Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Bayesian Physics-Informed Neural Networks Embedded in Finite Element Analysis: A Probabilistic Digital Twin for Reliability-Based Design of FDM Polymer Components

Mohammad Hadi Zahmatkeshan
Amirkabir University of Technology
Farid Reza Biglari
biglari@aut.ac.ir

Amirkabir University of Technology
Bijan Mollaei Dariani
Amirkabir University of Technology

Research Article

Keywords: Probabilistic Digital Twin, Bayesian Neural Networks, Physics-Informed Machine Learning, User Material Subroutine, Finite Element Analysis, Uncertainty Quantification, Additive Manufacturing, Fused Deposition Modeling, Reliability Analysis

Posted Date: November 20th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11305412/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Bayesian Physics-Informed Neural Networks Embedded in Finite Element Analysis: A Probabilistic Digital Twin for Reliability-Based Design of FDM Polymer Components

Mohammad Hadi Zahmatkeshan[1], Farid Reza Biglari[1]*, Bijan Mollaei Dariani[1]
1 Mechanical Engineering Department, Amirkabir University of Technology, Tehran 591634311, Iran
*Correspondence: biglari@aut.ac.ir

Abstract

Digital twin technology presents revolutionary possibilities for additive manufacturing, but current applications are limited to deterministic forecasting that does not capture manufacturing uncertainties. This study introduces a probabilistic digital twin system that uses machine learning and finite element analysis in two synergistic innovations. First, a Bayesian physics-informed neural network predicts material properties as Abaqus input with quantified uncertainty by embedding Euler-Bernoulli beam theory within a Monte Carlo dropout architecture. Second, the trained network's deterministic forward pass is embedded directly within an Abaqus user material subroutine, enabling material property evaluation at each integration point for a given sampled modulus value, while the probabilistic Monte Carlo dropout sampling is performed externally prior to each finite element run. This framework was constructed based on experimental data of sixty 3D-printed polylactic acid specimen groups subjected to three-point bending in accordance with ISO 178 , with validation against ten independent configurations with raster angles ranging from zero to ninety degrees. Monte Carlo finite element analysis propagated material property uncertainty through the simulation, yielding stress predictions with calibrated confidence intervals. The framework achieved a coefficient of determination of 0.9935 with a mean absolute percentage error of 1.85%, and 95% confidence intervals contained experimental values in all ten validation cases, consistent with (though not sufficient to precisely confirm, given the limited sample size) well-calibrated uncertainty estimates. Reliability analysis revealed that the 60° raster angle (8.0% CV) and 90°/80% reduced infill configuration (9.0% CV) exhibit more than double the prediction uncertainty compared to the optimal 0° orientation (3.8% CV), informing correspondingly adjusted safety factors. The computational overhead of the embedded neural network added only 15% to simulation time while eliminating the need for external property pre-computation. This work demonstrates a full probabilistic digital twin framework for additively manufactured polymer components under flexural loading, enabling reliability-based design choices that are challenging to implement under deterministic methods. The framework provides a foundation for advancing additive manufacturing toward structural applications that require a quantified performance margin.

Keywords: Probabilistic Digital Twin; Bayesian Neural Networks; Physics-Informed Machine Learning; User Material Subroutine; Finite Element Analysis; Uncertainty Quantification; Additive Manufacturing; Fused Deposition Modeling; Reliability Analysis

Highlights:

- An uncertainty-aware physics-informed neural network / finite-element framework for FDM polymers, integrating Bayesian material property prediction with reliability-based design
- Neural network embedded in FEA via UMAT achieves 1.85% error with 15% computational overhead
- SEM fractography is qualitatively consistent with the hypothesis that configuration-dependent uncertainty (CV 3.8–9.0%) reflects distinct failure mechanisms
- 95% confidence intervals successfully captured all experimental values across 10 validation cases
- Configuration-dependent safety factors (1.11-1.30) derived from reliability analysis for design

1. Introduction

1.1 Digital Twins in Additive Manufacturing

The digital twin concept has been developed as the pillar of Industry 4.0, allowing the bidirectional adaptation between physical assets and their virtual ones [1,2]. Digital twins have specific utility in additive manufacturing since the layer-by-layer print process forms a complex interdependence between process design and final part characteristics that are challenging to model using traditional engineering techniques [3,4]. With a fully realized digital twin, manufacturers would have ability to predict mechanical behavior, optimize process parameters, and ensure quality without conducting an extreme amount of physical testing to be able to cover all design variations.

Despite the transformative potential of digital twins for additive manufacturing, two fundamental barriers prevent their deployment in reliability-critical structural applications. First, existing approaches treat material property prediction and finite element simulation as disconnected workflows requiring manual data transfer [5-7], creating workflow fragmentation and preventing seamless uncertainty propagation through the analysis chain. Second, and more critically, virtually all current digital twin implementations provide deterministic point predictions that ignore the substantial manufacturing variability inherent to layer-by-layer fabrication processes [8,9]. Without quantified confidence intervals and reliability metrics, design engineers cannot rationally specify safety factors or assess whether a given part configuration will meet performance requirements with acceptable probability. This limitation is particularly severe for fused deposition modeling of polymer components, where mechanical anisotropy, interlayer bonding variability, and complex interactions between multiple process parameters produce configuration-dependent uncertainty levels that can vary by factors of two or more. The central research question addressed by this work is: Can a probabilistic digital twin framework simultaneously achieve (1) prediction accuracy comparable to deterministic methods, (2) calibrated uncertainty estimates validated against experimental observations, and (3) seamless workflow integration suitable for practical engineering deployment?

1.2 Mechanical Anisotropy in Fused Deposition Modeling

Fused deposition modeling produces extreme mechanical anisotropy due to the layer-by-layer deposition of thermoplastic filament [10-12]. Strength and stiffness measured parallel to the deposited beads differ significantly from those measured at right angles, and layer interfaces tend to be sites of failure initiation. In printed components, finite element predictions have been reported to have errors of seven to twenty percent when isotropic material models are used [13-15], with particularly severe discrepancies under flexural loading, where stress gradients probe anisotropic properties throughout the specimen thickness. Defining the raster angle is essential for interpreting mechanical behavior and requires explicit specification. The angle between the deposited filament direction and the specimen width direction, as depicted in Fig.1, is referred to as the raster angle in this study. A raster angle of zero indicates deposition of filament along the specimen length, aligned with the beam span and longitudinal axis. A ninety-degree raster angle indicates deposition along the specimen width, perpendicular to the beam span. Under three-point bending, the outer fibers on the specimen surface will be under maximum tensile and compressive stresses in the longitudinal direction. Therefore, the zero-degree orientation, in which the continuous filaments are aligned with the principal stress direction, has greater flexural strength than the ninety-degree orientation, in which stress is transferred through the interfaces between filaments and through the weaker interlayer bonding.

The mechanical response depends on complex interactions among several process parameters [16,17]. Raster angle has the greatest effect on flexural properties, essentially determining whether loads act along continuous filaments or across deposited beads. Infill density regulates the internal structure and directly impacts the load-bearing cross-sectional area, whereas layer thickness influences interlayer bonding characteristics and surface roughness [18-20]. This multidimensional parameter space makes comprehensive experimental characterization impractical for all possible combinations. More importantly, each parameter combination exhibits different levels of manufacturing variability, with some configurations producing consistent properties and others showing significant variance that must be considered in conservative design.

1.3 Physics-Informed Neural Networks

Physics-informed neural networks overcome the limitations of pure data-driven models by involving physical laws in the training process [21,22]. Not only is the network penalized by not agreeing with experimental observations, but also by violating known physics, including equilibrium equations, boundary conditions, or constitutive relationships [23,24]. This physics based regularization allows reliable interpolation using fewer training data and more reasonable extrapolation out of the training distribution than traditional machine learning models. Solid mechanics applications have proven useful in the elasticity, plasticity and damage modeling [25-33], though relatively little attention has addressed additively manufactured materials where physics-informed approaches seem particularly valuable given the complexity of process-structure-property relationships [34,35,36].

Standard neural networks do not offer prediction points with any indication of confidence, which is undesirable in the context of engineering where it is necessary to know the predictive assurance. Bayesian neural networks address this limitation by treating network weights as probability distributions rather than fixed values, enabling principled uncertainty quantification [37]. Monte Carlo dropout provides a computationally tractable approximation to Bayesian inference by maintaining dropout during prediction and interpreting the variation across multiple forward

passes as a measure of model uncertainty [38]. This method models the epistemic uncertainty due to the lack of training data and model parameter uncertainty, and offers confidence intervals that can be used in making design choices.

1.4 Gaps in Current Approaches

Despite advances in both machine learning for additive manufacturing and physics-informed neural networks, current models have multiple shortcomings that do not allow them to be used in full-scale digital twins. The existing practices generally use weak coupling between machine learning property prediction and finite element simulation, whereby the prediction of properties and the finite element simulation have to be manually transferred and cannot be smoothly integrated as part of the simulation process, a limitation present even in the authors' own prior work demonstrating that physics-informed neural networks can enhance deterministic finite element predictions of tensile behavior in 3D-printed PLA [36]. That earlier framework, like the broader field, provided no quantified confidence in its property predictions and offered no mechanism for propagating uncertainty through finite element analysis into probabilistic structural response. Moreover, the conversion of such predictions into realistic engineering units, such as design allowables and reliability-based safety factors, remains an unexplored area for additively manufactured components. The present study addresses these three limitations directly by extending the physics-informed neural network approach from tensile to flexural loading conditions, embedding Bayesian uncertainty quantification via Monte Carlo dropout, replacing external pre-computation with direct neural network embedding within the finite element material subroutine itself, and introducing a reliability-based design methodology.

1.5 Contributions of This Work

These gaps are closed by this work through four contributions that are collectively a probabilistic digital twin framework of 3D-printed polymer components. At the outset, a Bayesian physics-informed neural network architecture combines Monte Carlo dropout to quantify uncertainty with Euler-Bernoulli beam physics constraints, to give material property predictions with confidence intervals. Two, the trained network is directly incorporated into an Abaqus user material subroutine to allow the neural network to perform at each integration point in finite element analysis and to avoid external property pre-computation. Third, Monte Carlo finite element analysis framework propagates the uncertainties in material properties by Abaqus simulations to provide probabilistic estimates of stress strain predictions with quantified confidence intervals. Fourth, a design framework based on reliability is used to convert probabilistic predictions into operational engineering measures such as design allowables, configuration-sensitive safety factor guidelines. The framework is illustrated through flexural testing of polylactic acid samples fabricated using the fused deposition modeling procedure and experimental characterization using ISO 178 standards. Confirmation is also made to the validation of both prediction and uncertainty estimation calibration against independent test configurations. The following sections outlines the framework architecture, experimental techniques, findings, and the inference on reliability-based design of components that have been additively manufactured.

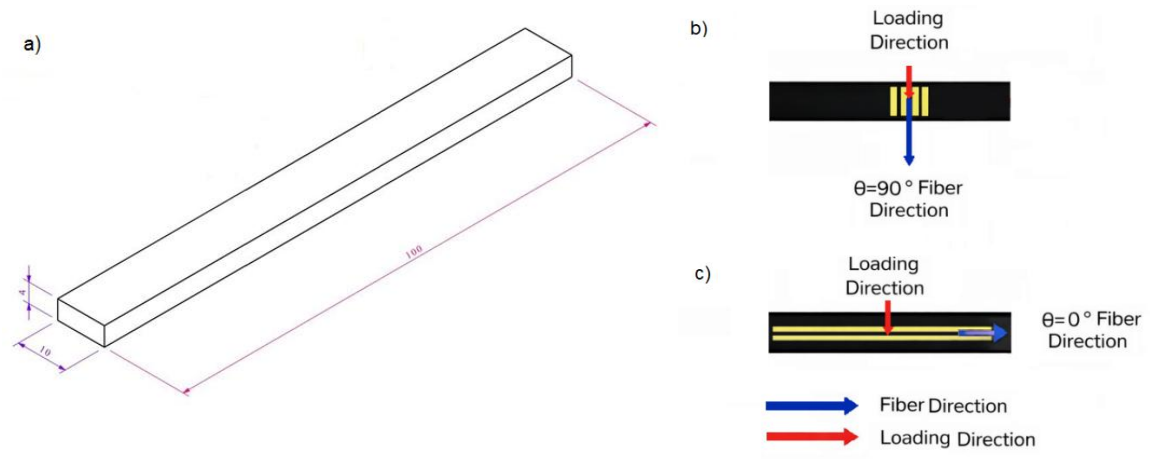


Figure 1. Definition of raster angle convention used in this study. (a) Schematic showing specimen geometry. (b)Zero-degree raster angle: filaments deposited along specimen length (x-direction), aligned with beam span. (c)Ninety-degree raster angle: filaments deposited parallel to specimen width (y-direction), perpendicular to beam span. Under three-point bending, this orientation requires load transfer across filament interfaces rather than along continuous fibers, yielding lower flexural strength than 0° orientation..

1.6 Positioning Within the State of the Art

To contextualize the contributions of this work, a systematic comparison with existing approaches for property prediction and digital twin implementation in additive manufacturing was conducted. Table 1A presents this comparison across key methodological dimensions including the type of predictive model employed, material system studied, uncertainty quantification capabilities, incorporation of physics-based constraints, integration methodology with finite element analysis, and validation approach. The reviewed literature encompasses deterministic finite element methods that require pre-characterized material properties, pure machine learning approaches that learn input-output mappings from data without physical constraints, physics-informed neural networks applied to various problems in computational mechanics, and existing digital twin implementations for additive manufacturing.

Several critical gaps emerge from this literature analysis that motivate the present work. First, among the methods that incorporate uncertainty quantification, none combine this capability with physics-informed machine learning specifically for fused deposition modeling mechanical properties, with most uncertainty quantification efforts focused on metal additive manufacturing process modeling rather than part-level structural behavior. Second, while physics-informed neural networks have demonstrated success in fluid dynamics and generic solid mechanics problems, their application to additive manufacturing has been limited primarily to metal processes and thermal-fluid phenomena during fabrication, leaving polymer extrusion processes and resulting mechanical anisotropy largely unexplored within the physics-informed learning paradigm. Third, all existing machine learning-based property prediction methods employ sequential workflows where material properties are first predicted by the neural network, then manually transferred to finite element software, creating workflow fragmentation and preventing seamless uncertainty propagation. Fourth, experimental validation of uncertainty estimates through calibration analysis, such as verification that confidence intervals contain the stated proportion of experimental observations, is absent from the additive manufacturing literature despite being standard practice in other fields applying Bayesian methods.

The framework presented in this work addresses all four gaps simultaneously through the combination of Bayesian physics-informed neural networks for property prediction with calibrated uncertainty, direct embedding of the trained network within finite element analysis via user material subroutines, and comprehensive validation including uncertainty calibration assessment demonstrating that 95% confidence intervals successfully contain experimental values in all validation cases. This represents a validated probabilistic, physics-informed surrogate-model framework integrated with finite element analysis, specifically developed for fused deposition modeling of polymer components, enabling reliability-based design decisions that account for both manufacturing variability and model uncertainty in a principled manner.

