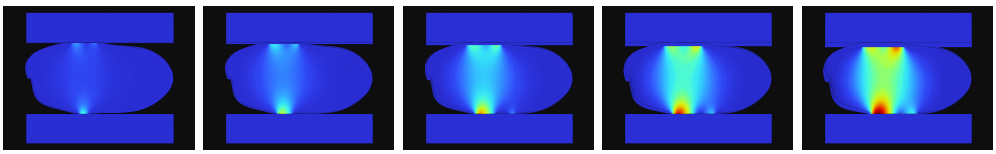
Fig. 1: Microscopic images of the real hemp fibers used to extract cross-sectional contours, showing examples of fibers with and without lumen. Reprinted from Composites Part A: Applied Science and Manufacturing, Vol 175, Jason Govilas, Cédric Clévy, Johnny Beaugrand, Vincent Placet, Investigating the influence of plant fiber geometry on apparent transverse elastic properties through finite element analysis, 107789, Copyright (2023), with permission from Elsevier [30].

where E_{L} and E_{T} represent the longitudinal and transverse Young's moduli, ν_{LT} and ν_{TT} are the Poisson's ratios in their respective planes, and G_{TT} is the transverse shear modulus. The underlying nonlinear equilibrium equations were resolved via the Multifrontal Massively Parallel Sparse Direct Solver (MUMPS). All stationary computations were performed on a high-performance computing cluster (512 GB RAM, 2.5 GHz processor architecture), requiring an average computation time of approximately 5 minutes per fiber.

Following the forward simulations, the apparent transverse Young's modulus ($E_{\mathrm{T,app}}$) of each fiber was determined via an inverse identification procedure based on experimental protocols. The compression force per unit length (F_{L}) was evaluated at each step by integrating the vertical component of the contact pressure across the contact boundary between the fiber outer boundary and the upper platen. The total radial deformation (U) corresponds directly to the prescribed platen displacement. These variables, alongside the geometric characteristics of the fiber cross-sections, were used to populate Jawad's analytical formulation [32] for the Single Fiber Transverse Compression Test (SFTCT):

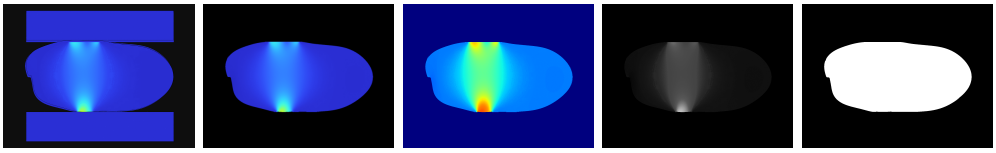
$$U = \frac{4F_{\mathrm{L}}}{\pi}\left[S_{11}\left(\sinh^{-1}\left(\frac{r}{b}\right)+\ln(2)\right)-\frac{1}{2}\left(S_{11}+S_{12}\right)+S_{12}\frac{R}{b}\sqrt{1+\left(\frac{r}{b}\right)^2}-\frac{r}{b}\right] \quad (1)$$

where b is the contact semi-width, R and r denote the major and minor radii of an ellipse fitted to the outer fiber contour, and $S_{11} = 1/E_{\mathrm{T,app}}$ and $S_{12} = -(\nu_{\mathrm{TT}}/E_{\mathrm{T,app}})$



Step 4 Step 8 Step 12 Step 16 Step 20

Fig. 2: The progressive evolution of the internal stress field output from the FE transverse compression simulations at selected loading intervals.



(a) Raw image (b) RGB image (c) JET image (d) Stress field (e) Deformation

Fig. 3: The pre-processing steps applied to the FE output sequences.

represent the components of the compliance matrix. A least-squares regression optimization was implemented to fit this analytical expression to the force-displacement data generated by the finite element model, enabling the unique identification of $E_{\mathrm{T,app}}$ for each sample.

3.2 Pre-processing

Prior to model training, we apply a series of pre-processing operations to the generated dataset to ensure its suitability for deep learning. Fig. 3 illustrates an example applied on step 8 of the FE output sequence. The raw compression image sequence from the simulation contains several undesired visual elements, including the upper and lower platens of the simulation setup and a black contour surrounding the fiber specimen (Fig. 3a). These components do not contribute to useful visual information and instead introduce noise that can hinder the learning process.

To address this, basic image processing techniques are employed to automatically remove these regions, i.e., boundary-based masking for the platens and average pooling for the black contour pixels, resulting in a cleaner and more representative dataset for subsequent model training (Fig. 3b). Next, the resulting stress fields are visualized using the JET colormap, where stress increases from blue to red (Fig. 3c). To convert these images to grayscale while preserving the relative ordering of stress values within each sample, a per-sample normalisation strategy is adopted. For each sample, all distinct RGB values from the JET colormap appearing across the entire stress field sequence are first extracted and ordered according to their position in the colormap, from blue (low stress) to red (high stress), yielding an ordered set D. The grayscale intensity range [0,1] is then discretized into 256 equally spaced bins. Each pixel is

assigned a grayscale value based on the index of its RGB value within D, so that lower-stress (more blue) pixels map to darker intensities and higher-stress (more red) pixels map to brighter intensities. This procedure yields grayscale images that represent the normalized stress field throughout the compression sequence (Fig. 3d). Following grayscale conversion, the normalized stress fields are scaled to physical units using a global min-max scaling procedure. The minimum value is defined as 0, and the maximum value is derived from the peak stress observed across the entire simulation dataset. This preserves the physical scale of the stress values within a normalized $[0, 1]$ range optimized for model training. Lastly, binary masks of the fiber are extracted from the non-zero pixels in each grayscale frame, capturing the deformation of the fiber throughout the compression sequence (Fig. 3e).

