

We would like to thank the reviewers for carefully reading the manuscript and their comments for improving the quality of this paper. The comments and responses are listed below, and the changes are shown in the marked version of the revised manuscript.

Responses to the **2nd reviewer's** comments:

- 1. The title and graphical abstract report the development of a digital twin with an active feedback loop for controlling the multi-zone heating system. In practice, the authors implement a digital model for forward prediction of resin cure without data feedback or control. The digital twin nature of the approach should be discussed.**

The digital twin referenced in this work is a digital representation of the VARIM process that establishes a mapping between process parameters and curing responses. The use of a machine learning approach is motivated by the high computational cost of physics-based simulations and the fast inference capability of trained ML models. As noted by the reviewer, the model is intended to serve as a rapid forward prediction tool, which is then used for curing optimization. Additional clarifications have been added in section 2.2 *A Machine Learning process model for the VARIM process* and section 4.1 *Machine Learning and Optimization*

- 2. The manuscript lacks critical details about the ML implementation, including dataset size, model architecture, training procedure, and hyperparameter selection. Since the code is not provided, it makes it challenging to reproduce the results. Additionally, potential limitations, such as generalization to unseen geometry or arbitrary initial conditions, are not discussed.**

Thank you for the comments. Detailed information about the ML implementation as well as the limitations were added to the methods section, which are also copied below. A key limitation of the proposed machine learning model is its dependence on the distribution of the training data. Model performance may degrade when evaluated under initial conditions, boundary conditions, or geometric configurations that were not represented during training, limiting its generalization to previously unseen scenarios. However, the digital twin framework can generate physically consistent reference solutions through its calibrated multiphysics model. This capability enables the system to produce reliable ground-truth data under new conditions and continuously refine the machine learning model, thereby improving predictive accuracy over time.

“The dataset employed for training the forward machine learning (ML) model is composed of physics-based, numerically simulated time-series data that describe the coupled evolution of temperature and resin degree of cure during a multizone composite curing process. Each simulation represents a complete curing cycle and is stored as a CSV file containing synchronized temporal histories of zone-specific temperatures and their corresponding curing responses. This physics-based simulation framework ensures that the data reflect realistic thermochemical behavior governed by heat transfer and cure kinetics, while allowing systematic exploration of process parameters that would be costly or impractical to obtain experimentally.

For the validation study, a heated bed consisting of three independently controlled zones (denoted A, B, and C) is considered. Each zone is assigned a constant temperature setpoint, treated as an independent design variable. The setpoints for all zones vary from 50 °C to 80 °C in increments of 1 °C, resulting in 29,791 unique combinations and corresponding simulations. The selected temperature range reflects manufacturing constraints: temperatures above 80 °C are known to accelerate tool wear and degrade equipment longevity, whereas temperatures below 50 °C are insufficient to reliably initiate or sustain resin curing reactions. This parameter space therefore captures the operationally relevant regime for process optimization.

The raw simulation outputs consist of time-series data with variable sequence lengths due to differences in cure progression. To enable batch-based sequence learning, all sequences are parsed, padded to a uniform length, and assembled into three-dimensional tensors of the form (samples × time steps × features). Prior to training, temperature inputs from all zones are normalized to a common numerical scale, which reduces bias introduced by differing magnitudes across features and improves numerical conditioning during optimization. The output variables, representing the degree of cure, are scaled using an identical normalization strategy to maintain consistency between input and output spaces. The complete dataset is randomly partitioned into training (85%) and testing (15%) subsets to enable unbiased model validation and performance evaluation.

A key limitation of the proposed machine learning model is its dependence on the distribution of the training data. Model performance may degrade when evaluated under initial conditions, boundary conditions, or geometric

configurations that were not represented during training, limiting its generalization to previously unseen scenarios. However, the digital twin framework can generate physically consistent reference solutions through its calibrated multiphysics model. This capability enables the system to produce reliable ground-truth data under new conditions and continuously refine the machine learning model, thereby improving predictive accuracy over time.

The forward ML model is implemented in Python using the Keras high-level API with TensorFlow as the computational backend. The model is formulated as a sequential neural network explicitly designed to capture both static, nonlinear relationships between temperature and degree of cure and the temporal dependencies inherent in cure kinetics. The architecture begins with a fully connected (dense) layer consisting of 10 neurons, incorporating L1 regularization and a LeakyReLU activation function. This layer serves to project the normalized temperature inputs into a higher-level feature space while mitigating overfitting through sparsity-promoting regularization.

Following this initial feature extraction, a reshaping operation expands the input representation along the temporal dimension to align with the desired output sequence length. This enables subsequent recurrent layers to process the data as a time-dependent signal. At the core of the network is a Long Short-Term Memory (LSTM) layer with 100 hidden units and L1 regularization. The LSTM architecture is critical for modeling cure behavior because it explicitly retains information from previous time steps through gated memory cells, allowing the network to learn how prior temperature histories influence the current and future state of cure. This capability is essential for accurately capturing the path-dependent nature of thermoset curing, which cannot be represented adequately by purely feedforward networks.