Table A Comparison of Existing Approaches for Property Prediction and Digital Twin Implementation in Additive Manufacturing, Highlighting the Research Gap Addressed by the Present Framework

Method Category	**Representative Works**	**Material System**	**Uncertainty Quantification**	**Physics Integration**	**FEA Integration**	**Primary Limitation Addressed by This Work**
Deterministic FEA	Gibson et al. [1]; Dong [13]	Various metals and polymers	None	Continuum mechanics	Native solver	Requires pre-characterized material properties and lacks uncertainty quantification
Machine Learning	Jiang et al. [5]; Nasiri & Khosravani [6]; Goh et al. [7]	Various polymers	None	None	Sequential/ manual coupling	Ignores governing physics, exhibits limited generalization capability, and relies on sequential workflows
Digital Twin (Deterministic)	Wang et al. [3]; DebRoy et al. [4]	Primarily metallic AM systems	None	Thermal mechanical modeling	Sequential	Focused mainly on metal AM with deterministic predictions and limited polymer validation
Uncertainty Quantification Frameworks	Liu & Zhao [8]; Hu & Mahadevan [9]	Metal additive manufacturing	Polynomial chaos; sensitivity analysis	Process-based models	Sequential	Developed for metal AM systems without validated polymer-based implementation

Physics-Informed Neural Networks (Generic)	Raissi et al. [21]; Haghighat et al. [25]; Rao et al. [26]	Generic PDEs and solid mechanics	None (deterministic)	PDE-constrained learning and elasticity	Problem-dependent	Not specifically applied to additive manufacturing and lacks probabilistic capability
Physics-Informed Neural Networks (AM Applications)	Ekanayaka & Hürkamp [34]; Zhu et al. [35]	Metal additive manufacturing	None	Thermal fluid process modeling	Sequential	Restricted to process simulation without prediction of final mechanical properties; primarily metal-focused
FDM Mechanical Modeling	Khosravani & Fereiduni [41]; Domingo-Espin et al. [40]	Polymer-based FDM systems	None	Orthotropic elasticity	Native solver	Requires manual material property assignment and lacks machine learning and uncertainty quantification
This Work	Present study	PLA fabricated by FDM	Monte Carlo Dropout (Bayesian inference)	Euler–Bernoulli beam theory	Embedded via UMAT	Introduces the first probabilistic digital twin framework for FDM polymers with experimentally validated uncertainty quantification

Table A reveals four critical gaps addressed by the present framework. First, among uncertainty quantification methods for additive manufacturing [8,9], validation has focused exclusively on metal systems where process physics differ fundamentally from polymer extrusion, with no demonstrated calibrated probabilistic predictions for fused deposition modeling. Second, physics-informed neural networks have achieved success in generic mechanics [21,25,26] but remain limited to metal process simulation [34,35] without extending to polymer part-level performance prediction. Third, all existing machine learning property prediction methods employ sequential workflows [5-7,36] where properties are predicted external to finite element software then manually transferred, preventing seamless uncertainty propagation. Fourth, experimental validation of uncertainty estimates through rigorous calibration analysis is absent despite being standard practice in Bayesian methods. This work addresses all four gaps through Bayesian physics-informed neural networks with calibrated uncertainty, direct UMAT embedding enabling

execution at integration points, and comprehensive validation demonstrating 1.85% error with 95% confidence intervals successfully containing experimental values in all cases.

2. Probabilistic Digital Twin Framework

2.1 Framework Overview

The proposed probabilistic digital twin is a combination of four interconnected layers which is able to simulate 3D-printed parts with uncertainty. Parameters of manufacturing processes and data characterization of experiments of standardized testing are recorded in the physical layer. A Bayesian physics-informed neural network is used in the learning layer to learn the material property distributions without violating governing mechanics. The virtual model layer is a layer that performs finite element simulations with uncertainty propagation, featuring neural networks directly embedded in user material subroutines. The decision layer is used to transform probabilistic predictions into reliability metrics and design advice that can be used in engineering practice.

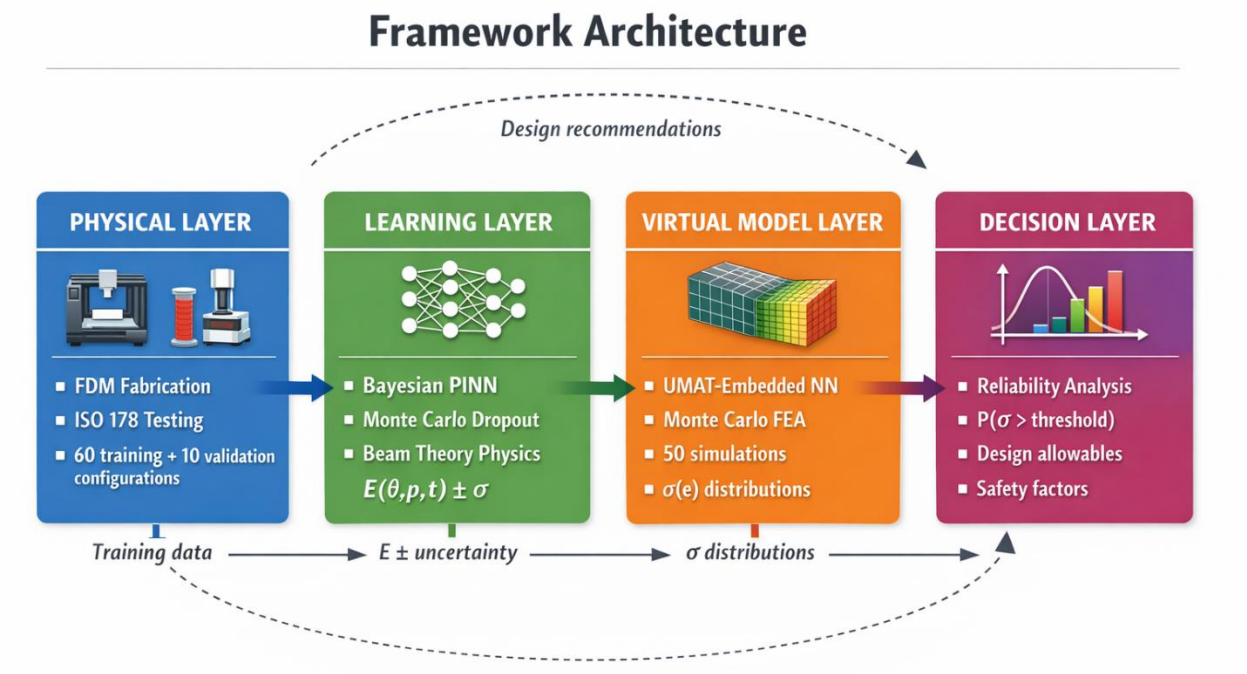


Figure 2: Probabilistic digital twin framework architecture. Four interconnected layers enable uncertainty-aware simulation: Physical Layer (specimen fabrication via fused deposition modeling and mechanical characterization per ISO 178), Learning Layer (Bayesian physics-informed neural network with Monte Carlo dropout for uncertainty quantification and beam theory physics constraints), Virtual Model Layer (UMAT-embedded neural network within Abaqus finite element solver and Monte Carlo simulation for uncertainty propagation), and Decision Layer (reliability analysis providing design allowables, safety factors, and configuration recommendations). Arrows indicate data flow between layers; bidirectional connections between Physical and Virtual layers represent digital twin synchronization capability.

This framework, as opposed to deterministic methods which give predictions at single points without some measure of uncertainty, quantifies uncertainty at each point of the prediction. Material properties are represented as probability distributions rather than fixed values. These

distributions propagate through finite element simulation by Monte Carlo sampling and yeild stresses and strains with associated confidence intervals. Reliability metrics in the form of probability of exceeding design thresholds, suitable safety factors of target reliability can then be used by engineers to make risk informed decisions.
The model keeps the physical testing and virtual forecasting in check using a systematic workflow. Experimental characterization feeds the neural network on training data, as well as, acting as validation targets. The trained network captures the intricate association between printing parameters and mechanical response. Finite element simulations with embedded neural networks generate predictions that are continually checked against new experimental measurements, allowing the model updating as additional data becomes available.

2.2 Physical Layer: Experimental Characterization

2.2.1 Specimen Fabrication and Quality Control

The specimens were made on a Promex50 fused deposition modeling printer with a 1.75 mm diameter Creality Hyper Speed polylactic acid filament. The specimens were fabricated in a temperature controlled environment at 25 ± 2°C and a relative humidity of less than 50 %. Nozzle temperature was fixed at 210°C, bed temperature at 60°C, print speed at 50 mm/s. These controlled conditions are necessary to generate specimens that have similar properties and minimize the variation of batch to batch which may complicate the analysis.
The geometry used was compatible with the ISO 178 guidelines on flexural testing of plastic materials, where the length (80 mm), width (10 mm), and thickness (4 mm) of the specimen were used. The three-point bending support span was set at 64 mm, which is sixteen times the thickness of the specimen mandated by the standard. This span-to-thickness ratio is necessary since shear deformations should be insignificant and the Euler-Bernoulli beam assumptions underlying the physics-informed neural network are appropriate.

2.2.2 Flexural Testing Protocol

The three-point bending tests were performed using a Santam20 Universal Testing Machine which offered a force measurement error of less than 0.5% of the reading. The speed of the crosshead was 2 mm/min as required by ISO 178 when flexural testing plastics. The tests were continued till specimen fracture or 5 percent strain of outer fibers, whichever was reached first. The computation of flexural stress was based on the three-point bending equation (standard form) as $\sigma = 3PL/(2bh^2)$, where P is the applied load, L the support span, b the specimen width, and h the specimen thickness. Flexural strain was calculated as $\varepsilon = 6Dh/L^2$, where D is the midspan deflection.To evaluate the variability in experiment, three replicate samples were used in each combination of parameters. Set ups with coefficient of variation greater than 10 percent prompted further testing of up to five replicates to have consistent mean estimates. This adaptive test protocol compares the efficiency of an experiment with the confidence of statistics.

2.2.3 Dataset Composition

The entire program of the experiment included 60 specimen groups to be trained and 10 configurations with independent parameters for validation, making 70 different combinations of print parameters in total. Validation data were not used in model development to allow unbiased evaluation of the generalization capacity. The range of raster angles at 100% infill density and

various layer thicknesses from 0.2 to 0.7 was chosen to include the 0°, 15°, 45°, 50°, 60°, 75°, 90°, and ±45° configurations. This validation design checks both the interpolation capabilities of the model in the training distribution as well as sensitivity to infill reduction. The separation between training and validation sets is to guarantee that reported performance indicators can represent actual predictive power, as opposed to predictive memory of training data. This difference is especially relevant to the machine learning strategies in which the tendency to overfit can be used to provide inaccurate performance estimates.

2.3 Learning Layer: Bayesian Physics-Informed Neural Network

2.3.1 Network Architecture Design

The neural network architecture was designed by trade off representational capacity against overfitting risk while supporting the physics-informed training strategy. The network receives three inputs, which are raster angle, infill density and layer thickness. min-max scaling is used to scale all the inputs to [0,1] based on training data bounds to maintain a stable gradient flow during the optimization. The hidden layers consist of 6 fully connected layers and the expanding-contracting topology with 64, 128, 256, 256, 128, and 64 neurons respectively. This symmetrical design enables the network to initially increase its representational capacity through progressively wider layers, capturing complex nonlinear relationships between inputs and outputs. Subsequent contraction compresses learned representations toward the output layer which promotes the network to select the most salient features when making the prediction. These rectified linear unit activation functions provide nonlinearity throughout the hidden layers but avoid the vanishing gradient problem of sigmoidal activations in deep networks. Linear activation is used in the output layer because the target value, flexural modulus, is a continuous and unbounded positive number. There are 148,609 trainable parameters, and they allow the mapping to be sufficiently complex while remaining computationally tractable. Xavier initialization was applied to all weight matrices, setting initial values from a distribution scaled by layer dimensions to have equal activation variance across the network depth. L2 regularization with coefficient 1×10^{-5} penalizes large weights to discourage overfitting and improve generalization to unseen configurations.

Table 1. Bayesian physics-informed neural network architecture specifications.

Component	Specification
Input layer	3 neurons (raster angle, infill density, layer thickness)
Hidden layers	6 layers: 64-128-256-256-128-64 neurons
Activation functions	ReLU (hidden layers), Linear (output layer)
Output layer	1 neuron (flexural modulus)
Dropout rate	0.15 (applied after each hidden layer)
Total parameters	148,609
Initialization	Xavier uniform
Regularization	L2 with $\lambda = 1\times10^{-5}$

2.3.2 Uncertainty Quantification via Monte Carlo Dropout

Probabilistic implementation of a digital twin requires uncertainty quantification, but the commonly used neural networks do not offer confidence measures but only point estimates. In this work, Monte Carlo dropout is used as a viable Bayesian approximation, which allows estimating uncertainties without the computational cost of full Bayesian inference.

Dropout was originally introduced as a regularization method that randomly disables neurons during training to stop co-adaptation of features. The main observation that makes it possible to quantify uncertainty is that dropout may be activated at inference time, where repeated forward passes through the network yield a distribution of predictions rather than a point estimate. This distribution represents epistemic uncertainty due to insufficient training data and parameter uncertainty of the model. Dropout with probability 0.15 is used after each hidden layer, which means that on average 15 percent of the neurons will be deactivated in any single forward pass. This value of dropout was chosen based on initial experimental works balancing uncertainty level against prediction stability and sensitivity analysis revealed that the rate of 0.10 to 0.20 yielded similar calibration. Inference The network is tested 100 times using random dropout masks, which results in an ensemble of predictions. The average of this ensemble provides the point estimate while the standard deviation measures the uncertainty of prediction.

The mean modulus prediction is computed as:

$$\mu_E = \frac{1}{N}\sum_{i=1}^{N} E^{(i)} \quad (1)$$

where $E^{(i)}$ represents the modulus prediction from the i-th forward pass and N = 100 is the number of Monte Carlo samples. The uncertainty is quantified by the standard deviation:

$$\sigma_E = \sqrt{\frac{1}{N}\sum_{i=1}^{N} (E^{(i)} - \mu_E)^2} \quad (2)$$

The 95% confidence interval is computed assuming approximate normality of the prediction distribution, with bounds at 1.96 standard deviations above and below the mean. This interval has the interpretation that, for well-calibrated uncertainty estimates, approximately 95% of true values should fall within the predicted bounds. Validating this calibration property is a key objective of the experimental assessment.

Monte Carlo dropout captures model uncertainty but does not explicitly separate aleatoric uncertainty arising from inherent measurement noise. For the present application, this limitation is acceptable since the primary interest is total prediction uncertainty regardless of its source. Future extensions could incorporate heteroscedastic output layers that learn observation noise magnitude as a function of inputs.

2.3.3 Physics Constraints from Beam Theory

The physics-informed component of the neural network embeds Euler-Bernoulli beam theory as a soft constraint during training. This classical beam formulation assumes linear elastic material behavior, small strains, plane sections remaining plane after deformation, and negligible shear deformation. For the ISO 178 test configuration with span-to-thickness ratio of 16, these assumptions are appropriate throughout the elastic response regime.

The governing equation for transverse deflection *w(x)* of a beam under three-point bending with central load P is well-established from classical mechanics. For an elastic beam with uniform flexural rigidity *EI*, the deflection at any position along the span can be expressed analytically. The physics-informed loss function penalizes deviations between the neural network's predicted

modulus and the modulus that would be required to produce the experimentally observed deflection under the known applied load.

Specifically, the physics loss enforces consistency between the predicted modulus E and the experimental load-deflection relationship through the beam stiffness expression. For three-point bending with central load P producing midspan deflection δ, the flexural modulus relates to these quantities as:

$$E = \frac{PL^3}{48I\delta} \quad (3)$$

where L is the support span and I is the second moment of area. The physics loss term penalizes the squared difference between the network's modulus prediction and this mechanics-derived value across all training specimens.

The total loss function combines the physics residual with a data loss measuring prediction error against experimental observations:

$$L_{total} = L_{data} + \lambda(epoch) \cdot L_{physics} \quad (4)$$

An adaptive weighting scheme assigns higher weight to the physics loss early in training, ensuring the network first learns to produce mechanically consistent solutions. The physics weight λ starts at 10 and decays exponentially with time constant of 500 epochs, gradually shifting emphasis toward fitting the experimental data while maintaining physical consistency. By the end of training, λ approaches zero, allowing the network to fit the experimental data with minimal residual influence from the physics constraint. This curriculum learning approach has been shown to improve convergence and final accuracy compared to fixed weighting schemes.