Let the sequence of physical stress field images be $\mathbf{S} = \{S_i\}_{i=0}^{N-1}$ and the sequence of fiber deformation masks be $\mathbf{B} = \{B_i\}_{i=0}^{N-1}$, where N is the number of steps in each compression sequence. These image sequences, $\mathbf{S}$ and $\mathbf{B}$, are used for self-supervised training of an image-to-image model, which is described in detail later in this work. To further increase the diversity of the training data, online data augmentation is applied via a stochastic spatial transformation operator, $\mathcal{T}_{\text{flip}}$. During training, horizontal flips are applied with a 50% probability such that $\mathbf{S}_{\text{aug}} = \mathcal{T}_{\text{flip}}(\mathbf{S})$ and $\mathbf{B}_{\text{aug}} = \mathcal{T}_{\text{flip}}(\mathbf{B})$. Crucially, the operator is applied uniformly across all N frame steps within a given sequence to preserve the exact spatial and temporal continuity between the fiber masks and their corresponding evolving stress fields. In addition, since the simulation-generated images are stored at a fixed pixel resolution that does not encode the physical dimensions of the fibers, fibers with substantially different true sizes may appear visually similar in the image domain. As the physical dimensions strongly influence the mechanical response of the fiber under compression, they constitute essential information that must be explicitly incorporated into the prediction pipeline. To this end, the major and minor radii (R, r) of each fiber sample are normalized using a global min-max scaling procedure similar to the one used for the stress fields, mapping them to a $[0, 1]$ range based on the dataset-wide minimum and maximum values. These normalized scalars are retained as part of the sample representation and are later fed directly into the image-to-image model, as described in the following section.

Alongside this, a regression model is trained to predict the mechanical properties, specifically the transverse Young's modulus (E_{T}), as derived from the analytical model [32]. In the proposed unified framework, the network takes as input the final predicted stress field generated by the stress image-to-image model. This formulation allows the learned representation of the full compression response to directly inform the estimation of effective material properties.

3.3 Modeling Framework

The overall modeling framework consists of three components, each specialized for a distinct aspect of the fiber mechanics, which together enable accurate prediction of both the evolving stress field and the resulting mechanical properties. Fig. 4 illustrates how the three models work together: the Deformation U-Net predicts the fiber deformation sequence step by step from the initial fiber mask and the Stress U-Net predicts

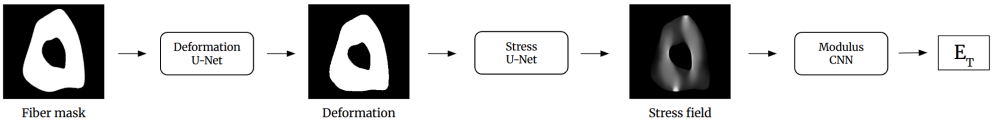


Fig. 4: Schematic representation of the deep learning pipeline. Starting from the initial fiber mask, the Deformation U-Net predicts the evolving fiber geometry step by step. At each step, the Stress U-Net predicts the corresponding physical stress field, conditioned on the predicted deformation. The final predicted stress field is then passed to the Modulus CNN, which estimates the transverse Young's modulus (E_T).

the physical stress field sequence over the same steps, while the Modulus CNN predicts the transverse Young's modulus from the final predicted stress field.

Image-to-Image Model

To model the coupled evolution of fiber deformation and internal stress fields across the compression sequence, we utilize a recursive U-Net-based image-to-image model. Its architecture follows the standard U-Net formulation, consisting of an encoder–decoder structure with skip connections. The encoder is composed of convolutional blocks with batch normalization and ReLU activation, interleaved with max-pooling to reduce spatial resolution, while the decoder symmetrically upsamples the feature maps using transposed convolutions and concatenates them with encoder features to preserve spatial detail. A final 1×1 convolution followed by a sigmoid activation produces a pixel-wise output image.

Two instances of this model are trained. The first instance, referred to as the Deformation U-Net (f), is used for deformation prediction. It takes as input the fiber mask B_i at time step i, and learns to predict the future deformation state $B_{i+\Delta}$, where Δ is a fixed stride in the compression sequence. The second instance, referred to as the Stress U-Net (g), is used for stress prediction. It takes as input the triplet $(S_i, B_i, \hat{B}_{i+\Delta})$, consisting of the current stress field, the current fiber mask, and the predicted next-step mask from f, together with the normalized ellipse axis lengths. It learns to predict the corresponding future physical stress field $S_{i+\Delta}$. Formally, the recursive prediction process is expressed as:

$$\hat{B}_{i+\Delta} = f_\theta(B_i), \qquad \hat{S}_{i+\Delta} = g_\phi(S_i, B_i, \hat{B}_{i+\Delta})$$

where θ and ϕ, denote the learnable parameters of the two network instances f and g, respectively. The normalized radii (R, r) are incorporated into g by concatenating them with the latent embedding produced by the bottleneck fully connected layer of the encoder. The resulting $(d{+}2)$-dimensional feature vector, where d denotes the bottleneck dimension, is then provided as input to the corresponding fully connected layer in the decoder before feature reconstruction, enabling g to distinguish between fibers that share similar image representations but exhibit different stress responses due to their physical dimensions.