The temporal features extracted by the LSTM layer are further refined through a dense layer with 50 units and L1 regularization, followed by an additional dense layer with 100 units and a LeakyReLU activation function. These layers enhance the nonlinear representational capacity of the network. To preserve the sequential structure of the data, all post-LSTM dense layers are wrapped in a TimeDistributed framework, which applies the same transformation independently at each time step. This design enables the model to generate a continuous curing trajectory rather than a single aggregate prediction. The final TimeDistributed dense layer outputs zone-wise degree-of-cure predictions at every time step for all three heating zones.

Model training is conducted using the Adam optimizer with a mean squared error (MSE) loss function, which is well suited for continuous regression tasks involving time-series outputs. L1 regularization is applied uniformly across layers with a coefficient of 0.1 to reduce model complexity and improve generalization. The optimizer is initialized with a learning rate of 0.01, which is adaptively reduced using a ReduceLROnPlateau strategy to improve convergence as training saturates. Specifically, the learning rate is reduced by a factor of 0.5 after 100 epochs without improvement, with a minimum allowable learning rate of 0.001.

To further mitigate overfitting and unnecessary computation, early stopping is employed with a minimum improvement threshold of 0.1 and a patience of 500 epochs. Training is performed with a batch size of 32 and a maximum cap of 100,000 epochs. In practice, the combined use of adaptive learning rate reduction and early stopping results in consistent convergence after approximately 15,000 epochs, indicating stable learning behavior across different initializations and training runs.”

3. The Nelder–Mead algorithm is used for optimization, but its limitations (e.g., local search, sensitivity to initialization) are not discussed. No comparison is provided with alternative approaches, such as simulation-only optimization or modern derivative-free methods like Bayesian optimization.

Thank you for the comment. We added a section discussing the choice of optimization method and this section is copied below.

“This objective function is minimized using the Nelder–Mead algorithm, a gradient-free optimization method well suited for non-differentiable or noisy problems. Compared with alternatives such as Bayesian optimization, which incurs substantial computational overhead due to probabilistic surrogate construction and kernel updates, Nelder–Mead offers an efficient and fast approach that is particularly well suited for the low-dimensional design space of three temperature variables and is crucial when the objective function will be scaled up to optimize wind turbine blades with greater than 100 heating zones. Since each evaluation of the ML digital twin takes only fractions of a second, the advantages of Bayesian methods in minimizing the number of expensive evaluations do not apply here, while methods like evolutionary algorithms or particle swarm optimization would introduce unnecessary stochasticity and significantly higher computational cost. Consequently, the Nelder–Mead algorithm provides the most efficient balance of robustness, simplicity, and speed for exploring the curing design space in this ML-accelerated optimization framework.”

4. The manuscript lacks a quantitative assessment of the results compared to appropriate baselines. For example, it would be valuable to compare the time gains and mechanical properties achieved with ML-based optimization versus those achieved with the conventional optimization approach.

Thank you for your comments. As noted in our response to the previous comment, a major limitation of conventional optimization approaches, most of which are gradient-based, is the need to evaluate gradients. This requirement is not feasible for the present application. This limitation motivates the use of a gradient-free approach, which offers significant computational savings while still delivering strong optimization performance.

The baseline used to assess the benefits of the proposed method is based on the conventional heating profile dictated by the slowest curing cycle.

5. Abstract, line 20: the 12.5 % curing time reduction and improved curing uniformity are not supported by experimental evidence in the manuscript, nor discussed. The authors should consider removing this claim or providing the supporting evidence.

Thank you for your comments. We have added supporting evidence based on the experiments. The following section is added:

“4.2 Curing time improvement

The optimized multizone method allows for a unique temperature setpoint in each zone, enabling the use of higher temperature setpoint in thicker and less thermally conductive regions. This targeted heating improves the overall curing cycle time. Using the multiphysics simulations as a benchmark, it was observed that the conventional methods take 631.8 mins to reach saturation while the multizone approach requires 552.8 minutes which amount to roughly 12.5% increase in curing cycle time. Since the cure degree value approaches to 100% asymptotically, saturation is considered to be reached when the rate change of cure degree is less than 0.01% per minute. The individual graphs for each zone in conventional and multizone methods along with the curing time are shown side by side for comparison in Figure 8. The red dashed line indicates the saturation point.