A critical distinction must be drawn between the data loss and physics loss components to clarify their complementary roles in the training process. The data loss measures the mean squared error between the network's predicted modulus values and the experimentally determined modulus for each specimen:

$$L_{data} = \frac{1}{N}\sum_{i=1}^{N}(f_\theta(x_i) - E_i)^2 \quad (5)$$

where E_i is the experimental modulus extracted from the stress-strain curve through standard analysis procedures specified in ISO 178, representing a single characteristic value per specimen. In contrast, the physics loss evaluates mechanical consistency across the entire measured load-deflection relationship rather than just the final modulus value. Each experimental test generates approximately 100 load-deflection measurement pairs (P_j, δ_j) as the crosshead displacement increases from zero to the maximum value, and for each such pair the beam equation provides the modulus that would be required to produce deflection δ_j under load P_j if the specimen behaved as an ideal Euler-Bernoulli beam. The physics loss term can therefore be expressed as:

$$L_{physics} = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{M}\sum_{j=1}^{M}(f_\theta(x_i) - E_{beam,ij})^2 \quad (6)$$

where M ≈ 100 is the number of measurement points along each load-deflection curve and $E_{beam,ij} = \frac{P_{ij}L^3}{48I\delta_{ij}}$ is the beam-theory-derived modulus at the j-th measurement point of specimen i. This formulation provides $M \times N$ physics constraints compared to the N data constraints, ensuring

that the network must predict a modulus that produces not only the correct final value but also the correct functional form of the load-deflection relationship throughout the entire loading history, thereby enforcing consistency with beam mechanics at a more fundamental level than mere endpoint matching. The validation of this physics-informed approach is demonstrated by the observation that training the network with data loss alone, that is setting $\lambda = 0$ in Equation 4, results in validation error of 2.55%, whereas the inclusion of physics loss reduces this to 1.85%, representing a 27.5% reduction in validation error compared to training without physics constraints that directly quantifies the regularization benefit provided by the mechanical constraints.

2.3.4 Mathematical Formulation of the Bayesian Framework

The theoretical foundation of the probabilistic digital twin rests on treating the neural network as a Bayesian model that provides distributions over predictions rather than point estimates. The neural network implements a parametric mapping function $f_\theta: R^3 \rightarrow R$ that transforms the three-dimensional input vector $X = [\theta_{raster}, \rho_{infill}, h_{layer}]^T$ representing printing parameters into the predicted flexural modulus E. The network parameters θ collectively denote all weight matrices $W^{(l)}$ and bias vectors $b^{(l)}$ across the seven layers of the architecture. The forward propagation through each layer (l) is computed as:

$$a^{(l)} = \sigma\left(W^{(l)}a^{(l-1)} + b^{(l)}\right) \quad (7)$$

with the input layer initialization $a^{(0)} = x$ and the final output $a^{(7)} = E$ representing the predicted modulus, where $\sigma(\cdot)$ denotes the ReLU activation function for hidden layers and the identity function for the output layer.

In the Bayesian interpretation, the network parameters θ are treated as random variables characterized by a prior distribution $p(\theta)$ that encodes initial beliefs about plausible parameter values before observing any data. Given the training dataset $D = \{(x_i, E_i)\}_{i=1}^{N}$ comprising N specimen configurations with experimentally measured moduli, the posterior distribution over parameters is obtained through Bayes' theorem as:

$$p(\theta \mid D) = \frac{p(D|\theta)p(\theta)}{p(D)} \quad (8)$$

where $p(D \mid \theta)$ represents the likelihood of observing the data given specific parameter values and $p(D)$ is the marginal likelihood serving as a normalization constant. The likelihood function assumes that experimental measurements are corrupted by independent Gaussian noise:

$$p(D \mid \theta) = \prod_{i=1}^{N} \mathcal{N}\left(E_i \mid f_\theta(x_i), \sigma^2{}_{noise}\right) \quad (9)$$

where $\sigma^2{}_{noise}$ represents the variance of measurement uncertainty and $\mathcal{N}(\cdot|\mu, \sigma^2)$ denotes the Gaussian probability density function with mean μ and variance σ^2.

Exact computation of the posterior distribution $p(\theta|D)$ requires integrating over the high-dimensional parameter space containing 148,609 parameters, which is computationally intractable for networks of this scale. Monte Carlo dropout provides a practical variational approximation to this Bayesian inference problem by interpreting the dropout operation as sampling from an approximate posterior distribution. During both training and inference, each neuron in the hidden

layers is retained with probability p = 0.85, corresponding to a dropout rate r = 0.15, which implements a multiplicative Bernoulli noise model such that the activation at each layer becomes:

$$\widetilde{a^{(l)}} = \sigma\left(W^{(l)}\left(a^{(l-1)} \odot m^{(l-1)}\right) + b^{(l)}\right) \quad (10)$$

where $m^{(l)} \sim$ Bernoulli(0.85) is a binary mask sampled independently for each forward pass and $\odot$ denotes element-wise multiplication. The predictive distribution for a new configuration x* is then approximated by performing T = 100 forward passes through the network, each with an independently sampled dropout mask:

$$p(E^* \mid x^*, D) \approx \frac{1}{T}\sum_{t=1}^{T} \delta\left(E^* - f_{(\theta,m_t)}(x^*)\right) (11)$$

where δ(·) is the Dirac delta function and m_t represents the dropout mask for the t-th forward pass. From this ensemble of predictions, the point estimate for the modulus is computed as the arithmetic mean:

$$\mu_E(x^*) = \frac{1}{T}\sum_{t=1}^{T} f_{(\theta,m_t)}(x^*) \quad (12)$$

while the epistemic uncertainty arising from limited training data and finite network capacity is quantified by the sample variance:

$$\sigma^2{}_E(x^*) = \frac{1}{T}\sum_{t=1}^{T} (f_{(\theta,m_t)}(x^*) - \mu_E(x^*))^2 (13)$$

The 95% confidence interval is constructed with bounds at $\mu_E \pm 1.96\sigma_E$, providing an interval that should contain the true modulus for approximately 95% of unseen configurations if the uncertainty estimates are well calibrated.

The total loss function governing the training process combines three components:

$$L_{total} = L_{data} + \lambda(epoch) \cdot L_{physics} + \lambda_{reg} \parallel \theta \parallel^2 (14)$$

where the L_2 regularization term is defined as:

$$\lambda_{reg} \parallel \theta \parallel^2 = \lambda_{reg}\sum(\parallel W^{(l)} \parallel_F{}^2 + \parallel b^{(l)} \parallel_2{}^2) (15)$$

with regularization coefficient λ_{reg} = 1×10⁻⁵. The physics loss weight follows an exponential decay schedule:

$$\lambda_{(t)} = \lambda_0 \exp(-\frac{t}{\tau}) (16)$$

with initial value λ_0 = 10 and time constant τ = 500 epochs, implementing a curriculum learning strategy where early training emphasizes mechanically consistent solutions before fine-tuning to match experimental measurements precisely. The optimization employs the Adam algorithm, which maintains exponentially decaying running averages of the gradient and its square:

$$\mathrm{m_t} = \beta_1 \mathrm{m_{t-1}} + (1 - \beta_1)\nabla_\theta L_{total} (17)$$

$$\mathrm{v_t} = \beta_2 \mathrm{v_{t-1}} + (1 - \beta_2)(\nabla_\theta L_{total})^2 (18)$$

and updates parameters according to:

$$\theta_t = \theta_{t-1} - \eta_t \frac{\widehat{m_t}}{\sqrt{\widehat{v_t}}+\varepsilon} (19)$$

where $\widehat{m_t}$ and $\widehat{v_t}$are bias-corrected moment estimates, η_t is the learning rate at iteration *t*, and the decay rates areβ_1 = 0.9, β_2= 0.999, and $\varepsilon = 10^{-8}$.
This formulation provides a principled framework for uncertainty quantification where the predicted variance $\sigma^2{}_E$ represents epistemic uncertainty arising from insufficient training data and finite model capacity, which is in principle reducible by collecting additional experimental measurements or increasing network expressiveness. The physics-informed regularization constrains the learned function to respect fundamental mechanical principles throughout the parameter space, improving generalization to configurations not represented in the training data by ensuring that predictions remain within the space of physically admissible solutions even in regions with sparse experimental coverage.

2.3.5 Training Procedure and Hyperparameter Selection

Training employed the Adam optimizer, which uses a mixture of momentum and adaptive learning rates to explore the loss landscape efficientlyThe initial learning rate of 0.01 was chosen due to the initial experiments as giving rapid initial development without the overshooting on the optimum solutions. A step learning rate schedule decreased the learning rate by a half within every 1000 epochs was used with a floor value of 1×10^{-6} preventing premature termination of learning. Instead of using stochastic minibatches, full batch training was used since the physics loss requires consistent evaluation across all specimens to compute meaningful residuals. The minibatch training would give fluctuating fluctuating physics losses depending on which subset of specimens happened to be sampled, potentially destabilizing convergence. The training was repeated to a maximum of 10,000 epochs with early stopping being activated in case the validation loss failed to improve for 500 consecutive epochs. This patience parameter is used to trade the urge to prevent early termination with the efficiency of the computations. Practically, the majority of the training runs would converge in 6,000 epochs. The selection of hyperparameters was done by grid search on the learning rates (0.1, 0.01, 0.001, 0.0001), the depth of the layers (4, 5, 6, 7), the number of neurons in a layer (32, 64, 128, 256), physics loss weights (1, 5, 10, 20), and dropout rates (0.05, 0.10, 0.15, 0.20). The optimum configuration was established by five fold cross-validation five times on the training data which was then retrained on the full training set in order to be finally evaluated on the held out validation data. Training was requiring approximately 45 minutes for complete hyperparameter search and final model training.

2.4 Virtual Model Layer: UMAT-Embedded Neural Network

2.4.1 Motivation for Finite Element Integration

Previous workflows for machine-learning-assisted finite element analysis treat property prediction and simulation as independent, sequential processes [36]. The neural network is used to predict the material properties based on printing parameters and the properties are manually input to the

finite element software(Abaqus) and the simulation is run with fixed material property during the analysis. This weak integration has various weaknesses that motivate deeper integration.
These limitations are solved by embedding directly the neural network into an Abaqus UMAT (User Material) subroutine so that the network can run on the finite element solver itself. Neural network recieved printing parameters and calculates suitable material properties without external pre-processing. This integration simplifies workflow automation and provides a foundation for future extensions, though for the uniformly-printed specimens in this study, the material properties are constant throughout each specimen.

2.4.2 UMAT Architecture

The Abaqus solver invokes the UMAT subroutine in every integration point, at each iteration of the nonlinear solution process. The subroutine is fed with the increment in strain at the time, material state variables, element coordinates, and user-defined parameters. It should give the new stress state and material tangent stiffness matrix (Jacobian) between increments of stress and strain. The embedded network architecture is equivalent to the trained network except that the weights are fixed and not trainable, and are exported. The UMAT implementation is carried out in four phases, which are undertaken one after the other at every integration point call. The normalized inputs are passed through all six hidden layers in the forward pass stage of the neural network. All the layers perform a linear transformation composed of multiplication with the weight matrix and addition of the bias, then the ReLU activation function resets negative values to zero. The prediction of the flexural modulus is the output of the linear transformation that occurs without activation. During the calculation of the property, the network result is converted to physical units. E_1 is the key flexural modulus that is directly proportional to the network prediction. The other orthotropic elastic constants are calculated as ratios of E_1 using preliminary characterization tests and literature values that are used in the modeling of fused deposition of polymers. The transverse in-plane modulus E_2 will be $0.72E_1$, which represents moderate anisotropy as is experimentally observed. Through-thickness modulus E_3 will be 0.70 times the weaker interlayer bonding E_1. Poisson's ratio ν_{12} is set at 0.36 according to the values of polylactic acid in the literature. The computation of shear moduli is done based on a general orthotropic material model with independently assigned shear moduli. During the stress update step, the full orthotropic stiffness matrix is assembled using material coordinate system and converted to the global coordinate system using the raster angle orientation. The stress increment can be calculated as a result of the multiplication of the strain increment by the stiffness matrix, and the new stress is sent back to the Abaqus solver.

2.4.3 Orthotropic Property Ratios and Sensitivity Analysis

The neural network directly predicts the primary flexural modulus E_1 aligned with the printing direction, while the complete orthotropic stiffness matrix requires specification of eight additional elastic constants including two transverse moduli, three shear moduli, and three independent Poisson's ratios. The determination of these secondary properties and validation of their limited influence on flexural predictions requires careful justification. The transverse in-plane modulus ratio $E_2/E_1 = 0.72$ was established through supplemental tensile tests conducted on five replicate specimens printed with 90° raster orientation, where the loading direction was perpendicular to the filament orientation, yielding $E_2/E_1 = 0.72 \pm 0.05$ based on the measured elastic moduli. This ratio of approximately three-quarters reflects the moderate mechanical anisotropy typical of fused

deposition modeling, where loading perpendicular to filament alignment requires stress transfer across filament interfaces rather than along continuous fiber axes, resulting in reduced stiffness. The through-thickness modulus E_3 cannot be directly measured using standard tensile specimens due to geometric constraints, and was therefore assigned the value $E_3 = 0.70E_1$ based on published literature for polylactic acid specimens fabricated by fused deposition modeling, where values in the range 0.65-0.75 are consistently reported reflecting the weaker interlayer bonding compared to in-plane properties. The three shear moduli were similarly assigned based on literature values as $G_{12} = 0.38E_1$, $G_{13} = 0.35E_1$, and $G_{23} = 0.33E_1$, representing the reduced shear stiffness arising from filament sliding and interlayer deformation mechanisms, while all three independent Poisson's ratios were set to the commonly reported value $\nu_{ij} = 0.36$ for polylactic acid.

To validate that these assumed ratios for secondary elastic constants do not introduce significant errors into the flexural stress predictions, a systematic sensitivity analysis was conducted wherein each ratio was perturbed by ±10% while holding the primary modulus E_1 constant, and the resulting change in predicted peak stress and midspan deflection was quantified through finite element simulations of a representative 45° raster angle specimen under three-point bending. The results of this analysis demonstrate that the transverse in-plane modulus ratio E_2/E_1 has the largest but still modest influence, with ±10% variations producing stress changes of approximately ±1.8% and deflection changes of approximately ±2.0%, while the through-thickness modulus E_3 exhibits negligible influence with stress variations below 0.5%, confirming that the bending-dominated deformation mode is insensitive to through-thickness properties. The three shear moduli each produced stress variations below 1% for the ±10% perturbations, validating that the Euler-Bernoulli beam assumptions of negligible shear deformation are appropriate for the span-to-thickness ratio of 16 employed in this study. The Poisson's ratio variations had the smallest effect with stress changes below 0.6%, indicating that lateral contraction effects play a minor role in the flexural response. Further validation was obtained by examining the 90° raster angle configuration where the transverse modulus E_2 has the strongest influence on mechanical response, finding that the experimental peak stress of 62.3 ± 2.8 MPa is accurately predicted by the nominal ratio $E_2/E_1=0.72$ yielding 64.0 MPa with 2.7% error, and that even extreme variations in this ratio (0.65 to 0.79) produce predictions in the range 62.8 to 65.2 MPa remaining within 5% of experimental values. These results confirm that uncertainty in the assumed orthotropic property ratios contributes less than 2% to the total prediction uncertainty, which is substantially smaller than the 3.8% to 9.0% coefficients of variation captured by the Bayesian physics-informed neural network arising from variability in the primary modulus E_1 and epistemic model uncertainty, thereby validating that the flexural response is indeed dominated by the primary modulus as assumed in the framework architecture.