Both networks are trained in a self-supervised manner using the simulation-derived deformation and stress sequences $\mathbf{B} = \{B_i\}_{i=0}^{N-1}$ and $\mathbf{S} = \{S_i\}_{i=0}^{N-1}$. For f, the training objective is the pixel-wise binary cross-entropy (BCE) loss between the predicted and ground-truth deformation masks:

$$\mathcal{L}_{\text{def}} = -\frac{1}{M(N-\Delta+1)} \sum_{j=1}^{M} \sum_{i=0}^{N-\Delta} \sum_{p} \left[B_{i+\Delta}^{j}(p) \log \hat{B}_{i+\Delta}^{j}(p) + (1 - B_{i+\Delta}^{j}(p)) \log(1 - \hat{B}_{i+\Delta}^{j}(p)) \right] \tag{2}$$

where M is the number of fiber samples and p indexes pixels. For g, training is performed by minimizing the MSE between predicted and ground-truth stress fields at the corresponding future time step:

$$\mathcal{L}_{\text{str}} = \frac{1}{M(N-\Delta+1)} \sum_{j=1}^{M} \sum_{i=0}^{N-\Delta} \left\| \hat{S}_{i+\Delta}^{j} - S_{i+\Delta}^{j} \right\|_2^2 \tag{3}$$

It is important to note that during training of g, the predicted masks $\hat{B}_{i+\Delta}$ from f are used as inputs rather than the ground-truth FE-derived masks $B_{i+\Delta}$. This ensures that the input distribution of g at training time matches the distribution it will encounter during inference, where only predicted future masks are available.

Regression Model

To predict the transverse Young's modulus E_{T}, we train a regression model, referred to as the Modulus CNN (h), which consists of a sequence of convolutional blocks, each containing two convolutional layers followed by batch normalization and ReLU activation. Max-pooling is applied after each block to progressively reduce spatial resolution, while the number of filters doubles across layers, forming a pyramidal encoder. The output of the final convolutional block is flattened into a bottleneck representation and passed through a hidden and an output fully connected layer to produce the mechanical property prediction. Model training is performed in a supervised setting using the final predicted stress field S_{N-1}^{j} as input, the apparent transverse Young's modulus $E_{\text{T,app}}$ as ground-truth target, and is guided by a Mean Squared Error (MSE) loss:

$$\mathcal{L}_{\text{reg}} = \frac{1}{M} \sum_{j=1}^{M} \left(\hat{E}_{\text{T}}^{j} - E_{\text{T,app}}^{j} \right)^2 \tag{4}$$

where $\hat{E_{\text{T}}^{j}}$ denotes the output prediction.

4 Results

We evaluate the proposed framework along two complementary axes that mirror its predictive tasks: the fidelity of the predicted stress field sequences and the accuracy of the estimated transverse Young's modulus. In both cases, the framework is benchmarked against representative baseline architectures drawn from the stress

Table 1: Hyperparameters of the deep learning models

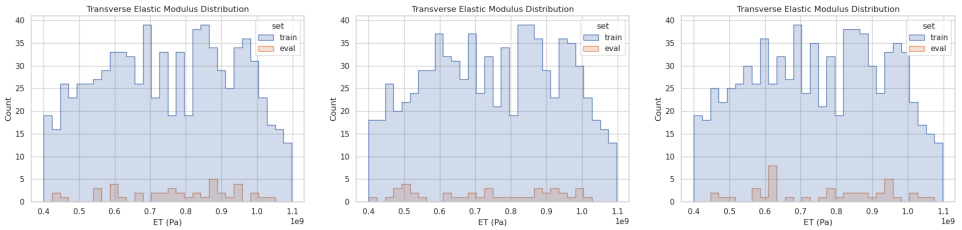
Hyperparameter	**Regression**	**Image-to-Image**
Input size	$(256, 256)$	$(256, 256)$
Number of layers	4	4
Number of initial filters	64	32
Convolution kernel	3×3	3×3
Pooling size	4	2
Bottleneck dimension	32	128
Hidden dimension	16	-
Output dimension	1	$(256, 256)$

field and mechanical property prediction literature, using quantitative metrics alongside qualitative comparisons against the finite element ground truth. To ensure that the observed performance is not an artifact of a particular initialization, all experiments are repeated across three random seeds and reported as means with standard deviations. Beyond predictive accuracy, we assess the computational efficiency of the framework relative to the finite element simulations it is designed to replace, and conduct an ablation study to isolate the individual contributions of the recursive formulation, the deformation branch, and the choice of stride.

4.1 Implementation Details

The hyperparameters chosen for the models are summarized in Table 1. All models were optimized using the Adam optimizer [33]. The Deformation U-Net was trained for 10 epochs with a batch size of 16 and a learning rate of $1{\times}10^{-2}$, while for both the Stress U-Net and Modulus CNN, 50 epochs were required at a learning rate of $3{\times}10^{-4}$. A train–validation split of 90–10 was used, where all steps generated from a given fiber simulation were kept within the same subset to prevent data leakage. The split was performed independently on the lumen and non-lumen samples and the resulting subsets were then concatenated, ensuring an even distribution of both fiber types across the training and validation sets. All experiments were conducted using three different random seeds to ensure model generalization. Figure 5 illustrates the distribution of the transverse Young's modulus across the training and evaluation splits for the different experimental seeds, showing consistent coverage of E_{T} range across all seeds. Lastly, the stride Δ was fixed at 4. Preliminary experiments showed that smaller values (e.g., $\Delta = 1$) resulted in negligible visual differences between consecutive frames, preventing the recursive U-Net from effectively learning the temporal evolution of the stress field. A stride of 4 introduces sufficiently pronounced changes to support stable learning, while still enabling prediction across five compression steps, thereby preserving an adequate level of temporal resolution.