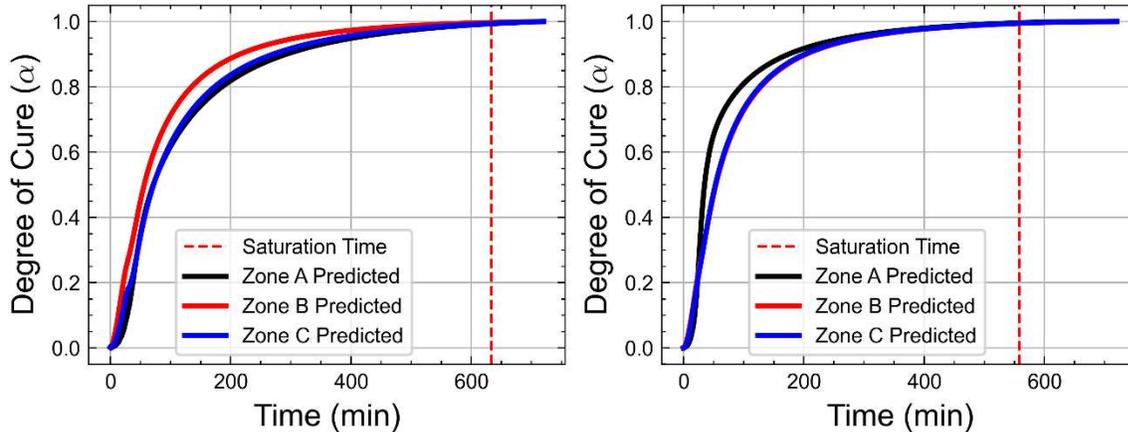


Figure 8: Multiphysics simulations with conventional (Left) and multizone (Right) setups along with degree of cure saturation point showing an improvement in the curing cycle time.

The curing cycle time improvement was also validated experimentally. The temperature values from the thermocouples were input in equation (7) to output the curing progress, then normalized to provide the degree of cure shown below in Fig 9.

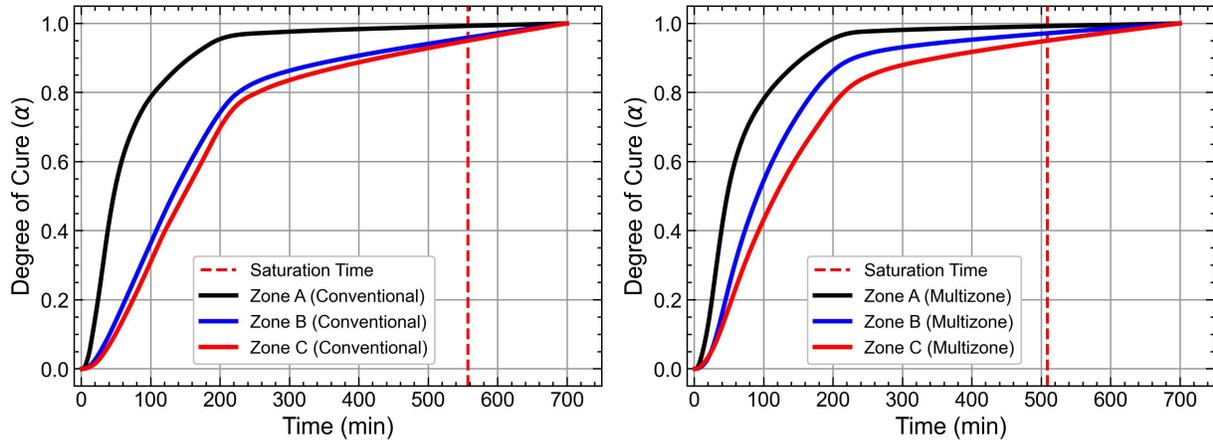


Figure 9: Degree of cure curves from composite manufacturing with conventional (Left) and multizone (Right) setups along with degree of cure saturation point showing an improvement in the curing cycle time.

In the experiment the curing cycle improves from 202.5 min to 185.5 min which results in an 8.4% improvement in the curing cycle time. The difference between the ideal value of 12.5% from the multiphysics simulation and the experimental value can be accounted for by experimental and human-induced variability in the manufacturing procedure. Additionally, the experimental graphs do not consider the vacuum setup time and the bed heating time which may account for the time difference between the simulation and the experiment during which the resin is assumed to have no significant curing.

6. Line 108: "To accurately simulate these coupled processes...". What is the definition of an accurate result? The author could provide references to supporting literature.

To ensure accuracy, the model is first calibrated based on known material database and additional experiments performed by the group. The model is then further validated by comparing to measurements (such as infusion front and temperature) made from VARIM experiments. Accuracy is considered when the model prediction matches the experimental values within $\pm 10\%$ error range.

7. Table 1: Consider reporting the values used in this work.

Table 1 was updated as follows to include all values. These parameters are sufficient for anyone who has an interest in replicating this work. Note that the values of the pre-exponential coefficient and the activation energy are functions of the degree of cure and are shown in Figure 1a and b.