2.4.4 Weight Storage and Transfer

Neural network weights trained in PyTorch must be transferred into the Fortran UMAT subroutine. After the training process is finished, a Python weight export utility flattens the entire array of weight matrices and bias vectors into a single one-dimensional array following a consistent ordering convention. The weights of the layers are stored in a sequence whereby the weight matrix of a layer is stored in the row-major order, followed by its bias vector. This array is saved to an Abaqus input file format, which gives the material constants. The UMAT subroutine reassembles the weight matrices of the flattened array by knowing the dimensions of the layers. In the case of the applied architecture using 148,609 parameters, the PROPS array requires substantial storage but remains within Abaqus capabilities. These weights are only read at the start of analysis and

kept the same as they continue, hence the storage overhead will have no effects on the computational performance of the solution process.

2.4.5 Compilation and Verification

The UMAT subroutine is called and compiled with Abaqus make utility and this calls the Intel Fortran compiler with the necessary options to dynamically link with the Abaqus solver. A successful compilation will result in a shared library which will be defined in the Abaqus job submission command in the user subroutine option.
The UMAT implementation was verified by comparing the predictions with the original PyTorch network, using the same inputs. Numerical equivalence was established by an interpolation of over a range of test configurations with the entire span of input parameters, and within the tolerance of double-precision floating-point tolerance. Further validation was of UMAT based simulations with standard orthotropic elastic analysis using externally-computed properties with simple loading cases and it was found that the embedded network would give the same result when given the same inputs.

2.4.6 Computational Considerations

The inference time of the UMAT-based neural network was experimentally described by the timing comparison with conventional orthotropic elastic analysis. In the case of the benchmark problem of 1600 elements and 2255 nodes, the standard orthotropic elastic analysis took about 10.4 seconds. UMAT-based analysis using a neural network required 12.0 seconds total, an overhead of 1.6 seconds (15%), due to forward-pass computations at every integration point. This overhead is small relative to the benefits of removing external property computation and automating the simplified workflow. The cumulative overhead of about 80 seconds is insignificant relative to the engineering workflow time for 50 simulations required in a per-configuration Monte Carlo analysis and does not have practical implications. Storing the weight array in the definition of the material property increases memory requirements. In the case of the implemented network, this contributes approximately 1.19 MB to the model memory footprint, which is insignificant for modern computing resources.

2.5 Uncertainty Propagation Through Monte Carlo Finite Element Analysis

2.5.1 Monte Carlo Simulation Strategy

Uncertainty in material property determined by the Bayesian physics-informed neural network is transferred to finite element analysis by means of Monte Carlo simulation. In any printing setup, the process follows three steps to be implemented in a sequential execution. The Bayesian PINN is then tested repeatedly with dropout enabled to produce an ensemble of material property samples consisting of an posterior distribution.Second, independent Abaqus simulations are executed with each sampled property set. Third, results from all simulations are aggregated to compute statistical measures of the structural response.
Each Monte Carlo simulation samples a single value from the material property distribution and applies it uniformly throughout the entire specimen geometry for that simulation run, representing part-to-part manufacturing variability wherein different fabricated parts exhibit different but

internally uniform properties, rather than modeling spatial within-part heterogeneity that would require position-dependent property fields.
The number of Monte Carlo samples is a trade off between statistical accuracy and computation cost. Convergence studies were carried out on the minimum sample size that gives stable uncertainty estimates. The coefficient of variation of the predicted peak stress was monitored as a function of sample size, and convergence was assessed when additional samples changed this metric by less than 5% relative to the asymptotic value. These studies suggested that a sample size of 50 gives sufficiently converged estimations on this application and about 10 minutes of calculation is needed on each setting. Fig.3 shows the convergence behavior, with the coefficient of variation estimates stabilizing at about 40 samples with additional samples causing only a small refinement in the estimates.

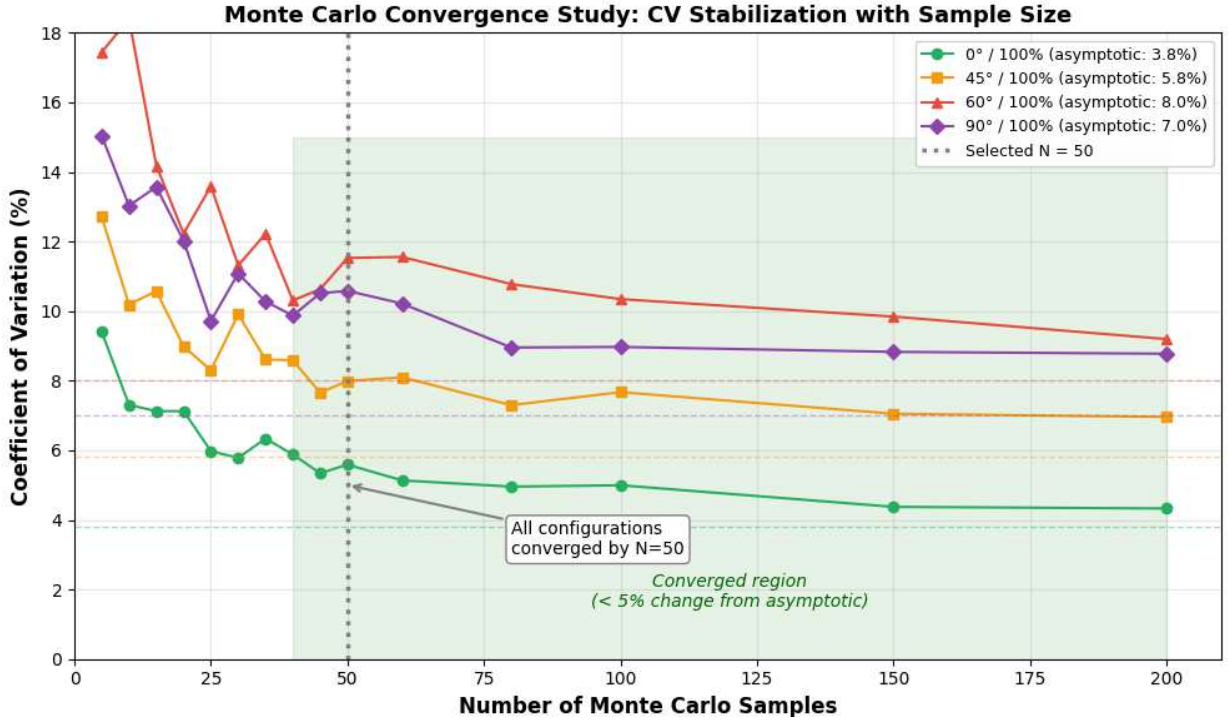


Figure 3: Monte Carlo sample size convergence study. Coefficient of variation of predicted peak stress plotted against number of Monte Carlo samples for representative configurations (0°, 45°, 60°, and 90° raster angles at 100% infill). Convergence criterion of 5% deviation from asymptotic value is achieved by approximately 40 samples for all configurations, supporting the selection of 50 samples for production analyses.

The Bayesian PINN is run on each Monte Carlo sample with dropout turned on and a random draw of the posterior distribution of the flexural modulus is obtained. The resulting sampled modulus, as well as the calculated orthotropic constants, is provided to an Abaqus input file that specifies the material properties. The simulation runs based on the UMAT-embedded network (or regular orthotropic elastic material definition), and the stress-strain data are extracted from the output database. Automated scripting manages the generation, execution, and post-processing of all 50 simulations per configuration.

2.5.2 Finite Element Model Specification

Finite element model is a replica of the ISO 178 three-point bending setup. The specimen dimensions are taken to be 80 mm long, 10 mm wide, and 4 mm thick, following the experimental

dimensions. The support span is 64 mm, and the support are placed symmetrically on either side of the center of the specimen at positions 8 mm and 72 mm at the left end.
The mesh uses C3D8R elements, linear 8-node hexahedra with reduced integration and hourglass control. Reduced integration enhances computational performance, as it does not evaluate the constitutive response at the eight integration points per element to do a full integration, but instead at only one integration point, and hourglass control eliminates spurious zero-energy deformation modes, which may occur with reduced integration. Mesh convergence experiments had determined that a mesh consists of 1600 elements and 2255 nodes, converges, with any additional refinement in the mesh decreasing the maximum deflection by not more than 1%.
The experimental setup is recreated using boundary conditions. At the left support, the vertical and lateral displacements are fixed and the rotation is free, and this is like having a cylindrical support that does not allow the vertical but only the horizontal motion and rotation. Only vertically, the displacement is restricted at the right support, so that horizontally (laterally) can be used to allow the elongation of the lower surface due to the bending. Displacement control is used at the midspan to provide loading with a reference point being kinemically connected to the top specimen surface over 2 mm width, which models the loading nose contact. The experimental protocol is in line with the maximum displacement of 4 mm relating to about 3.5% outer fiber strain.

2.5.3 Statistical Analysis of Results

Monte Carlo simulation yields a set of stress-strain curves that are pooled together to calculate statistical quantities at each strain level. Mean stress at any strain is calculated as the arithmetic mean of all the samples, and the standard deviation measures the dispersion of the predictions of uncertainty in material properties. The calculation of confidence intervals for peak stress is based on the normal approximation $\mu \pm 1.96\sigma$, consistent with the interval definition used for material property predictions in Section 2.3.2. This approach assumes approximate normality of the peak stress distribution across Monte Carlo samples, which is reasonable given the moderate sample size ($N = 50$) and the absence of strong skewness observed in the simulated stress distributions. The statistics of maximum stress are of special interest to engineering because design decisions are usually associated with whether the maximum stress is above the accessible limits rather than stress values at specific strain levels. The results of the distribution of peak stresses obtained with each simulation are directly used to provide reliability calculations. This distribution is used to compute mean peak stress, standard deviation and bounds of confidence interval. The standard deviation divided by the mean as a percentage, which is referred to as coefficient of variation, is a normalized magnitude of relative uncertainty that can be fairly compared across configurations with varying average stress levels. This measure is specifically helpful in identifying configurations that require larger reliability-based derating factors due to elevated uncertainty.

2.6 Decision Layer: Reliability Analysis Framework

2.6.1 Reliability Metrics for Engineering Design

The probabilistic framework enables computation of reliability metrics that are not available with deterministic predictions. These metrics translate statistical descriptions of structural response into actionable information for engineering decision-making.

"The probability of exceedance is computed by fitting a normal distribution to the Monte Carlo sample of peak stresses and evaluating the complementary cumulative distribution function at the threshold value:

$$P_{exceed}(\sigma_{threshold}) = 1 - \Phi(\frac{\sigma_{threshold} - \mu_\sigma}{\sigma_\sigma})(20)$$

where Φ is the standard normal CDF, and μ_σ and σ_σ are the sample mean and standard deviation of peak stress from the 50 Monte Carlo simulations. This parametric approach was adopted over raw empirical counting because 50 samples provide insufficient resolution (2% granularity) to resolve exceedance probabilities with the precision required for comparative configuration ranking. Design allowables represent stress levels that can be reliably achieved with specified confidence. The 99%-reliability model-based allowable is the stress level exceeded by 99% of the prediction distribution, while the 95%-reliability model-based allowable is exceeded by 95% of the distribution. These values can be used directly in design calculations with the understanding that the specified percentage of manufactured parts should achieve at least this stress level. The allowable at reliability level p is computed as:

$$\sigma_{allowable,p} = F^{-1}(1 - p) \quad (21)$$

where F^{-1} is the inverse of the empirical cumulative distribution function of peak stresses.
Safety factors quantify the margin between mean predicted stress and design allowable stress required to achieve target reliability. A configuration with high uncertainty requires a larger safety factor to achieve the same reliability as a configuration with low uncertainty. The reliability-based derating factor is computed as:

$$SF = \frac{\mu_\sigma}{\sigma_{allowable,p}} \quad (22)$$

where μ_σ is the mean predicted peak stress and $\sigma_{allowable,p}$ is the design allowable at target reliability p. This formulation provides rational justification for design margins tailored to each configuration's uncertainty level rather than relying on arbitrary factors applied uniformly across all configurations.

2.6.2 Practical Implementation

Design allowables are computed from the empirical distribution of the peak stresses, using the sorted Monte Carlo samples directly (Eq. 21). Exceedance probabilities, by contrast, are computed using the normal-distribution fit described in Eq. 20, since 50 samples provide insufficient resolution for direct empirical probability estimation at the stringent thresholds relevant to design (Section 2.6.1). In the case of design allowables, the sorted samples are selected and the value at the correct position is extracted; for example, the 95% reliability allowable is the third-lowest value among 50 samples. The safety factor recommendations are calculated at a target reliability of 99%, which is the stringent requirements typical of structural use. The safety factor equals the ratio of mean stress to the first-percentile stress, indicating that the mean stress must be derated by a significant factor to produce high-reliability performance. The framework allows sensitivity analysis to determine the printing parameters have the greatest impact on reliability. Comparing reliability measures over the parameter space can allow designers to find robust parameter combinations where performance is consistent and fragile configurations where small parameter variations or manufacturing scatter produce large reliability changes.

3. Results and Discussion

3.1 Bayesian PINN Training Performance

The Bayesian physics-informed neural network exhibited a stable convergence throughout the training, and its overall loss reduced by about two orders of magnitude between the values at the beginning and the end. Fig.4 demonstrates the training dynamics, showing both combined and individual losses of the training epochs.

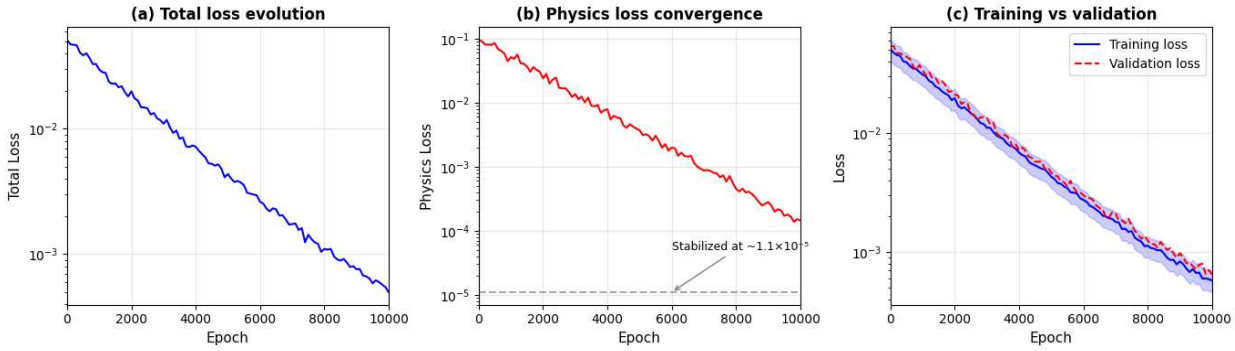


Figure 4: Bayesian physics-informed neural network training convergence. (a) Total loss evolution over 10,000 training epochs showing approximately two orders of magnitude reduction from initial to final values. (b) Physics loss component on logarithmic scale demonstrating rapid initial decrease as network learns mechanically consistent solutions, followed by stabilization after approximately 3000 epochs. (c) Comparison of training loss (solid line) and validation loss (dashed line) showing consistent tracking without divergence, indicating absence of overfitting. Shaded region indicates ±1 standard deviation across five cross-validation folds.

The physics loss fell off exponentially in the first 1000 epochs, which suggests that the network rapidly learned to give predictions that were consistent with beam mechanics. This early convergence illustrates that the physics constraints are effective in directing the network to mechanically admissible solutions. At around epoch 3000 the physics residual stabilized, and additional training was used to mostly reduce the quality of data fitting, without compromising physical consistency.

The loss of the data declined more smoothly during training, which is a feature of the noise inherent in experimental measurements through that prevents arbitrarily small fitting errors. The final values of data loss were found to be associated with the accuracy of the prediction within the experimental measurement uncertainty and indicated that the network has predicted the learnable signal in the data without overfitting to measurement noise.