To demonstrate the efficiency of the proposed deep learning pipeline, its computational requirements were compared against the FE method. While the initial simulation required a high-performance computing system (512 GB RAM, 2.5 GHz CPU) and averaged 5 minutes for a single fiber, the deep learning models were trained on a standard desktop NVIDIA RTX 3060 GPU, taking only 60 seconds per epoch.



Seed 1 Seed 2 Seed 3

Fig. 5: The E_{T} distribution of the training-evaluation split with different experiment seeds.

Once trained, the models achieve a significant reduction in computational overhead, reducing inference time to under a second.

[b!]

4.2 Evaluation Metrics

For the quantitative evaluation of our proposed method, a set of complementary metrics was employed to assess both image sequence prediction and regression performance. All reported metrics represent the mean values across the different experimental seeds, with the corresponding standard deviations.

For the image-to-image models, evaluation was performed using the Structural Similarity Index Measure (SSIM) [34] ranging from 0 to 1, where 1 indicates perfect structural similarity, computed between the predicted and ground-truth images. In addition, the Root Mean Squared Error (RMSE) was used to quantify the pixel-wise deviation between the predicted and ground-truth stress fields. Because the stress fields are scaled by a single dataset-wide maximum, fibers with low peak stress occupy only a small portion of the [0, 1] range, and an unnormalized RMSE would be dominated by the few high-stress samples. To place all fibers on a comparable footing and isolate the fidelity of the predicted spatial distribution from differences in absolute magnitude, each field is rescaled to its own [0, 1] range before the error is computed. RMSE is thus defined as:

$$\mathrm{RMSE} = \frac{1}{M} \sum_{j=1}^{M} \sqrt{\sum_{i=0}^{N-\Delta} \left\| c^j \cdot \hat{S}^j_{i+\Delta} - c^j \cdot S^j_{i+\Delta} \right\|_2^2} \tag{5}$$

where c^j is the sample-specific scaling factor, defined by the ratio of the dataset's global maximum to the fiber's maximum stress, which rescales each globally-scaled field to a per-sample [0, 1] range. The deformation masks are already binary and require no such rescaling. Lower RMSE values indicate better agreement between the predicted and ground-truth images.

For the regression model, the performance is evaluated using three complementary metrics. The coefficient of determination (R^2), ranging from 0 to 1, where 1 indicates

a perfect fit, measures the proportion of variance in the target values explained by the model:

$$\mathrm{R}^2 = 1 - \frac{\sum_{j=1}^{M} \left(y^j - \hat{y}^j\right)^2}{\sum_{j=1}^{M} \left(y^j - \bar{y}\right)^2} \tag{6}$$

where y^j and $\hat{y}^j$ denote the ground-truth and predicted values, respectively, and $\bar{y} = \frac{1}{M}\sum_{j=1}^{M} y^j$ is the mean of the ground-truth values. The Mean Absolute Error (MAE), expressed in physical units, quantifies the average magnitude of prediction errors:

$$\mathrm{MAE} = \frac{1}{M}\sum_{j=1}^{M} \left|\hat{y}^j - y^j\right| \tag{7}$$

The Mean Absolute Percentage Error (MAPE), expressed in percentage (%), provides a scale-independent measure of relative error:

$$\mathrm{MAPE} = \frac{100}{M}\sum_{j=1}^{M} \left|\frac{\hat{y}^j - y^j}{y^j}\right| \tag{8}$$

4.3 Experimental Results

The performance of the proposed framework is evaluated on both the image-to-image and regression components, compared against representative baseline models from the literature.

Image-to-Image Model. Table 2 reports the stress field prediction performance of the proposed Stress U-Net (g) alongside four baseline architectures drawn from the stress field prediction literature. The Stress U-Net achieves an SSIM of 0.9574±0.0074 and an RMSE of 0.0345±0.0079, with low variance across experimental seeds indicating robust generalization. The compared models include StressNet [22], an encoder–decoder network originally proposed for stress prediction in linear elastic materials. We also consider a StressNet variant used in [13, 21], referred to here as StressNet-mod, which omits the skip connections between encoder and decoder. In addition, we benchmark against U-Net-mod [10], a slightly modified version of the standard U-Net applied to composite stress field prediction, and SCSNet [22], a lighter segmentation-based network from the same work. The proposed model outperforms

Table 2: Comparative evaluation of stress field prediction models

Model	**SSIM** (0–1) ↑	**RMSE** (0–1) ↓
Stress U-Net (ours)	**0.9574±0.0074**	**0.0345±0.0079**
StressNet [22]	0.9372±0.0048	0.0438±0.0022
StressNet-mod [13, 21]	0.9172±0.0240	0.0479±0.0094
U-Net-mod [10]	0.9170±0.0218	0.0563±0.0090
SCSNet [22]	0.8076±0.0290	0.0894±0.0023