Material	Model parameters	Value	Source of data
Epoxy resin	Density	1150 kg/m ³	Vendor material datasheet
	Coefficient of thermal conductivity	0.35 W/m.K	Vendor material datasheet
	Newtonian viscosity	0.28 N.s/m ²	Vendor material datasheet
	Specific heat	2500 J/kg.K	DSC measurements
	Enthalpy	450000 J/kg	DSC measurements
	Pre-exponential coefficient (C) for kinetic cure model	Fig 1 (a)	DSC measurements

	Activation energy/Universal Gas constant (E/R)	Fig 1 (b)	DSC measurements
Glass fiber	Density	2520 kg/m ³	Vendor material datasheet
	Permeability along primary axis	2×10 ⁻⁹ m ²	In-house lab experiments
	Permeability perpendicular to primary axis	8×10 ⁻¹⁸ m ²	In-house lab experiments
	Thermal conductivity	0.6 × W/m.K	Vendor material datasheet
	Specific heat	840 J/kg.K	Vendor material datasheet
Foam	Density	60 kg/m ³	Vendor material datasheet
	Thermal conductivity	0.029 W/m.K	Vendor material datasheet

Table 1: A summary of values used in the multiphysics solver and their sources

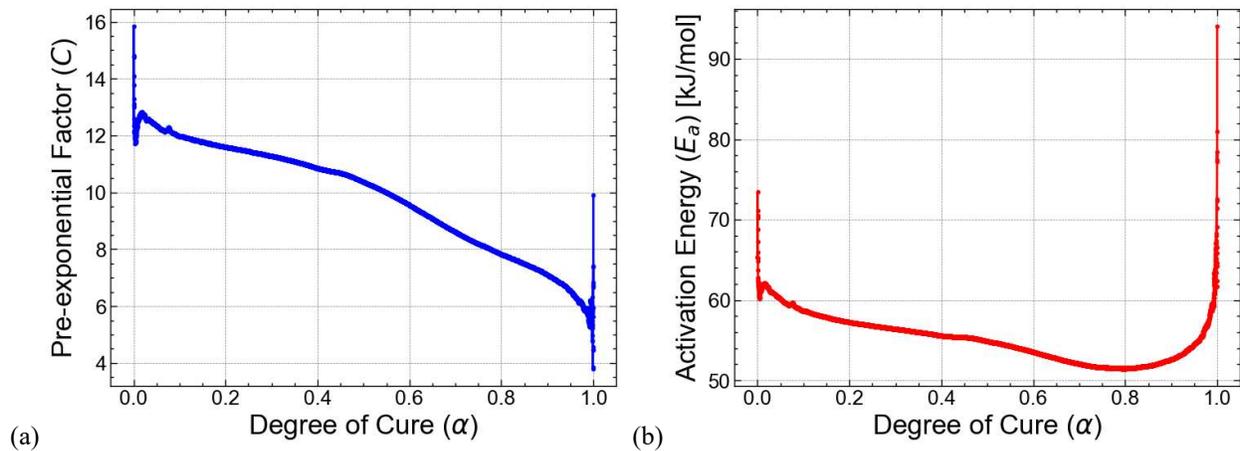


Figure 1: Calibrated kinetic cure parameters and model validation (a) Pre-exponential factor vs degree of cure, (b) Activation energy vs degree of cure.

8. Line 154: "calibrated multiphysics model". The authors reported the calibration of the cure kinetics only.

Table 1 now contains all the values used in the multiphysics simulation.

9. Line 157: The manuscript mentions that CNN-RNN models were chosen, but the authors do not report a CNN.

Edited out. CNN was initially a part of the model but was later removed.

10. Line 191: Consider providing a schematic to help the reader understand the data structure.

The following image along with description was added to help the reader understand the data and the flow of information in the ML model. We also added additional comments in the text under section 2.2 *A Machine Learning model for the VARIM process*, which are copied below.