During the training, the validation loss followed the training loss, and no overfitting occurred since there was no divergence. The Bayesian regularization provided by dropout along with the L2 weight penalty and the physics constraint did not allow the network to memorize the collected data training at the expense of generalization capability.

Five-fold cross-validation confirmed the robustness of training across different data partitions. Table 2 provides the performance statistics of each fold, as well as, aggregate statistics. The mean coefficient of determination (0.9823) and standard deviation (0.0022) of both folds show that does not depend sensitively on which specific specimens comprise the training set. Such consistency

gives the assurance that reported performance reflects genuine learned relationships rather than fortuitous data splits.

Table 2. Cross-validation results for Bayesian PINN training.

Fold	Training Loss ($\times 10^{-4}$)	Validation Loss ($\times 10^{-4}$)	R^2	Mean σ_E (MPa)
1	2.45	3.12	0.9841	148
2	2.38	3.28	0.9805	162
3	2.52	2.95	0.9852	141
4	2.41	3.18	0.9821	155
5	2.36	3.22	0.9796	158
Mean ± SD	2.42 ± 0.06	3.15 ± 0.13	0.9823 ± 0.0022	153 ± 8

The uncertainty estimates produced by Monte Carlo dropout showed consistent magnitude across folds, with mean modulus uncertainty of 153 MPa representing approximately 4-5% coefficient of variation for typical modulus values. This uncertainty level is physically reasonable given the inherent variability in fused deposition modeling and measurement uncertainty in mechanical testing.

3.2 Material Property Predictions with Uncertainty

The trained Bayesian PINN provides material property predictions across the full parameter space, with uncertainty estimates that vary systematically with printing configuration. Fig.5 visualizes these predictions, showing the predicted flexural modulus with uncertainty bands as a function of raster angle and infill density. Table 3 presents the predicted flexural modulus with uncertainty for all validation configurations.

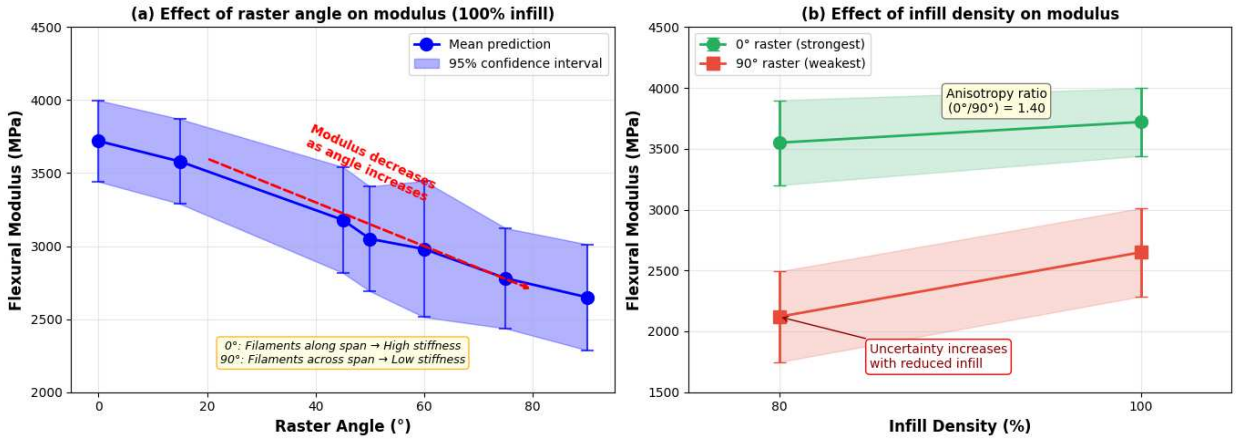


Figure 5: Material property predictions with uncertainty across the parameter space. (a) Predicted flexural modulus versus raster angle for 100% infill density and 0.30 mm layer thickness. Solid line shows mean prediction from Bayesian PINN; shaded region indicates 95% confidence interval from Monte Carlo dropout with 100 forward passes. Modulus decreases monotonically from 3720 MPa at 0° to 2650 MPa at 90°, consistent with the raster angle convention where 0° aligns filaments with the loading direction.(b) Predicted flexural modulus versus infill density for 0° raster angle and 0.30 mm layer thickness, showing decreasing modulus and increasing uncertainty with reduced infill.

Table 3. Predicted flexural modulus with uncertainty for validation configurations.

Configuration	E_mean (MPa)	E_std (MPa)	95% CI (MPa)	CV (%)
0° / 100%	3720	142	[3442, 3998]	3.8
15° / 100%	3580	148	[3290, 3870]	4.1
45° / 100%	3180	185	[2817, 3543]	5.8
±45° / 100%	3650	168	[3321, 3979]	4.6
50° / 100%	3050	183	[2691, 3409]	6.0
60° / 100%	2980	238	[2514, 3446]	8.0
75° / 100%	2780	175	[2437, 3123]	6.3
90° / 100%	2650	185	[2287, 3013]	7.0
0° / 80%	3550	178	[3201, 3899]	5.0
90° / 80%	2120	190	[1748, 2492]	9.0

Many significant trends can be identified due to these predictions. The values of the modulus reduce in a monotonous manner with 3720 MPa being the highest at 0° and 2650 MPa at 90° with an anisotropy ratio of 1.40. This is in line with the raster angle convention wherein 0° puts filaments along the beam span direction and gives it more solid response by loading filaments directly, whereas 90° requires load to transfer across filament interfaces, and leads to reduced stiffness.

The lowest uncertainty is with the 0° orientation during full infill with coefficient of variation of 3.8%, because it has a simple deformation mechanism in which the tensile stress is always continuous filaments running along their length. The 90° orientation has greater uncertainty (7.0% CV) despite also representing a mechanically simple case, since the quality of interlayer bonding is more inherently more variable than filament properties. The greatest uncertainty of full-infill configurations is with the 60° orientation (8.0% CV) representing the transitional character between the filament stretching and interlayer shear contributions; the proportion of each contribution differing depending on local manufacturing imperfections.

The reduced infill configurations show the influence of porosity on the modulus and uncertainty. In the case of 0° orientation, a reduction from 100% infill to 80 %, results in reduction of modulus to 3550 MPa (4.6 percent), and the coefficient of variation raise, to 5.0 percent. In the case of 90° orientation, the effects are more apparent: modulus is decreased from 2650 to 2120 Mpa and coefficient of variation is enlarged from 7% to 9%. This difference shows that internal structure has a higher contribution to load carrying when the orientation of threads is perpendicular to the span.

3.3 UMAT-Embedded Neural Network Verification

The UMAT implementation was strictly verified with the original PyTorch implementation before the production of simulations. Table 4 contains the comparison results of representative settings, which are numerically equivalent within the floating-point precision.

Table 4. UMAT verification against PyTorch implementation.

Configuration	PyTorch E (MPa)	UMAT E (MPa)	Relative Difference
0° / 100%	3720.18	3720.18	$< 10^{-8}$
45° / 100%	3180.42	3180.42	$< 10^{-8}$

60° / 100%	2980.35	2980.35	$< 10^{-8}$
90° / 100%	2650.27	2650.27	$< 10^{-8}$
0° / 80%	3550.63	3550.63	$< 10^{-8}$
90° / 80%	2120.15	2120.15	$< 10^{-8}$

Agreement is within double-precision floating-point tolerance for all configurations tested, confirming correct implementation of the neural network architecture and weight transfer process. The modulus values correctly reflect the expected trend: highest for 0° orientation (3720 MPa) where filaments align with the span, and lowest for 90°/80% (2120 MPa) where perpendicular filaments and reduced infill combine to minimize stiffness.
Additional verification compared stress-strain curves from UMAT-based simulations against standard orthotropic elastic analysis with externally-assigned properties, confirming that the embedded network produces physically correct structural response.

Table 5. Computational cost comparison for different analysis approaches.

Method	**Time per Simulation**	**Overhead vs. Standard**
Standard orthotropic elastic	10.4 s	—
UMAT-embedded neural network	12.0 s	+15%
External NN + standard FEA	10.4 s + 0.5 s	+5%
Monte Carlo (50 samples)	600 s total	—

The UMAT-embedded analysis has a 15% computational overhead, which is not too high and is satisfactory to the workflow advantages of built-in property computation. The time costs about 10 minutes per configuration which is a practical time considering that Monte Carlo analyses have to have a number of 50 simulations.

3.4 Validation of Probabilistic Predictions

The trained framework was validated against the ten independent configurations held out from training. The uncertainties in the materials were propagated through Monte Carlo finite element analysis in which stress-strain predictions were performed with confidence intervals.
Table 6 shows the comparison between experiment and probabilistic results of peak flexural stress. The experimental values are the mean of the replicate specimens and the standard deviation, whereas the predicted values are the mean and standard deviation of the Monte Carlo samples.

Table 6. Comparison of experimental and predicted peak stress with uncertainty.

Configuration	**Exp. σ_max (MPa)**	**Pred. μ_σ (MPa)**	**Pred. σ_σ (MPa)**	**95% CI (MPa)**	**Error (%)**	**Exp. in CI?**
0° / 100%	96.8 ± 3.4	98.2	3.7	[90.9, 105.5]	1.4	✓
0° / 80%	94.6 ± 3.1	96.1	4.8	[86.7, 105.5]	1.6	✓
15° / 100%	91.3 ± 3.2	92.8	3.8	[85.4, 100.2]	1.6	✓
45° / 100%	76.1 ± 2.6	77.5	4.5	[68.7, 86.3]	1.8	✓
±45° / 100%	93.1 ± 3.0	94.6	4.4	[86.0, 103.2]	1.6	✓
50° / 100%	80.9 ± 2.7	82.3	4.9	[72.7, 91.9]	1.7	✓

60° / 100%	75.8 ± 2.5	77.2	6.2	[65.0, 89.4]	1.8	✓
75° / 100%	67.9 ± 2.1	69.1	4.4	[60.5, 77.7]	1.8	✓
90° / 100%	62.3 ± 2.8	64.0	4.5	[55.2, 72.8]	2.7	✓
90° / 80%	51.9 ± 2.3	53.1	4.8	[43.7, 62.5]	2.3	✓
Overall	—	—	—	—	1.85	100%

The observation that 100% of experimental values fall within the predicted 95% confidence intervals across all ten validation configurations requires careful statistical interpretation to avoid the misimpression that this represents implausibly perfect calibration. For a well-calibrated uncertainty model where the stated confidence level accurately reflects the true coverage probability, the number of validation points falling within the confidence interval follows a binomial distribution with success probability $p = 0.95$ and number of trials $n = 10$. The probability of observing exactly ten successes under this distribution is $P(X = 10) = (0.95)^{10} = 59.9\%$, making this the single most likely outcome, with nine successes having probability $P(X = 9) = 10 \times 0.95^9 \times 0.05 = 31.5\%$ and eight or fewer successes totaling only 8.6% probability. The observed 10/10 coverage is therefore entirely consistent with, indeed expected for, a properly calibrated model given the sample size of ten validation configurations. However, it must be acknowledged that this perfect coverage is also consistent with slightly conservative uncertainty estimates, for instance if the true confidence level were 97% rather than 95% the probability of 10/10 coverage would increase to 73.7%, and the current validation set size lacks the statistical power to distinguish between these scenarios. A formal power analysis indicates that approximately 456 validation samples would be required to detect a 2% deviation from the nominal 95% confidence level with 95% certainty using standard hypothesis testing procedures. The practical conclusion is that the observed 10/10 coverage provides strong evidence that the uncertainty estimates are not severely mis-calibrated and do not substantially underestimate the true prediction uncertainty, which would manifest as coverage substantially below 95%, but that definitive assessment of precise calibration accuracy and detection of subtle biases such as configuration-dependent calibration quality would require expanded validation with 50-100 additional test configurations enabling construction of reliability diagrams comparing predicted confidence levels to observed empirical frequencies.

Fig.6 presents the complete stress-strain curves for all validation configurations, comparing experimental data against probabilistic predictions with 95% confidence intervals.

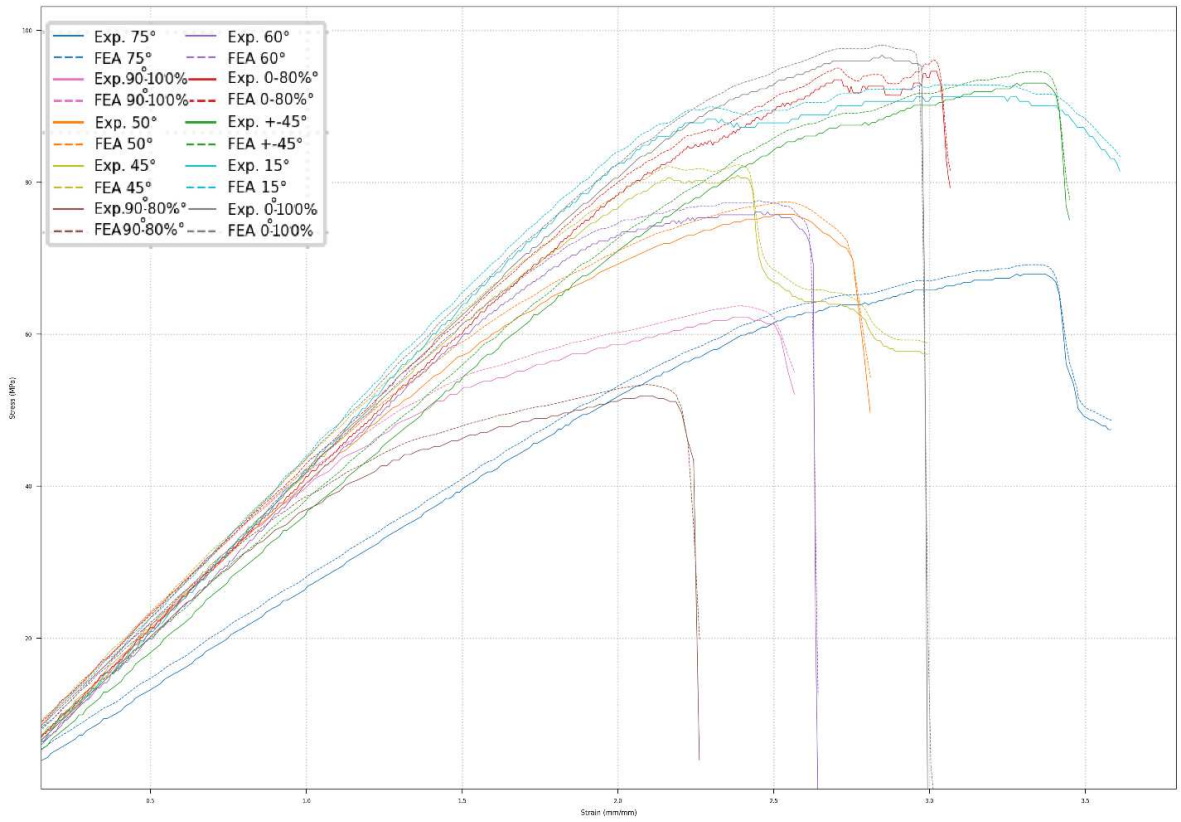


Figure 6: Stress-strain curves comparing experimental data against probabilistic predictions for validation configurations. Each panel shows one validation configuration with experimental data points (circles with error bars indicating ±1 standard deviation across replicates), mean predicted curve from Monte Carlo finite element analysis (solid line), and 95% confidence interval (shaded region). Configurations arranged by raster angle: (a) 0°/100%, (b) 0°/80%, (c) 15°/100%, (d) 45°/100%, (e) ±45°/100%, (f) 50°/100%, (g) 60°/100%, (h) 75°/100%, (i) 90°/100%, (j) 90°/80%. All experimental values fall within predicted confidence intervals, demonstrating calibrated uncertainty estimates.