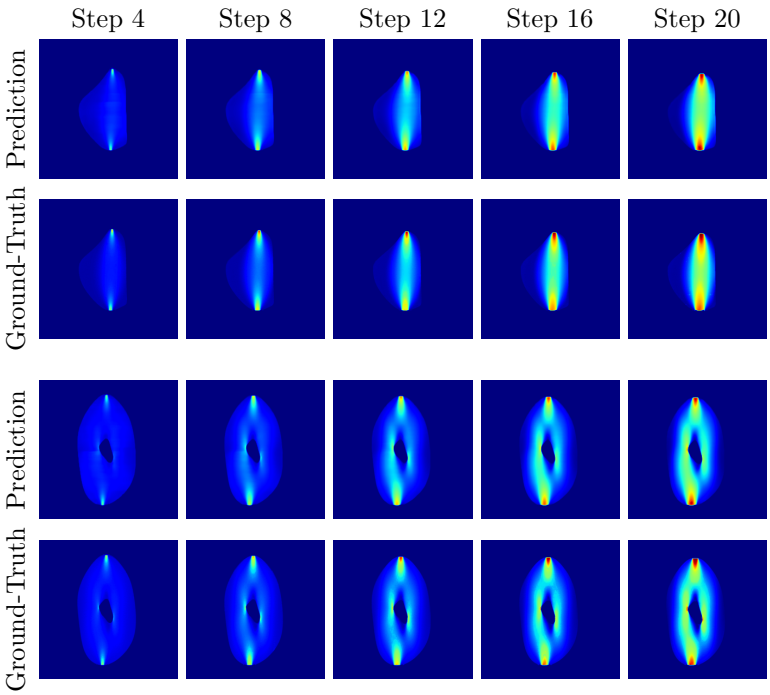


Fig. 6: Qualitative evaluation of the proposed framework for stress field prediction of individual pseudo-realistic plant fibers under transverse compression.

all other models on both metrics, demonstrating that its design is well-suited to the dynamic nature of the task. The final configuration was selected through extensive hyperparameter tuning, and the contribution of individual architectural choices to this performance is examined in detail in the ablation study section (4.4). In contrast, the baseline architectures were originally developed for idealized geometries or single-step stress prediction. In our dynamic setting, we operate on individual fibers with pseudo-realistic geometries derived from real plant-based fibers, which further increases the complexity of the task.

Figure 6 presents representative examples of the predicted versus ground-truth stress fields for individual pseudo-realistic plant fibers under transverse compression. Each pair of rows corresponds to a single fiber, with the top row showing the predicted stress field and the bottom row showing the ground-truth. Columns indicate sequential compression steps, highlighting the temporal evolution of the stress distribution. The first sample corresponds to a fiber without a lumen, while the second sample contains a lumen, illustrating the model's ability to capture structural variations. For visualization purposes, the physical stress fields predicted by g are mapped back to the JET colormap. Overall, the figure demonstrates that the models accurately reproduce both the spatial patterns and magnitudes of the stress field across different fiber geometries, illustrating the robustness of the framework.

Table 3: Comparative evaluation of regression models for transverse Young's modulus prediction

Model	$\mathbf{R^2}$ (0–1) ↑	**MAE** (MPa) ↓	**MAPE** (%) ↓
Modulus CNN (ours)	**0.938±0.024**	**24.31±1.23**	**5.25±0.91**
ResNet-34 [35]	0.895±0.059	30.48±6.84	6.89±1.70
AlexNet [36]	0.881±0.088	28.42±3.91	6.67±0.49
TI-CNN [26]	0.837±0.079	35.94±7.07	8.10±0.98
CNN-custom [18]	0.754±0.014	48.04±8.89	10.49±2.11

Regression Model. Table 3 reports the performance of the Modulus CNN (h) alongside four baseline models adapted to the same regression task. The proposed model achieves an R^2 of 0.938 ± 0.024, an MAE of 24.31 ± 1.23 MPa, and a MAPE of $5.25 \pm 0.91\%$, outperforming all baselines across every metric. The compared models include ResNet-34 [35] and AlexNet [36], well-established convolutional architectures that have shown strong performance across a broad range of tasks, including in the context of fiber composite analysis [13]. Other architectures are the TI-CNN [26], a temporal architecture originally proposed for predicting maximum internal stress during fracture in brittle materials, and a customised CNN [18], referred to here as CNN-custom, developed for Young's modulus prediction from composite microstructure images. The proposed model outperforms all baselines across every metric. This is primarily attributed to its use of the final predicted stress field as input rather than the raw fiber geometry, allowing the learned representation of the full compression response to directly inform the modulus estimate — a formulation that none of the baselines exploit.

4.4 Ablation Study

To assess the contribution of each design decision in the proposed framework, three ablation variants of the image-to-image model are evaluated and compared against the full model in Table 4.

Stride-1 This variant applies the full two-U-Net recursive architecture, but uses a stride of $\Delta = 1$ instead of the adopted $\Delta = 4$. Both f and g are trained and evaluated with unit-step transitions, so that $\hat{B}_{i+1} = f_\theta(B_i)$ and $\hat{S}_{i+1} = g_\phi(S_i, B_i, \hat{B}_{i+1})$, requiring 20 recursive rollout steps to reach the same final compression state.