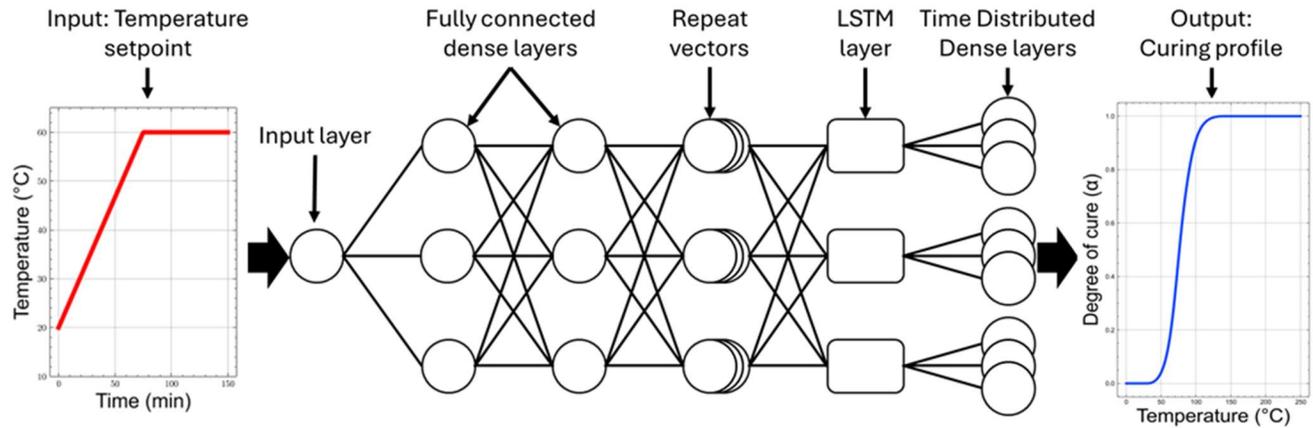


Figure 3: The ML model for the VARIM process. An ML-based prediction model based on trained LSTM neural network is employed. The model takes the temperature setpoint as a scalar input (left) and outputs the curing profile (right). The LSTM node is used in conjunction with fully connected dense layers, repeat vectors, and time-distributed dense layers to realize the prediction.

“The dataset employed for training the forward machine learning (ML) model is composed of physics-based, numerically simulated time-series data that describe the coupled evolution of temperature and resin degree of cure during a multizone composite curing process. Each simulation represents a complete curing cycle and is stored as a CSV file containing synchronized temporal histories of zone-specific temperatures and their corresponding curing responses. This physics-based simulation framework ensures that the data reflect realistic thermochemical behavior governed by heat transfer and cure kinetics, while allowing systematic exploration of process parameters that would be costly or impractical to obtain experimentally.

For the validation study, a heated bed consisting of three independently controlled zones (denoted A, B, and C) is considered. Each zone is assigned a constant temperature setpoint, treated as an independent design variable. The setpoints for all zones vary from 50 °C to 80 °C in increments of 1 °C, resulting in 29,791 unique combinations and corresponding simulations. The selected temperature range reflects manufacturing constraints: temperatures above 80 °C are known to accelerate tool wear and degrade equipment longevity, whereas temperatures below 50 °C are insufficient to reliably initiate or sustain resin curing reactions. This parameter space therefore captures the operationally relevant regime for process optimization.”

11. Line 206: Critical details are missing: what library is used to implement the network? Is the code available to ensure reproducibility? What is the dataset size? Is a validation dataset used? What is the data split?

Thank you for the comment. The following information was added to Section 2.2 *A Machine Learning model for the VARIM process* which includes the library, information of the dataset as well as further details of the hyperparameters of the ML model. The code can be provided upon request.

“The raw simulation outputs consist of time-series data with variable sequence lengths due to differences in cure progression. To enable batch-based sequence learning, all sequences are parsed, padded to a uniform length, and assembled into three-dimensional tensors of the form (samples × time steps × features). Prior to training, temperature inputs from all zones are normalized to a common numerical scale, which reduces bias introduced by differing magnitudes across features and improves numerical conditioning during optimization. The output variables, representing the degree of cure, are scaled using an identical normalization strategy to maintain consistency between input and output spaces. The complete dataset is randomly partitioned into training (85%) and testing (15%) subsets to enable unbiased model validation and performance evaluation.

A key limitation of the proposed machine learning model is its dependence on the distribution of the training data. Model performance may degrade when evaluated under initial conditions, boundary conditions, or geometric configurations that were not represented during training, limiting its generalization to previously unseen scenarios. However, the digital twin framework can generate physically consistent reference solutions through its calibrated multiphysics model. This capability enables the system to produce reliable ground-truth data under new conditions and continuously refine the machine learning model, thereby improving predictive accuracy over time.

The forward ML model is implemented in Python using the Keras high-level API with TensorFlow as the computational backend. The model is formulated as a sequential neural network explicitly designed to capture both static, nonlinear relationships between temperature and degree of cure and the temporal dependencies inherent in cure kinetics. The architecture begins with a fully connected (dense) layer consisting of 10 neurons, incorporating L1 regularization and a LeakyReLU activation function. This layer serves to project the normalized temperature inputs into a higher-level feature space while mitigating overfitting through sparsity-promoting regularization.

Following this initial feature extraction, a reshaping operation expands the input representation along the temporal dimension to align with the desired output sequence length. This enables subsequent recurrent layers to process the data as a time-dependent signal. At the core of the network is a Long Short-Term Memory (LSTM) layer with 100 hidden units and L1 regularization. The LSTM architecture is critical for modeling cure behavior because it explicitly retains information from previous time steps through gated memory cells, allowing the network to learn how prior temperature histories influence the current and future state of cure. This capability is essential for accurately capturing the path-dependent nature of thermoset curing, which cannot be represented adequately by purely feedforward networks.

The temporal features extracted by the LSTM layer are further refined through a dense layer with 50 units and L1 regularization, followed by an additional dense layer with 100 units and a LeakyReLU activation function. These layers enhance the nonlinear representational capacity of the network. To preserve the sequential structure of the data, all post-LSTM dense layers are wrapped in a TimeDistributed framework, which applies the same transformation independently at each time step. This design enables the model to generate a continuous curing trajectory rather than a single aggregate prediction. The final TimeDistributed dense layer outputs zone-wise degree-of-cure predictions at every time step for all three heating zones.