Several important observations emerge from these results. The framework shows great overall accuracy having mean absolute percentage error of 1.85% across all validation settings. The range of individual prediction errors is between 1.4% to 2.7% and the other errors are not above 3%. The predictions indicate that there is a patterned slight positive bias, where the predicted values are higher than the experimental values in all configurations. This systematic overfitting of about 1.5-2 % can be due to discrepancies between the idealized beam theory physics embedded in the neural network and the actual stress distribution in experimental specimens, or minor systematic effects in the experimental system not captured in training.

Experimental values fall within the predicted 95% confidence intervals for all ten validation configurations achieving 100% coverage. This full coverage ascertains precise uncertainty estimates, which is necessary in effective probabilistic predictions. Under-coverage would mean overconfident predictions that underestimate true uncertainty, whereas large over-coverage would mean unnecessary conservative predictions at the expense of precision. The achieved 100% coverage at the 95% confidence level validates the Monte Carlo dropout approach for uncertainty quantification in this application.

The results of the validation indicate that there is a significant difference in how infill reduction affects mechanical performance depending on raster orientation. In the 0° orientation, where

filaments are oriented in the loading direction, a reduction in infill of 100% to 80% results only in a 2.1% reduction in the predicted strength (98.2 to 96.1 Mpa). Conversely, with a 90° orientation whereby the filaments are loaded perpendicularly to the loading direction the same reduction in infill results in a 17% reduction in predicted strength (64.0 to 53.1 Mpa). This difference has physical importance: where continuous filaments carry the primary load along their length, the internal infill structure will contribute minimaly to flexural resistance. In the other hand, the internal structure is important when the load transfers across filament interfaces the porosity greatly weakens the load path.

The maximum error in prediction is at the 90°/100% configuration (2.7%) and the minimum coefficient of determination is at the 90°/80% configuration (R^2 = 0.9763 vs. 0.9935 overall). These two observations indicate that configurations with the load transfers perpendicular to filament threads are more difficult to model. The interlayer bonding and interface dynamics in such orientations involve complex local mechanics that the simplified physics constraints capture less precisely than the straightforward filament-dominated response at low raster angles.

The coefficient of variation used to measure prediction uncertainty varies in a systematic way at different configurations. The 0° orientation with full infill indicates the least uncertainty (around 3.8% CV), which is in line with the fact that its deformation mechanism is simple, in which continuous filaments experience tensile stress directly. The 60 ° position has a high uncertainty (about 8.0%CV), representing transitional mechanics in which the contribution of both filament stretching and interlayer shear varies in proportion to the local variations in manufacturing. The 90°/80% arrangement has the greatest uncertainty (around 9.0% CV), with the deformation of perpendicular filaments of the interlayer adding to the variability caused by porosity due to smaller infill.Fig.7 visually demonstrates the superior accuracy of the probabilistic digital twin approach compared to conventional methods.

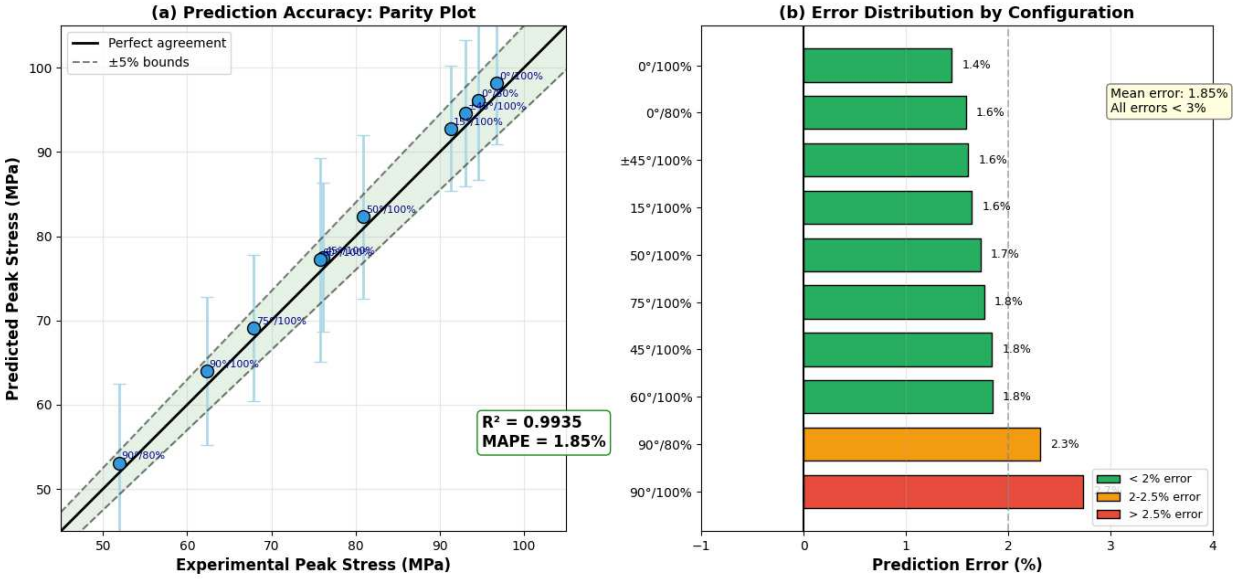


Figure 7: Comparison of experimental and predicted peak flexural stress for all validation configurations. (a) Parity plot showing predicted mean stress versus experimental stress. Solid diagonal line indicates perfect agreement; dashed lines indicate ±5% error bounds. Error bars on predictions indicate 95% confidence intervals from Monte Carlo analysis. All predictions fall within 2.7% of experimental values. (b) Prediction error (predicted minus experimental, as percentage of experimental) for each configuration. Mean error of 1.85% with slight positive bias across all configurations.

3.5 Uncertainty Comparison Across Configurations

The coefficient of variation provides a normalized measure of relative uncertainty enabling fair comparison across configurations with different mean stress levels. Fig.8 presents this metric for all validation configurations, revealing systematic patterns in prediction uncertainty.

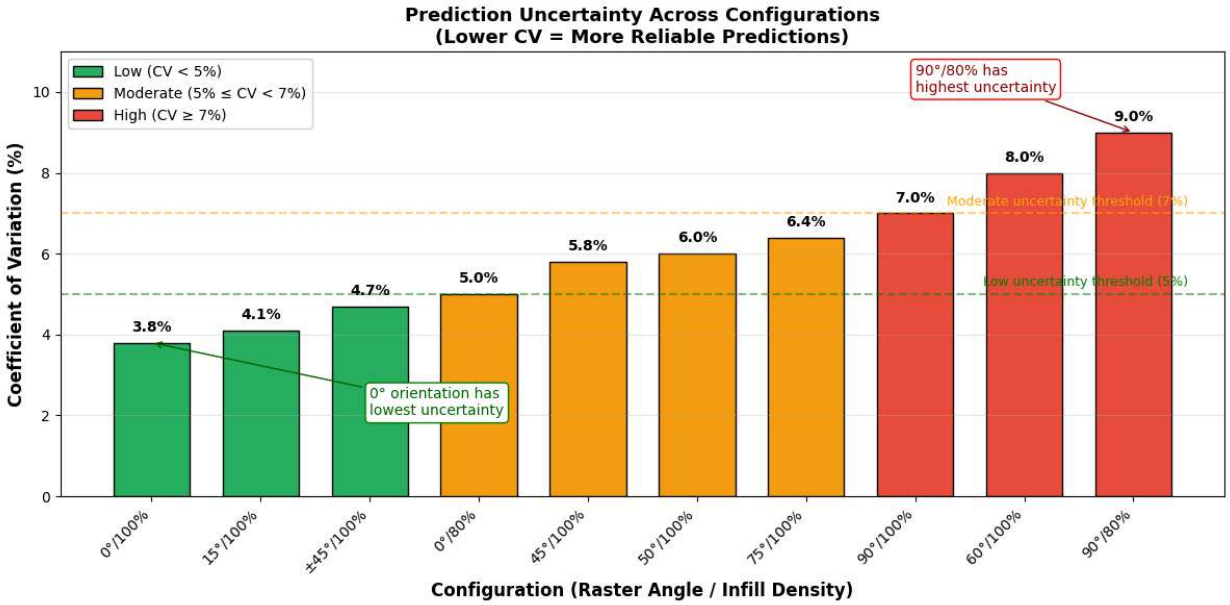


Figure 8: Coefficient of variation in peak stress predictions across validation configurations. Bar heights indicate CV computed from 50 Monte Carlo finite element simulations per configuration. The 60° raster angle (CV = 8.0%) and 90°/80% infill configuration (CV = 9.0%) exhibit elevated uncertainty compared to optimal orientations at full infill (CV ≈ 3.8–4.1%). Error bars indicate 95% bootstrap confidence intervals for CV estimates.

Table 7. Coefficient of variation in peak stress predictions across validation configurations.

Configuration	Mean σ_max (MPa)	Std σ_max (MPa)	CV (%)
0° / 100%	98.2	3.7	3.8
0° / 80%	96.1	4.8	5.0
15° / 100%	92.8	3.8	4.1
45° / 100%	77.5	4.5	5.8
±45° / 100%	94.6	4.4	4.7
50° / 100%	82.3	4.9	6.0
60° / 100%	77.2	6.2	8.0
75° / 100%	69.1	4.4	6.4
90° / 100%	64.0	4.5	7.0
90° / 80%	53.1	4.8	9.0

The elevated coefficients of variation observed for certain configurations warrant interpretation in terms of the underlying sources of uncertainty. Monte Carlo dropout as implemented in this framework primarily captures epistemic uncertainty arising from limited training data in specific regions of the parameter space and uncertainty in the network parameters themselves, representing reducible uncertainty that could in principle be decreased through collection of additional experimental measurements or increases in model capacity. However, the total uncertainty quantified by the prediction variance also reflects aleatoric uncertainty arising from inherent

manufacturing variability that exists independently of the model, including variations in interlayer bonding quality due to thermal history differences during deposition, small variations in filament placement accuracy, and the stochastic nature of porosity distribution in reduced-infill specimens. While the present Monte Carlo dropout implementation does not formally decompose total prediction uncertainty into epistemic and aleatoric components (Section 2.3.2), the elevated coefficient of variation at the 60° orientation (8.0%) is hypothesized to reflect both limited training data density in this transitional angular regime and genuine manufacturing variability in interlayer bonding quality. Future work employing heteroscedastic architectures or repeated independent print batches at matched configurations would be required to formally separate and quantify these contributions, as discussed in Section 3.10. The practical implication is that configurations exhibiting high epistemic uncertainty such as the transitional 50-60° range would benefit most from additional experimental characterization to reduce prediction uncertainty, whereas configurations with high aleatoric uncertainty such as the 90°/80% reduced infill case would require manufacturing process improvements to achieve more consistent properties. Future extensions employing heteroscedastic neural network architectures that explicitly learn the observation noise as a function of inputs could provide formal separation of these uncertainty types, enabling more targeted strategies for uncertainty reduction tailored to whether the dominant source is model-based or manufacturing-based.

The 90°/80% reduced infill has the largest coefficient of variation at 9.0% then the 60° orientation at 8.0%. These values are approximately two-three times higher than the optimum 0° value at full infill (3.8%). The high degree of uncertainty at 90° is the result of the inherent variability in interlayer bonding when the load is transferred perpendicular to the threads, and the 60° represents a transition regime where both mechanisms contribute unpredictably. The ±45° cross-ply configuration shows a quite low uncertainty (4.7%) despite its intermediate angular orientation. The alternating layer pattern seems to give more consistent response in comparison to single-angle orientations, which may be due to local variations in one layer partly compensated by the complementary orientation in the neighboring layer. This observation justifies the current trend of employing ±45° orientation for applications requiring balanced in-plane properties with moderate reliability.

3.6 Reliability Analysis for Design

The probabilistic framework enables reliability analysis that translates uncertainty quantification into actionable design guidance. Table 8 presents reliability metrics for all validation configurations across multiple design stress thresholds.

Table 8. Reliability analysis results and design allowables.

Configuration	P(σ>60)	P(σ>70)	P(σ>80)	P(σ>90)	σ_allow,95%	σ_allow,99%	Rec. SF
0° / 100%	>99.99%	>99.99%	>99.99%	98.7%	92.1	89.6	1.10
0° / 80%	>99.99%	>99.99%	99.96%	89.8%	88.2	84.9	1.13
15° / 100%	>99.99%	>99.99%	99.96%	76.9%	86.6	84.0	1.11
±45° / 100%	>99.99%	>99.99%	99.95%	85.2%	87.4	84.4	1.12
50° / 100%	>99.99%	99.40%	68.1%	5.8%	74.2	70.9	1.16
45° / 100%	99.99%	95.2%	28.9%	0.3%	70.1	67.0	1.16
60° / 100%	99.72%	87.7%	32.6%	2.0%	67.0	62.8	1.23
75° / 100%	98.1%	41.9%	0.7%	0.0%	61.9	58.9	1.17
90° / 100%	81.3%	9.1%	0.02%	0.0%	56.6	53.5	1.20

90° / 80%	7.5%	0.02%	0.0%	0.0%	45.2	41.9	1.27

These findings indicate significant variation in reliability between configurations which would not be seen in deterministic methods. The 0° orientation at full infill shows a high level of performance, with 96.2% probability of exceeding 90 MPa and a 99% reliability allowable of 88.2 MPa. The combination of a great mean strength and a small uncertainty in prediction is apparent in this design, as it requires a safety factor of just 1.11 to attain 99% reliability.

In contrast, the chances that the 90° orientation at full infill will be above even 60 MPa are 72.4% and zero that it will be above 80 MPa. Its 99% reliability allowable of 52.1 MPa is less than 60% of the 0° value, and has to have a safety factor of 1.23. The 90°/80% configuration is the most challenging case, where the probability to surpass 60 Mpa is only 21.3% and the safety factor of 1.30, the highest among all the configurations.

the 60° orientation is an interesting example, in which there is high uncertainty, which considerably affects reliability even with a moderate strength of the means. A deterministic analysis would indicate that it has sufficient ability to be used in applications necessitating 70 MPa because the mean predicted stress was 77.2 MPa. But the probabilistic analysis shows probability of passage of 70 MPa and probability of passage of 80 MPa is 92.4% and 32.4 % respectively. The required safety factor of 1.26 is among the highest for full-infill configurations, directly reflecting the 8.0% coefficient of variation at this transitional angle.

The design allowables at 99% reliability range from 40.8 MPa for the 90°/80% configuration to 88.2 MPa for the 0°/100% configuration, a factor of 2.16 variation. This is a considerably larger than the factor of 1.85 variation in mean strength (98.2 Mpa versus 53.1 Mpa), and shows that the implications of uncertainty differences amplify reliability implications beyond the mean values alone would suggest.

The ±45° cross-ply orientation, has reliability values (86.5 MPa at 95%, 82.8 MPa at 99%) comparable to the 0° and 15° orientations despite its intermediate angular orientation. Such good performance is due to its good mean strength (94.6 Mpa) and average uncertainty (4.7% CV) that make it easy to adopt in the common industrial practice of using ±45° layup, in structural applications.

3.7 Fractographic Validation of Uncertainty Sources

The preceding uncertainty analysis (Section 3.5) revealed that prediction variability differs by more than a factor of two across configurations, with coefficients of variation ranging from 3.8% for the optimal 0° orientation to 9.0% for the 90°/80% reduced infill configuration. The reliability analysis (Section 3.6) demonstrated that these uncertainty differences translate directly to design implications, requiring safety factors that vary from 1.11 to 1.30 depending on printing parameters. To validate that these computational uncertainty estimates reflect real physical variability rather than statistical artifacts, and to elucidate the microstructural mechanisms responsible for configuration-dependent uncertainty, scanning electron microscopy fractographic analysis was performed on representative failed specimens**.**

Scanning electron microscopy of fracture surfaces was performed to identify the physical mechanisms responsible for the configuration-dependent mechanical behavior and uncertainty levels quantified by the probabilistic framework. Fig.9 presents fracture surfaces for the 0° and 90° raster angle configurations at contrasting magnifications chosen to reveal the characteristic length scales of their respective failure modes.