Table 4: Ablation study of the image-to-image model

Model	**Target**	**SSIM** (0–1) ↑	**RMSE** (0–1) ↓
Full	Deformation	**0.9822±0.0015**	**0.0761±0.0037**
Stride-1	Deformation	0.9805 ± 0.0024	0.0820 ± 0.0053
Full	Stress	**0.9574±0.0074**	**0.0345±0.0079**
Stride-1	Stress	0.8923 ± 0.0083	0.0734 ± 0.0102
No-def	Stress	0.8235 ± 0.0290	0.0903 ± 0.0137
No-rec	Stress	0.5162 ± 0.0316	0.1333 ± 0.0014

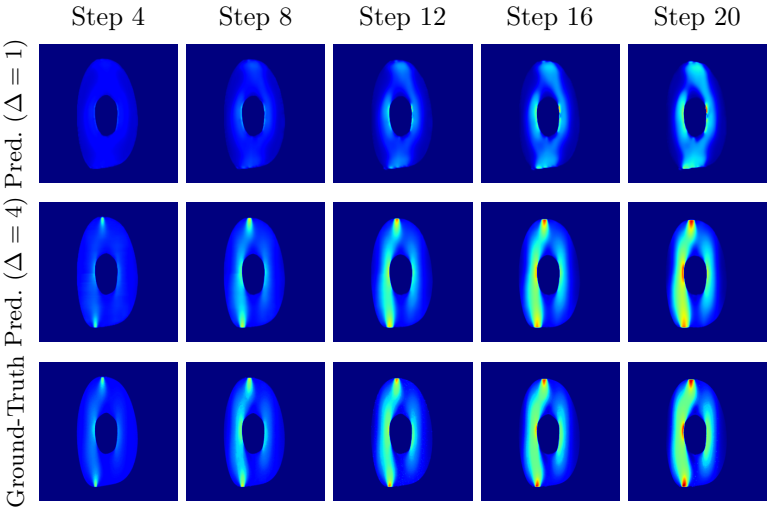


Fig. 7: Qualitative comparison of rollout predictions for stride values $\Delta = 1$ and $\Delta = 4$ across sequential evaluation steps.

No-def This variant replaces the two-U-Net cascade with a single stress U-Net g'_ϕ. Since the predicted deformation mask $\hat{B}_{i+\Delta}$ from f is no longer available, g'_ϕ receives only two input channels: the initial binary mask B_0 and the current stress field S_i. The model is still trained and evaluated recursively with stride $\Delta = 4$, so that $\hat{S}_{i+\Delta} = g'_\phi(S_i, B_0)$ at each step.

No-rec This variant removes both the deformation U-Net and the recursive rollout mechanism entirely. A single U-Net g''_ϕ is trained in a one-shot, non-recursive fashion: it receives the initial binary mask B_0 as its sole input and directly outputs the full predicted stress sequence as a five-channel image, $g''_\phi(B_0) = (\hat{S}_\Delta, \hat{S}_{2\Delta}, \hat{S}_{3\Delta}, \hat{S}_{4\Delta}, \hat{S}_{5\Delta})$.

The results in Table 4 demonstrate that each design choice contributes meaningfully to the model's performance. For the deformation target, Stride-1 and the Full model yield nearly equivalent SSIM scores (0.9805 vs. 0.9822). Because B_i is a binary mask, even a unit-stride transition produces a discrete, non-negligible change at the fiber boundary, i.e. pixels either belong to the fiber or they do not, providing a learnable signal regardless of Δ. The picture is markedly different for the stress target. The Full model achieves an SSIM of 0.9574 and an RMSE of 0.0345, whereas Stride-1 drops to 0.8923 and 0.0734 respectively, a substantial degradation despite its finer temporal resolution. This is further illustrated in Figure 7, which shows that $\Delta = 1$ rollout predictions exhibit noticeable loss of the stress flow compared to both the $\Delta = 4$ predictions and the ground-truth, consistent with the observation that unit-stride transitions produce near-identical consecutive frames, offering insufficient signal for the stress U-Net to learn meaningful temporal dynamics. Removing the deformation branch (No-def) leads to a further decline in stress prediction accuracy (SSIM = 0.8235, RMSE = 0.0903) and notably increased variance across seeds, highlighting the importance of the predicted future mask $\hat{B}_{i+\Delta}$ as an explicit geometric prior that guides the stress evolution at each step. The most severe degradation is

observed for No-rec (SSIM = 0.5162, RMSE = 0.1333), which collapses to near-chance structural fidelity, underscoring that the recursive formulation is critical. Compressing the entire compression sequence into a single feed-forward mapping discards the incremental temporal structure that the model relies on, and multi-step prediction from the undeformed mask alone is insufficient to capture the load-dependent nature of the internal stress fields.

5 Conclusion

In this study, a deep learning framework is proposed for the prediction of stress fields and mechanical properties in individual pseudo-realistic plant fibers under transverse compression. The framework consists of three specialized components: a Deformation U-Net, that predicts the evolving fiber geometry step by step, a Stress U-Net, that recursively predicts the physical stress field conditioned on the current stress state and the predicted deformation, and a Modulus CNN, that estimates the transverse Young's modulus from the final predicted stress field. The framework recovers the full physical stress field sequence with high fidelity, achieving an SSIM of 0.9574 ± 0.0074 and an RMSE of 0.0345 ± 0.0079, and accurately predicts the transverse Young's modulus with an R^2 of 0.938 ± 0.024 and a MAPE of $5.25 \pm 0.91\%$. These results are consistent across fibers with and without a lumen and across different experimental seeds, demonstrating robust generalization of the proposed approach.

Comparative evaluation against representative baseline architectures from the stress field and mechanical property prediction literature confirms that the proposed models outperform all compared methods on every metric. An ablation study further establishes that the recursive formulation, the deformation branch, and the choice of stride are all critical contributors to this performance.