Model training is conducted using the Adam optimizer with a mean squared error (MSE) loss function, which is well suited for continuous regression tasks involving time-series outputs. L1 regularization is applied uniformly across layers with a coefficient of 0.1 to reduce model complexity and improve generalization. The optimizer is initialized with a learning rate of 0.01, which is adaptively reduced using a ReduceLROnPlateau strategy to improve convergence as training saturates. Specifically, the learning rate is reduced by a factor of 0.5 after 100 epochs without improvement, with a minimum allowable learning rate of 0.001.

To further mitigate overfitting and unnecessary computation, early stopping is employed with a minimum improvement threshold of 0.1 and a patience of 500 epochs. Training is performed with a batch size of 32 and a maximum cap of 100,000 epochs. In practice, the combined use of adaptive learning rate reduction and early stopping results in consistent convergence after approximately 15,000 epochs, indicating stable learning behavior across different initializations and training runs.”

12. Line 275: How many simulations were performed on the geometry? Do the simulations include the optimal configuration?

Number of simulations for training is mentioned in section 2.2 *A Machine Learning model for the VARIM process*. The optimal configuration corresponds to the set temperature of 67.77, 62.77, and 60 °C. This configuration was not a part of the original dataset for training.

13. Figures 6 and 7: Check the curing time unit. If the unit is correct, the manufacturing settings seem unsuitable for an industrial process. Please, discuss.

The units were incorrect and are now fixed.

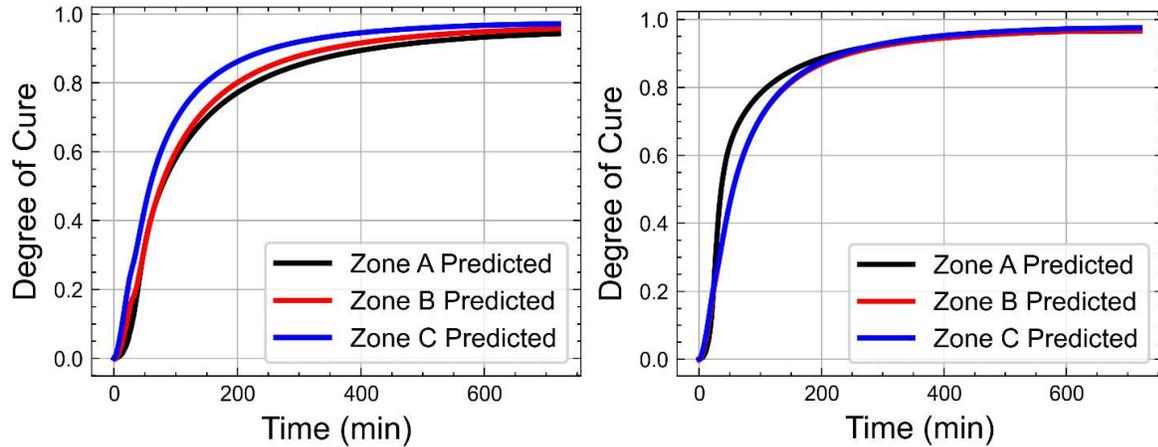


Figure 7: Degree of cure plots of the three zones for the conventional method (left) vs the optimized method (right).

14. Figure 6: The machine learning model matches well with the simulations, which is reasonable, especially if the optimized values are within the domain of the training data. The authors could strengthen the manuscript by characterizing the experimental degree of cure, for instance, by performing DSC on cutout samples and measuring the residual enthalpy.

The resin cure kinetics were validated using DSC data in Figure 1. For validating the degree of cure from the experiments, we retrieve the surface temperatures, input it into the calibrated resin cure kinetics model and normalize the data to acquire the degree of cure. This was done in post-processing to validate the curing cycle time improvement. There is no direct method to compare the ML model's predictions as the setup does not include a degree of cure sensor.

15. Figure 8: Please annotate the figure and describe the testing setup.

All figures have been annotated and the mechanical testing setup has been described with appropriate ASTM standards followed.

16. Table 2, Figure 10, line 335:

a. The manuscript requires additional data points to support the claims, particularly given the variability.

The mechanical validation section was rewritten to address the reviewers comments. To get a better understanding of the mechanical properties and from literature survey, we determined that bending tests are appropriate to determine the change in mechanical properties. Hence, new samples were created using 'conventional' methods and our optimized 'multizone' method and compared.

b. Please indicate how the yield strength, compressive modulus, and ultimate strength were derived from the experimental results.

We have now provided detailed explanation of the mechanical testing procedures, the ASTM standard followed as well as the equations used to determine the mechanical properties of the composite.

c. The cure uniformity was not characterized.