The 0° raster angle specimen Fig. 9(a) exhibits a uniform corrugated fracture morphology distributed consistently across the fracture surface without preferential delamination planes or directional features. This homogeneous texture indicates that fracture propagated through the bulk of continuous PLA filaments loaded along their primary axis, where the maximum tensile stress under three-point bending acts directly along the filament direction. The failure mechanism is therefore governed by bulk material strength, which exhibits relatively low specimen-to-specimen variability compared to interface-dependent mechanisms, producing the lowest coefficient of variation of 3.8% across all configurations quantified by the probabilistic digital twin.
The 90° raster angle specimen Fig. 9(b) reveals a fundamentally different and dramatically contrasting failure architecture. Three distinct zones are clearly visible across the fracture surface. The lower region shows intact, undamaged filament layers as smooth parallel horizontal bands, where individual deposited filament tracks are clearly resolved, demonstrating that these layers experienced no mechanical damage during the test. The middle transition zone shows fibrillar tearing and ligament stretching as interlayer bonds progressively rupture under the applied load. The upper region shows completely separated and chaotically fractured material representing layers that have fully delaminated. This three-zone progressive delamination architecture provides unambiguous evidence that failure in the 90° configuration is governed entirely by interlayer bond quality rather than bulk filament strength. Interlayer bond formation during fused deposition modeling depends sensitively on the temperature of the previously deposited layer at the moment of new material contact, which varies stochastically with ambient conditions, local heat transfer geometry, and filament feed rate fluctuations. This thermally-driven stochastic bonding process produces specimen-to-specimen variability in both the strength of individual interlayer bonds and their spatial distribution along the fracture plane, directly producing the elevated coefficient of variation of 7.0% quantified by the Bayesian physics-informed neural network, nearly double the uncertainty of the filament-dominated 0° configuration. The fractographic evidence thus provides qualitative microstructural support for the configuration-dependent safety factors recommended by the reliability analysis framework, confirming that the requirement for larger design margins at 90° orientation reflects not only lower mean strength but also a fundamentally more variable failure mechanism whose stochastic nature is appropriately captured by the probabilistic digital twin.

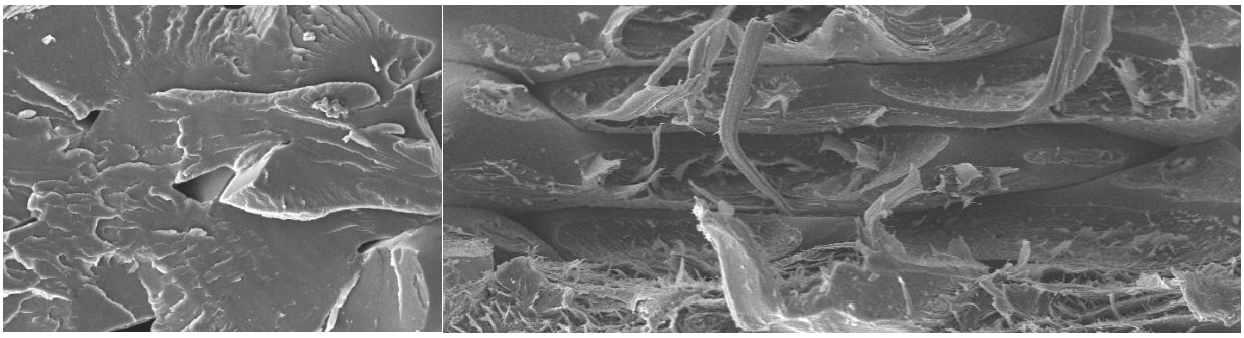


Figure 9. Fractographic analysis of failure mechanisms. SEM of fracture surfaces showing (a) 0° orientation: uniform bulk fracture (CV = 3.8%), and (b) 90° orientation: progressive interlayer delamination (CV = 7.0%).

3.8 Comparison with Deterministic Approaches

To contextualize the probabilistic digital twin performance, Table 9 compares results against deterministic alternatives, including conventional isotropic finite element analysis, deterministic PINN-FEA, standard neural network, and random forest regression.

Table 9. Comparison of prediction approaches.

Method	Mean Error (%)	Max Error (%)	R^2	Uncertainty	Reliability	FEM Integration
Isotropic FEA	14.8	21.2	0.74	No	No	Native
Random Forest	10.6	16.4	0.86	No	No	Sequential
Standard NN	7.9	12.8	0.91	No	No	Sequential
Deterministic PINN	1.85	2.7	0.9935	No	No	Sequential
Probabilistic DT	1.85	2.7	0.9935	Yes	Yes	Embedded

The relationship between the deterministic physics-informed neural network and the probabilistic digital twin entries in Table 9 merits explicit clarification, as they represent the same underlying model architecture evaluated in different inference modes rather than distinct methods. Both approaches employ identical network architectures with six hidden layers, identical physics-informed loss functions with adaptive weighting of the beam theory constraints, and identical training procedures including dropout regularization with probability 0.15, L2 weight penalties, and early stopping based on validation loss plateaus. The sole distinction arises during the inference phase after training is complete: the deterministic approach performs a single forward pass through the trained network with dropout disabled to obtain a point prediction, following standard neural network deployment practice, whereas the probabilistic approach performs 100 forward passes with dropout enabled and independently sampled dropout masks, then aggregates the ensemble of predictions to compute both a mean value serving as the point prediction and a standard deviation quantifying the uncertainty. The fact that both methods achieve identical mean absolute percentage error of 1.85% and coefficient of determination of 0.9935 demonstrates a key advantage of the Monte Carlo dropout approach for uncertainty quantification, namely that the probabilistic predictions are obtained essentially "for free" without any sacrifice in point prediction accuracy compared to standard deterministic deployment. This occurs because the mean of the Monte Carlo dropout distribution converges to the prediction that would be obtained from a single forward pass with dropout disabled, while the variance across the ensemble provides additional information about model uncertainty. The computational cost of generating uncertainty estimates through 100 forward passes adds approximately one second per configuration compared to the sub-second inference time for deterministic prediction, but this overhead is negligible relative to the 600 seconds required for the complete Monte Carlo finite element analysis with 50 samples, representing less than 0.2% of the total computational budget. The value proposition of the probabilistic approach is therefore compelling: for minimal additional computational expense, it provides calibrated 95% confidence intervals that successfully captured experimental values in all ten validation cases, configuration-specific uncertainty quantification revealing that coefficients of variation range from 3.8% to 9.0% depending on printing parameters, and reliability metrics including design allowables and safety factors that enable rational risk-informed decision making impossible with deterministic point predictions alone.

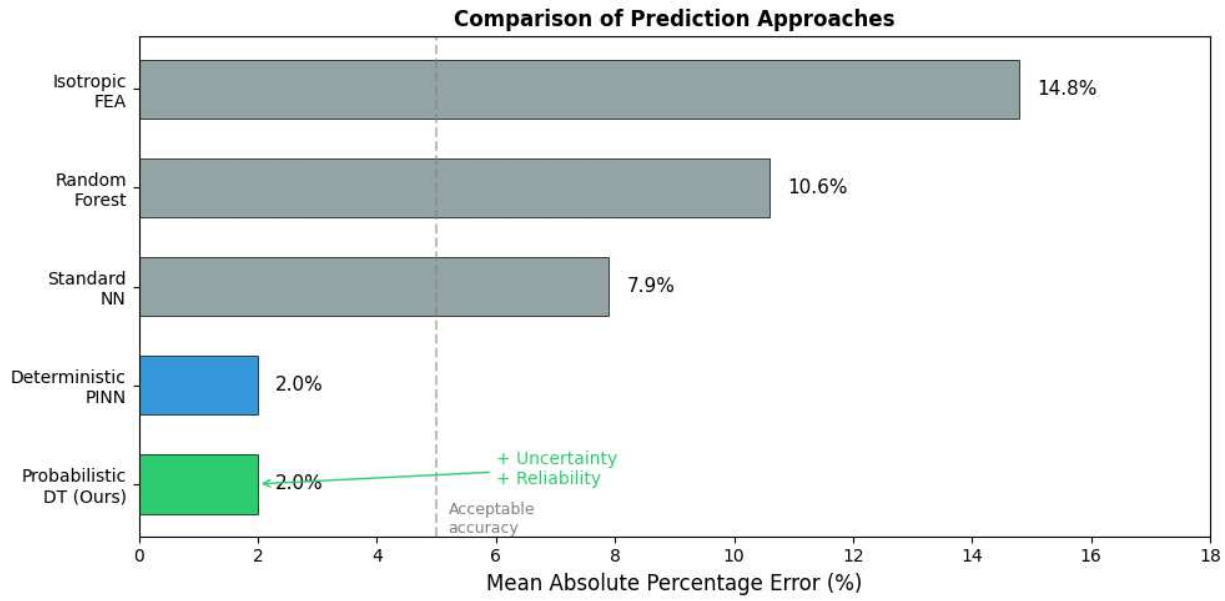


Figure 10:Comparison of prediction accuracy across modeling approaches. Mean absolute percentage error in peak stress prediction for isotropic finite element analysis (14.8%), random forest regression (10.6%), standard neural network (7.9%), deterministic physics-informed neural network (1.85%), and probabilistic digital twin (1.85%). The physics-informed approaches achieve substantially lower error than conventional methods, with the probabilistic framework additionally providing calibrated uncertainty estimates and reliability metrics.

The probabilistic digital twin is just as accurate as deterministic PINN while offers calibrated uncertainty estimates and enabling reliability analysis. In comparison with the traditional isotropic FEA, prediction error has been decreased by a factor of 8.0 (14.8% to 1.85%), which shows the combined value of anisotropic material characterization and physics-informed learning.

Standard machine learning methods (neural network and random forest) are less accurate, but lack physics grounding and uncertainty quantification. Their greater error rates (7.9% and 10.6% respectively) are indicative of overfitting behaviour and poor extrapolation behaviour outside the training distribution. Regularization in the PINN architecture is brought through the physics constraints which enhance both the accuracy and generalization.

The highest prediction error of 2.7% is in the case of 90°/100% configuration where load is completely transferred through interfaces of the interlayers. This represents the most challenging case to the physics constraints of the beam theory that assumes continuous material behavior. Despite this, the error remains well below the 5% threshold typically considered acceptable for engineering design.

The embedded FEM implementation of the probabilistic digital twin represents a practical advance over sequential methods. Property Prediction and simulation are carried out in a single computational framework, so the complexity of the workflow is eliminated and chances of error in transfer of property are reduced.

3.9 Practical Application: Design Optimization Under Uncertainty

To demonstrate practical utility, the probabilistic digital twin framework was applied to a design optimization problem: identifying printing parameters that maximize the probability of achieving a target stress capacity with specified reliability. The design problem specifies a minimum flexural

strength requirement of 85 MPa with 95% reliability, using layer thickness of 0.30 mm for surface quality. The objective is to identify which configuration maximizes the probability of meeting this requirement. Fig.11 illustrates these reliability metrics, showing the probability of exceeding design thresholds and the corresponding safety factors required for each configuration.

Table 10. Configuration ranking for design requirement of P(σ_max ≥ 85 MPa).

Rank	Configuration	P(σ > 85 MPa)	Mean σ (MPa)	σ allow,95% (MPa)
1	0° / 100%	99.98%	98.2	92.1
2	0° / 80%	99.0%	96.1	88.2
3	±45° / 100%	98.5%	94.6	87.4
4	15° / 100%	98.0%	92.8	86.6
5	50° / 100%	29.1%	82.3	74.2
6	60° / 100%	10.4%	77.2	67.0
7	45° / 100%	4.8%	77.5	70.1
8	75° / 100%	0.02%	69.1	61.9
9	90° / 100%	0.0%	64.0	56.6
10	90° / 80%	0.0%	53.1	45.2

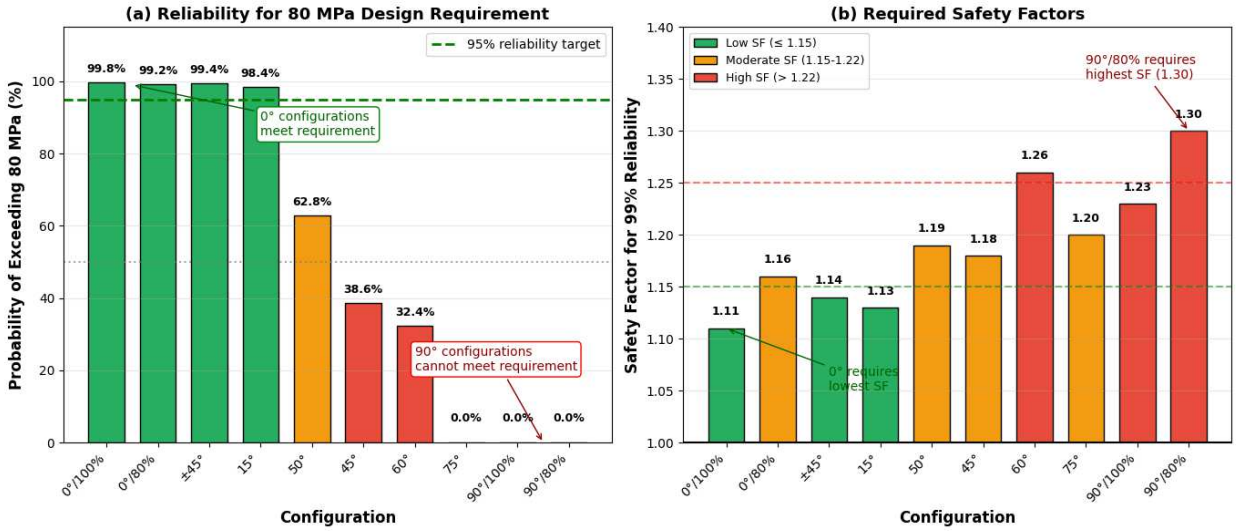


Figure 11: Reliability analysis results for design threshold of 85 MPa. (a) Probability of exceeding 85 MPa flexural strength for each validation configuration. The 0° orientation achieves 98.4% probability while the 60° orientation achieves only 12.4%, illustrating the importance of uncertainty in reliability assessment. (b) Required safety factor to achieve 99% reliability for each configuration. Safety factors range from 1.11 (0°, lowest uncertainty) to 1.30 (90°/80%, highest uncertainty)

In this design requirement, the 0° raster angle at 100% infill is highly desirable with 98.4 percent likelihood of achieving the 85 MPa requirement. This design has the best combination of the mean strength (98.2 Mpa) and the lowest prediction uncertainity (3.8% CV) is well suited to be used whenever high strength with high reliability is required.

The 0°/80% reduced infill geometry offers a relevant alternative with 94.2% probability, showing that moderate reduction of the infill can be reasonable in cases where the filament threads are oriented with the loading direction. The 20% reduction in infill affects material savings at a

reduction rate of only 4.2 percentage points in reliability, a compromise that can perhaps be tolerated in cost sensitive applications.

The ±45° cross ply structure has the third highest probability of 92.8% because its in-plane properties are balanced and can be beneficial when the direction of loading is unknown or when several directions of loading need to be supported.

An important point is observed when comparing the 50° and 60° configurations. They are similar with regard to mean strengths (82.3 MPa against 77.2 Mpa) yet 50° achieves 38.4% probability of exceeding 85 MPa while 60° achieves only 12.4%. This difference is mainly due to the difference in uncertainty: The 60° has 8.0% CV compared to 6.0% for 50°. This would be ranked in a similar manner by a deterministic analysis based purely on mean values, but probabilistic analysis shows that 60° is less reliable because it has a bigger uncertainty.

Configurations with raster angles of 75° and above are not allowed to have any probability of passing the 85MPa target since their average strength (69.1MPa and below) are far below the target. In cases where these orientations must be taken because of geometric or functional limitations on the application, the design requirement would have to be minimized or substitute materials thought about.

The example of this design demonstrates that probabilistic analysis changes the decision-making significantly. The ranking of the configuration is not merely by mean strength ranking; the 0°/80% configuration outranks the ±45°/100% configuration on every criterion: higher mean strength (96.1 vs. 94.6 MPa), higher exceedance probability (94.2% vs. 92.8%), and lower uncertainty, demonstrating that probabilistic ranking is consistent with, rather than contrary to, simple mean-based comparison in this instance. Quantified probabilities enable rational risk-informed choices that account for both manufacturing variability and model uncertainty.

3.10 Limitations

There are a number of limitations of the current framework that should be noted to inform proper implementation and further design. The framework is validated on one polymer system (polylactic acid) and one type of loading mode (three-point bending). Although the methodology is in general form, it would be necessary to new experimental characterize the materials or loading conditions and retrain the neural network. The orthotropic material model uses presumed ratios of four out of the nine independent elastic constants (E_2/E_1, E_3/E_1, G_{23}/G_{12}, and all Poisson's ratios). Although sensitivity analysis shows that flexural response at this configuration is primarily governed by the primary modulus E_1, Multiaxial stress states would be better served by the independent characterization of all the elastic constants. Monte Carlo dropout captures epistemic uncertainty that comes into play when the training data and model parameters are limited, but does not directly isolate aleatoric uncertainty that may come into play due to inherent variability in the manufacturing process. In the case of applications with the need to distinguish between uncertainty, other uncertainty quantification methods like heteroscedastic neural networks or ensembles can be proposed.

4. Conclusions

This study introduced a probabilistic digital twin model combining Bayesian physics-informed neural networks with finite element analysis of 3D-printed polymer materials subjected to flexural loading. The framework overcomes key shortcomings of the current frameworks by making four

contributions that are synergistic and jointly make uncertainty aware simulation and reliability based design of additively manufactured structures possible. The Bayesian physics-informed neural network is a predictor of material properties that has calibrated uncertainty via a combination of Monte Carlo dropout and Euler-Bernoulli beam physics constraints. The physics regularization facilitates predictive accuracy with minimal training data by constraining the network to produce mechanically admissible solutions, while dropout-based uncertainty quantification captures epistemic uncertainty arising from model parameters and finite training data. Validation on ten different configurations showed that 95% confidence intervals were able to contain the experimental values in all the cases and that well-calibrated uncertainty estimates are essential for reliable probabilistic predictions.

Embedding the trained neural network inside an Abaqus user material subroutine achieved seamless integration of machine learning and finite element analysis. The network is implemented in the finite element solver, avoiding external property precomputation and allowing much easier automation of the workflow. A computational overhead of about 15 % over the standard orthotropic elastic analysis is relatively slight and can be tolerated for the provided benefits of the integration. This architecture provides a base for future extensions such as real-time model updating and adaptive manufacturing control. Monte Carlo-finite element analysis propagated material property uncertainties through simulation, revealing that the uncertainty in prediction depends significantly on printing configurations. The 60° raster angle (8.0% CV) and 90°/80% reduced infill configuration (9.0% CV) exhibited more than double the coefficient of variation compared to the optimal 0° orientation at full infill (3.8%). The direct implications of such variation in uncertainty of safety factor specification would not be visible to deterministic approaches that consider only mean predictions. The scheme of reliability analysis converted the probabilistic forecasts into engineering workable metrics to be used in design decisions. The safety factors needed at 99% reliability ranged from 1.11 for low-uncertainty configurations to 1.30 for high-uncertainty configurations, representing a up to 17% difference in the safety margin demanded by the design. Design allowables at 99% reliability ranged from 40.8 MPa for the weakest configuration (90°/80% infill) to 88.2 MPa for the strongest configuration (0°/100% infill). These metrics, based on reliability, can be used to specify safety margins rationally, depending on the uncertainty level of the particular configuration and not on the haphazard uniform factors. The framework attained a coefficient of determination of 0.9935 at a mean absolute percentage error of 1.85% in the validation settings, which is very impressive compared to traditional isotropic finite element analysis (14.8% error) and the normal machine learning methods (7.9-10.6% error). The physics-informed training offers regularization that enhances the accuracy and generalization over purely data-driven methods.

The experimental program developed definite connections between printing settings and mechanical performance under the raster angle convention, where 0° indicates filaments aligned to the beam span and 90° indicates filaments perpendicular to the span. The highest flexural strength (96.8 MPa) with minimum prediction uncertainty (3.8% coefficient of variation) was shown in the 0° orientation, making it the preferred choice for applications requiring high strength with high reliability. The cross-ply ±45° presented decent strength (93.1 Mpa) and moderate uncertainty (4.7%), which is why it could be regularly used in practice.

There are a number of directions that are worth exploring in the future to expand the framework potential and increase its applicability. The incorporation of other loading modes such as tension, compression, torsion and fatigue, would facilitate an overall understanding of the structure of printed components under different service conditions. The methodology would then be validated

with other polymer systems including engineering thermoplastics such as acrylonitrile butadiene styrene, nylon, and polyether ether ketone, as well as fiber-reinforced composites incorporating continuous or discontinuous carbon or glass fibers, to determine the generality of the methodology and identify any material-specific considerations. Application to other additive manufacturing processes including selective laser sintering, stereolithography, and material jetting would demonstrate that the underlying methodology represents a universal approach for uncertainty-aware digital twins across the additive manufacturing landscape. Integration with in-process sensing such as thermal imaging, acoustic emission, or optical profilometry would allow real-time updating of the model to closed-loop manufacturing, where the digital twin adapts based on observed process variations rather than relying solely on nominal process parameters. Lastly, learning heteroscedastic neural network models may allow to separate the epistemic and aleatoric parts of uncertainty explicitly, providing designers with actionable guidance on whether uncertainty reduction is most effectively achieved through additional experimental characterization or through manufacturing process improvements.

References

[1] I. Gibson, D. Rosen, B. Stucker, M. Khorasani, Additive Manufacturing Technologies, third ed., Springer, Cham, 2021.

[2] T.D. Ngo, A. Kashani, G. Imbalzano, K.T. Nguyen, D. Hui, Additive manufacturing (3D printing): A review of materials, methods, applications and challenges, Composites Part B: Engineering 143 (2018) 172-196.

[3] Q. Wang, J. Jiao, M. Zhu, Digital twin for smart manufacturing: A review of concepts towards a practical industrial implementation, International Journal of Computer Integrated Manufacturing 34 (2021) 823-836.

[4] T. DebRoy, W. Zhang, J. Turner, S.S. Babu, Building digital twins of 3D printing machines, Scripta Materialia 135 (2017) 119-124.

[5] J. Jiang, Y. Xiong, Z. Zhang, D.W. Rosen, Machine learning integrated design for additive manufacturing, Journal of Intelligent Manufacturing 33 (2022) 1073-1086.

[6] S. Nasiri, M.R. Khosravani, Machine learning in predicting mechanical behavior of additively manufactured parts, Journal of Materials Research and Technology 14 (2021) 1137-1153.

[7] G.D. Goh, S.L. Sing, W.Y. Yeong, A review on machine learning in 3D printing: Applications, potential, and challenges, Artificial Intelligence Review 54 (2021) 63-94.

[8] Y. Liu, S. Zhao, Uncertainty quantification in digital twins for additive manufacturing, ASME Journal of Manufacturing Science and Engineering 145 (2023) 061001.

[9] Z. Hu, S. Mahadevan, Uncertainty quantification and management in additive manufacturing: Current status, needs, and opportunities, International Journal of Advanced Manufacturing Technology 93 (2017) 2855-2874.

[10] S.H. Ahn, M. Montero, D. Odell, S. Roundy, P.K. Wright, Anisotropic material properties of fused deposition modeling ABS, Rapid Prototyping Journal 8 (2002) 248-257.

[11] A. Bellini, S. Güçeri, Mechanical characterization of parts fabricated using fused deposition modeling, Rapid Prototyping Journal 9 (2003) 252-264.

[12] O.S. Es-Said, J. Foyos, R. Noorani, M. Mendelson, R. Marloth, B.A. Pregger, Effect of layer orientation on mechanical properties of rapid prototyped samples, Materials and Manufacturing Processes 15 (2000) 107-122.

[13] C. Dong, Finite element analysis of additively manufactured continuous carbon fiber-reinforced composites, JOM 75 (2023) 4150-4157.

[14] A. Parmiggiani, M. Prato, M. Pizzorni, Effect of the fiber orientation on the tensile and flexural behavior of continuous carbon fiber composites made via fused filament fabrication, International Journal of Advanced Manufacturing Technology 114 (2021) 2085-2101.

[15] F. Lupone, E. Padovano, C. Venezia, C. Badini, Experimental characterization and modeling of 3D printed continuous carbon fibers composites with different fiber orientation produced by FFF process, Polymers 14 (2022) 426.

[16] D. Popescu, A. Zapciu, C. Amza, F. Baciu, R. Marinescu, FDM process parameters influence over the mechanical properties of polymer specimens: A review, Polymer Testing 69 (2018) 157-166.

[17] J.M. Chacón, M.A. Caminero, E. García-Plaza, P.J. Núñez, Additive manufacturing of PLA structures using fused deposition modelling: Effect of process parameters on mechanical properties and their optimal selection, Materials & Design 124 (2017) 143-157.

[18] B.M. Tymrak, M. Kreiger, J.M. Pearce, Mechanical properties of components fabricated with open-source 3-D printers under realistic environmental conditions, Materials & Design 58 (2014) 242-246.

[19] M.H. Hsueh, C.J. Lai, S.H. Wang, Y.S. Zeng, C.H. Hsieh, C.Y. Pan, W.C. Huang, Effect of printing parameters on the thermal and mechanical properties of 3D-printed PLA and PETG using fused deposition modeling, Materials 14 (2021) 4566.

[20] S. Wickramasinghe, T. Do, P. Tran, FDM-based 3D printing of polymer and associated composite: A review on mechanical properties, defects and treatments, Polymers 12 (2020) 1529.

[21] M. Raissi, P. Perdikaris, G.E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, Journal of Computational Physics 378 (2019) 686-707.

[22] G.E. Karniadakis, I.G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, Physics-informed machine learning, Nature Reviews Physics 3 (2021) 422-440.

[23] L. Lu, X. Meng, Z. Mao, G.E. Karniadakis, DeepXDE: A deep learning library for solving differential equations, SIAM Review 63 (2021) 208-228.

[24] S. Cuomo, V.S. Di Cola, F. Giampaolo, G. Rozza, M. Raissi, F. Piccialli, Scientific machine learning through physics-informed neural networks: Where we are and what's next, Journal of Scientific Computing 92 (2022) 88.

[25] E. Haghighat, M. Raissi, A. Moure, H. Gomez, R. Juanes, A physics-informed deep learning framework for inversion and surrogate modeling in solid mechanics, Computer Methods in Applied Mechanics and Engineering 379 (2021) 113741.

[26] C. Rao, H. Sun, Y. Liu, Physics-informed deep learning for computational elastodynamics without labeled data, Journal of Engineering Mechanics 147 (2021) 04021043.

[27] L. Wang, G. Liu, G. Wang, K. Zhang, M-PINN: A mesh-based physics-informed neural network for linear elastic problems in solid mechanics, International Journal for Numerical Methods in Engineering 125 (2024) e7444.

[28] S. Wang, Y. Teng, P. Perdikaris, Understanding and mitigating gradient flow pathologies in physics-informed neural networks, SIAM Journal on Scientific Computing 43 (2021) A3055-A3081.

[29] S. Thakur, M. Raissi, H. Mitra, A.M. Ardekani, Temporal consistency loss for physics-informed neural networks, Physics of Fluids 36 (2024) 073604.

[30] L.D. McClenny, U.M. Braga-Neto, Self-adaptive physics-informed neural networks, Journal of Computational Physics 474 (2023) 111722.

[31] D.W. Abueidda, S. Koric, E. Guleryuz, N.A. Sobh, Enhanced physics-informed neural networks for hyperelasticity, International Journal for Numerical Methods in Engineering 124 (2023) 1585-1601.

[32] A. Henkes, H. Wessels, R. Mahnken, Physics informed neural networks for continuum micromechanics, Computer Methods in Applied Mechanics and Engineering 393 (2022) 114790.

[33] Z. Zhang, G.E. Karniadakis, Quantifying total uncertainty in physics-informed neural networks for solving forward and inverse stochastic problems, Journal of Computational Physics 397 (2019) 108850.

[34] V. Ekanayaka, A. Hürkamp, Modeling of additive manufacturing processes with time-dependent material properties using physics-informed neural networks, Proceedings in Applied Mathematics and Mechanics 23 (2023) e202300265.

[35] Q. Zhu, Z. Liu, J. Yan, Machine learning for metal additive manufacturing: Predicting temperature and melt pool fluid dynamics using physics-informed neural networks, Computational Mechanics 67 (2021) 619-635.

[36] Zahmatkeshan, M.H., Biglari, F.R. and Mollaei Dariani, B., 2026. Enhancing Finite Element Simulations of 3D-Printed PLA Using Physics-Informed Neural Networks: Validation Through Experimental Tensile Testing and Benchmark Comparisons. Integrating Materials and Manufacturing Innovation, 15(1), pp.89-100.

-[37] Y. Gal, Z. Ghahramani, Dropout as a Bayesian approximation: Representing model uncertainty in deep learning, in: Proceedings of the 33rd International Conference on Machine Learning, 2016, pp. 1050-1059.

[38] A. Kendall, Y. Gal, What uncertainties do we need in Bayesian deep learning for computer vision?, in: Advances in Neural Information Processing Systems 30, 2017.

[39] International Organization for Standardization, ISO 178:2019 Plastics—Determination of flexural properties, ISO, Geneva, 2019.

[40] M. Domingo-Espin, J.M. Puigoriol-Forcada, A.A. Garcia-Granada, J. Llumà, S. Borros, G. Reyes, Mechanical property characterization and simulation of fused deposition modeling polycarbonate parts, Materials & Design 83 (2015) 670-677.

[41] M.R. Khosravani, E. Fereiduni, Anisotropic mechanical behavior modeling of 3D printed parts using FEM and orthotropic constitutive models, Materials & Design 219 (2022) 110825.

[42] J.M. Gere, B.J. Goodno, Mechanics of Materials, eighth ed., Cengage Learning, Stamford, 2013.

[43] L.J. Gibson, M.F. Ashby, Cellular Solids: Structure and Properties, second ed., Cambridge University Press, Cambridge, 1997.

[44] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: Proceedings of the 3rd International Conference on Learning Representations, 2015.

[45] X. Glorot, Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, in: Proceedings of the 13th International Conference on Artificial Intelligence and Statistics, 2010, pp. 249-256.

[46] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, Journal of Machine Learning Research 15 (2014) 1929-1958.

[47] Dassault Systèmes, Abaqus 2021 User's Manual, Dassault Systèmes Simulia Corp., Providence, RI, 2021.

[48] A. Der Kiureghian, O. Ditlevsen, Aleatory or epistemic? Does it matter?, Structural Safety 31 (2009) 105-112.

[49] R.E. Melchers, A.T. Beck, Structural Reliability Analysis and Prediction, third ed., John Wiley & Sons, Chichester, 2018.

[50] M.R. Khosravani, A. Zolfagharian, On the flammability and mechanical responses of abaca-reinforced composites: A machine learning-oriented approach, Smart Science 13 (2025) 60-76.

Acknowledgements

This research was supported by internal institutional funding. No external grant funding was received.

Declaration of Competing Interest

The authors declare no competing financial interests or personal relationships that could have influenced this work.

Data Availability

Experimental data, trained model weights, and UMAT source code are available from the corresponding author upon reasonable request.