In addition to its predictive accuracy, an important advantage of the proposed framework is its computational efficiency. Whereas a single FE-based transverse compression analysis requires several minutes on a high-performance computing cluster, the trained models reduce inference time to under a second per fiber, enabling rapid characterization across large populations of fiber geometries. This speed-up is achieved without significant loss of accuracy, making the framework a practical surrogate for finite element simulations in the context of bio-based composite design.

A primary limitation of this work is that the framework is trained and evaluated entirely on synthetically generated fibers with finite element-derived ground truth. While the geometries are pseudo-realistic and the simulation setup follows an experimentally validated model, the predictions have not yet been validated against physical single-fiber compression measurements.

Future work could extend the framework to three-dimensional fiber geometries, other loading modes, or real microscopy images of plant fibers, further closing the gap between synthetic training data and experimental conditions. Incorporating uncertainty quantification into the predictions would additionally increase the trustworthiness of the framework for engineering applications.

Declarations

Funding. The work was supported by the European Union through the RAIDO Project under Grant No. 101135800.

Acknowledgements. Support from the EIPHI Graduate School (ANR-17-EURE-0002) is also acknowledged.

Conflict of interest. All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript.

References

[1] Rohit, K., Dixit, S.: A review - future aspect of natural fiber reinforced composite. Polymers from Renewable Resources **7**(2), 43–59 (2016) https://doi.org/10.1177/204124791600700202 https://doi.org/10.1177/204124791600700202

[2] Skosana, S.J., Khoathane, C., Malwela, T.: Driving towards sustainability: A review of natural fiber reinforced polymer composites for eco-friendly automotive light-weighting. Journal of Thermoplastic Composite Materials **38**, 754–780 (2024)

[3] Sarker, A., Ul Islam, T., Islam, M.R.: A review on recent trends of bioinspired soft robotics: Actuators, control methods, materials selection, sensors, challenges, and future prospects. Advanced Intelligent Systems **7**(3), 2400414 (2025) https://doi.org/10.1002/aisy.202400414 https://advanced.onlinelibrary.wiley.com/doi/pdf/10.1002/aisy.202400414

[4] Vignesh, P., Gopalarama Subramaniyan, G.: Advancements in biopolymer-based natural fiber composites: Synthesis, characterization, and applications. Surface Review and Letters **33**(04), 2530003 (2026) https://doi.org/10.1142/S0218625X25300035 https://doi.org/10.1142/S0218625X25300035

[5] Ramesh, M., Palanikumar, K., Reddy, K.H.: Plant fibre based bio-composites: Sustainable and renewable green materials. Renewable and Sustainable Energy Reviews **79**, 558–584 (2017) https://doi.org/10.1016/j.rser.2017.05.094

[6] Ardanuy, M., Claramunt, J., Toledo Filho, R.D.: Cellulosic fiber reinforced cement-based composites: A review of recent research. Construction and Building Materials **79**, 115–128 (2015) https://doi.org/10.1016/j.conbuildmat.2015.01.035

[7] Govilas, J., Chevallier, A., Akleh, W., Beaugrand, J., Clévy, C., Placet, V.: Mechanical behavior and apparent stiffness of flax, hemp and nettle fibers under single fiber transverse compression tests. Composites Part A: Applied Science and Manufacturing **185**, 108321 (2024) https://doi.org/10.1016/j.compositesa.2024.108321

[8] Placet, V., Blot, M., Weemaes, T., Bernollin, H., Laurent, G., Amiot, F., Clévy, C., Beaugrand, J.: Transverse compressive properties of natural fibers determined using micro mechatronic systems and 2d full-field measurements. Materials Today: Proceedings **31**, 303–308 (2020) https://doi.org/10.1016/j.matpr.2020.01.383 . 4th International Conference on Natural Fibers – Smart Sustainable Materials, ICNF 2019

[9] Khan, T., Sultan, M.T.B.H., Ariffin, A.H.: The challenges of natural fiber in manufacturing, material selection, and technology application: A review. Journal of Reinforced Plastics and Composites **37**(11), 770–779 (2018) https://doi.org/10.1177/0731684418756762 https://doi.org/10.1177/0731684418756762

[10] Bhaduri, A., Gupta, A., Graham-Brady, L.: Stress field prediction in fiber-reinforced composite materials using a deep learning approach. Composites Part B: Engineering **238**, 109879 (2022) https://doi.org/10.1016/j.compositesb.2022.109879

[11] Pathan, M.V., Ponnusami, S.A., Pathan, J., Pitisongsawat, R., Erice, B., Petrinic, N., Tagarielli, V.L.: Predictions of the mechanical properties of unidirectional fibre composites by supervised machine learning. Scientific Reports **9**(1), 13964 (2019) https://doi.org/10.1038/s41598-019-50144-w

[12] Malashin, I.P., Martysyuk, D., Nelyub, V., Borodulin, A., Gantimurov, A., Tynchenko, V.: Deep learning for property prediction of natural fiber polymer composites. Scientific Reports **15**(1), 27837 (2025) https://doi.org/10.1038/s41598-025-10841-1

[13] Shokrollahi, Y., Nikahd, M.M., Gholami, K., Azamirad, G.: Deep learning techniques for predicting stress fields in composite materials: A superior alternative to finite element analysis. Journal of Composites Science **7**(8) (2023) https://doi.org/10.3390/jcs7080311