To experimentally characterize cure uniformity, it would be necessary to interrupt the curing process at multiple time points, extract samples from the manufactured parts, and evaluate both the degree of cure and the associated mechanical properties. Such efforts are beyond the scope of this manuscript. Our calibrated physics-based model and ML-driven digital twin already provide predictions of non-uniformity, enabling process optimization to mitigate these effects.

d. Figure 10 seems to be truncated.

Figure 10 has been corrected.

The following section now replaces the mechanical testing and validation section to better present the effects of multizone curing:

“4.3 Experimental validation

The structural performance of the composites manufactured by the two approaches were evaluated to determine the influence of curing strategy subjected to bending loads representative of wind turbine blade structures. Flexural testing was selected because bending is the dominant loading condition experienced by such structures during operation. Four-point bending tests were conducted on foam sandwich and glass fiber laminate specimens from Zones A and C respectively were cut using water jet cutting to characterize their mechanical response under controlled loading conditions. All testing procedures followed the standard method specified by ASTM International. The specimens in their setup are shown in Fig 10.

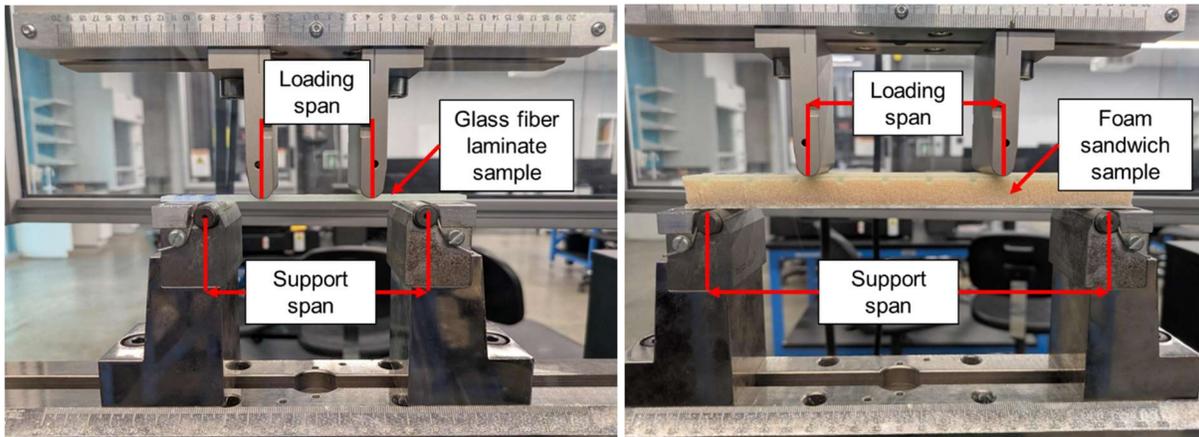
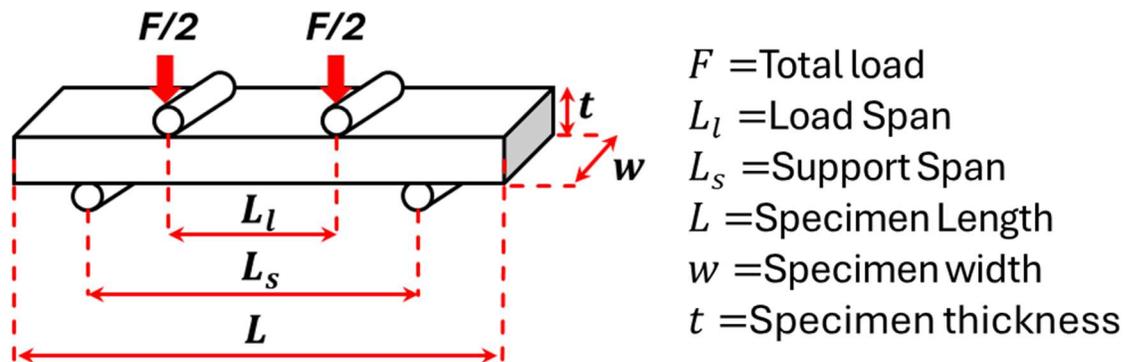


Figure 10: Schematic of the 4-point bending setup (Top). Test performed on glass fiber laminate (Bottom left) and foam sandwich (Bottom right) showing the loading and support spans to ensure pure bending.

The foam sandwich composites from Zone A were in accordance with ASTM standard D7249 (D30 Committee, ASTM International, 2020). The specimens were approximately 14.7 mm thick, 32.8 mm wide, and a support span of 176.4 mm with 20% length overhand on both sides. 4 such samples were cut from each specimen. The load span was half of the support span at 88.2 mm.

For glass fiber laminates from Zone C, 6 samples from conventional and multizone procedures each were cut. The specimen were approximately 12.8 mm wide, 2.3 mm thick, and 140 mm long samples were cut and tested in accordance with ASTM D7264 (D30 Committee, ASTM International, 2021) test method. A support span of 40:1 was chosen, i.e., 92 mm, to ensure failure due to pure bending and to observe compression on top laminate layers and tension in the bottom layers and with load span being half of the support span. The strain rate was determined using the ASTM standard D6272 (D20 Committee, ASTM International, 2017) formula:

$$R = \frac{0.167ZL^2}{d} \quad (13)$$

Where R is the rate of crosshead displacement in mm/min, L is the support span in mm, d is the depth of the beam in mm and Z is the rate of straining of the outer fibers which was kept at a minimum value of 0.01 mm/mm.

The flexure stiffness for both the specimens was calculated using equation (14) between strain values of 0.001 and 0.003 as specified by the ASTM standards.

$$E = \frac{\Delta\sigma}{\Delta\varepsilon} \quad (14)$$

E = flexural chord modulus of elasticity in MPa, $\Delta\sigma$ = difference in flexural stress between the two selected strain points in MPa, and $\Delta\varepsilon$ = difference between the two selected strain points.

Similarly, the formula for flexural strength is defined in equation (15)

$$\sigma = \frac{3PL}{4bh^2} \quad (15)$$

where: σ = stress at the outer surface in the load span region in MPa, P = applied force in N, L = support span in mm, b = width of beam in mm, and h = thickness of beam in mm.

The average value with errors for the conventional and multizone setup is shown in Fig 10. Table 2 presents the individual values of each specimen.

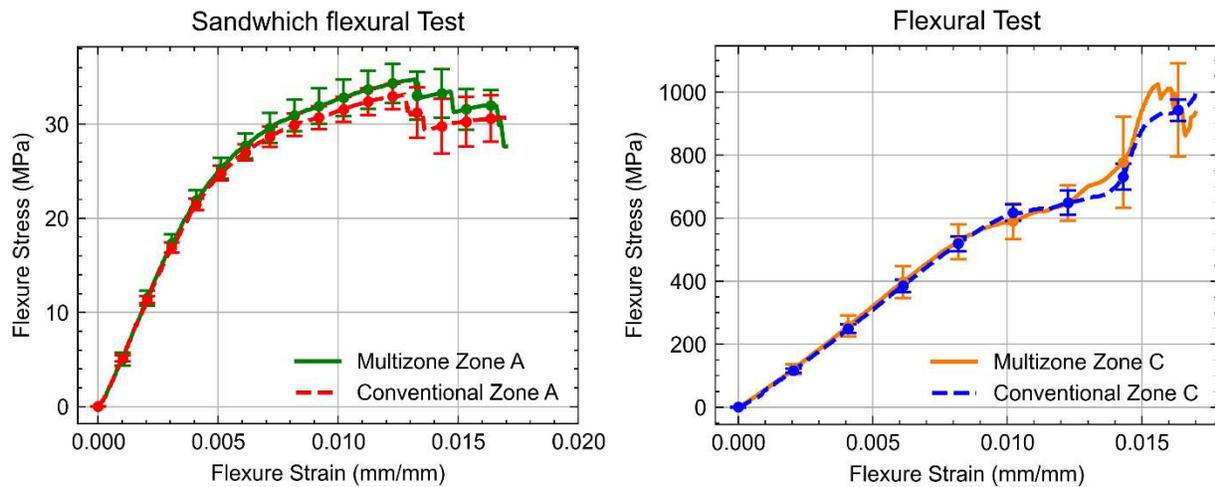


Figure 11: Average and standard deviation of stress vs strain curves for foam sandwich composites (Left) and glass fiber laminates (Right).

	Flexural (MPa)	Stiffness	Percentage difference	Flexural strength (MPa)	Percentage difference
Foam sandwich composite using conventional method	$5,693.85 \pm 146.77$		4.68 %	179.66 ± 72.02	4.68 %

Foam sandwich composite using multizone method	5,960.60 ± 189.97		188.06 ± 75.79	
Glass fiber composite using conventional method	61,747.16 ± 3,476.47	4.97 %	1246.32 ± 4.20	3.21 %
Glass fiber composite using multizone method	64,815.55 ± 9,213.66		1286.32 ± 4.51	

Table 2: A summary of flexural modulus and flexural strength calculated from the mechanical testing

The flexural test results indicate a consistent 4–5% increase in both the modulus of elasticity and flexural strength for specimens produced using the optimized multizone curing strategy. This improvement suggests that the multizone approach not only matches but can exceed the mechanical performance achieved through the conventional uniform-temperature process. The enhancement is likely attributable to the reduced thermal gradients and improved cure uniformity associated with multizone temperature control, which mitigates defect formation and promotes more homogeneous material properties. As a result, the composites produced under optimized conditions exhibit greater resistance to bending-induced stresses, delaying the onset of failure and improving overall structural reliability.”

17. Line 113: "and $\Delta T(t)$ The". Uppercase.

Fixed.

18. Line 126, Figure 1: the SI symbol for Kelvin is K, and °C for degree Celsius.

Fixed

19. Line 312: Compliance calibration is a common practice and may not require detailed information.

This section was removed and the appropriate ASTM standard was mentioned instead.