[14] Karuppusamy, M., Thirumalaisamy, R., Palanisamy, S., Nagamalai, S., El Sayed Massoud, E., Ayrilmis, N.: A review of machine learning applications in polymer composites: advancements, challenges, and future prospects. J. Mater. Chem. A **13**, 16290–16308 (2025) https://doi.org/10.1039/D5TA00982K

[15] Huang, H., Hadigheh, S.A., Aghabalaei Baghaei, K.: Influences of fibre shape on the transverse modulus of unidirectional fibre reinforced composites using finite element and machine learning methods. Composite Structures **312**, 116872 (2023) https://doi.org/10.1016/j.compstruct.2023.116872

[16] Singh, A.P., Asgar, M.E., Singholi, A.K.S., Gautam, U., Ranjan, R.: Machine learning approaches for predicting mechanical properties of composite materials (2025) https://doi.org/10.21467/proceedings.178.5 . Online ISSN: 2582-3922

[17] Gholami, K., Ege, F., Barzegar, R.: Prediction of composite mechanical properties: Integration of deep neural network methods and finite element analysis. Journal of Composites Science **7**(2) (2023) https://doi.org/10.3390/jcs7020054

[18] Raaghav, V., Bikos, D., Rago, A., Toni, F., Charalambides, M.: Explainable Prediction of the Mechanical Properties of Composites with CNNs (2025). https://arxiv.org/abs/2505.14745

[19] Su, H., Guan, T., Liu, Y.: A three-dimensional prediction method of stiffness properties of composites based on deep learning. Computational Mechanics **71**(3), 583–597 (2023) https://doi.org/10.1007/s00466-022-02253-z

[20] Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation (2015). https://arxiv.org/abs/1505.04597

[21] Sun, Y., Hanhan, I., Sangid, M.D., Lin, G.: Predicting mechanical properties from microstructure images in fiber-reinforced polymers using convolutional neural networks. Journal of Composites Science **8**(10) (2024) https://doi.org/10.3390/jcs8100387

[22] Nie, Z., Jiang, H., Kara, L.B.: Stress field prediction in cantilevered structures using convolutional neural networks. Journal of Computing and Information Science in Engineering **20**(1) (2019) https://doi.org/10.1115/1.4044097

[23] Jiang, H., Nie, Z., Yeo, R., Farimani, A.B., Kara, L.B.: Stressgan: A generative deep learning model for two-dimensional stress distribution prediction. Journal of Applied Mechanics **88**(5) (2021) https://doi.org/10.1115/1.4049805

[24] Yang, Z., Yu, C.-H., Buehler, M.J.: Deep learning model to predict complex stress and strain fields in hierarchical composites. Science Advances **7**(15), 7416 (2021) https://doi.org/10.1126/sciadv.abd7416 https://www.science.org/doi/pdf/10.1126/sciadv.abd7416

[25] Gupta, S., Mukhopadhyay, T., Kushvaha, V.: Microstructural image based convolutional neural networks for efficient prediction of full-field stress maps in short fiber polymer composites. Defence Technology **24**, 58–82 (2023) https://doi.org/10.1016/j.dt.2022.09.008

[26] Wang, Y., Oyen, D., Guo, W.G., Mehta, A., Scott, C.B., Panda, N., Fernández-Godino, M.G., Srinivasan, G., Yue, X.: Stressnet - deep learning to predict stress with fracture propagation in brittle materials. npj Materials Degradation **5**(1), 6 (2021) https://doi.org/10.1038/s41529-021-00151-y

[27] Schuster, M., Paliwal, K.K.: Bidirectional recurrent neural networks. IEEE Transactions on Signal Processing **45**(11), 2673–2681 (1997) https://doi.org/10.1109/78.650093

[28] Chen, Z., Yacouti, M., Shakiba, M., Wang, J.-X., Luo, T., Varshney, V.: Predicting stress and damage in carbon fiber-reinforced composites deformation process using composite u-net surrogate model. Theoretical and Applied Mechanics Letters **16**(1), 100634 (2026) https://doi.org/10.1016/j.taml.2025.100634

[29] Mattrand, C., Béakou, A., Charlet, K.: Numerical modeling of the flax fiber morphology variability. Composites Part A: Applied Science and Manufacturing **63**, 10–20 (2014) https://doi.org/10.1016/j.compositesa.2014.03.020

[30] Govilas, J., Clévy, C., Beaugrand, J., Placet, V.: Investigating the influence of plant fiber geometry on apparent transverse elastic properties through finite element analysis. Composites Part A: Applied Science and Manufacturing **175**, 107789 (2023) https://doi.org/10.1016/j.compositesa.2023.107789

[31] Govilas, J.: Plant fiber mechanical characterization with high precision micromechatronic means : investigation of single fiber transverse behavior and interfiber adhesion. Theses, Université Bourgogne Franche-Comté (February 2023). https://theses.hal.science/tel-04472025

[32] Jawad, S.A., Ward, I.M.: The transverse compression of oriented nylon and polyethylene extrudates. Journal of Materials Science **13**(7), 1381–1387 (1978)

[33] Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization (2017). https://arxiv.org/abs/1412.6980

[34] Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE Transactions on Image Processing **13**(4), 600–612 (2004) https://doi.org/10.1109/TIP.2003.819861

[35] He, K., Zhang, X., Ren, S., Sun, J.: Deep Residual Learning for Image Recognition (2015). https://arxiv.org/abs/1512.03385

[36] Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet Classification with Deep Convolutional Neural Networks. Curran Associates, Inc. (2012). https://proceedings.neurips.cc/paper_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf