Modeling Dataset Development for Machine Learning Prediction of the Thermo-Chemical Curing Process in Advanced Composite Structures

Arash Ramian[1][0000-0001-8329-4937], J.M. Ashfiqur Rahman[1][0009-0006-7208-5439], Tian Zhao[1][0000-0001-6456-9763], and Rani Elhajjar[1][0000-0003-4778-5186]

[1]College of Engineering & Applied Science, University of Wisconsin-Milwaukee, Milwaukee, WI 53211, USA
{aramian,rahmanj,tzhao,elhajjar}@uwm.edu

Abstract. This study examines the curing behavior of composite laminates by combining finite element simulations with data-driven surrogate modeling. A cure kinetics model is implemented in ANSYS Composite Cure Simulation to generate a large parametric set of cure histories for representative laminate configurations. The design space covers a wide range of processing, material, and structural parameters, including heating ramps, dwell temperatures and times, laminate thickness, number of plies, total heat of reaction, symmetric and unsymmetric stacking sequences, and intentionally non-optimal cure cycles. In total, 1016 three-dimensional simulations are carried out. For each cure cycle, time-dependent histories of temperature, degree of cure, and maximum principal stress are extracted and collected into a comprehensive dataset. The input space is represented using compact descriptors of the cure cycle, laminate geometry, material properties, and layup statistics, while the model outputs the corresponding thermal, chemical, and stress evolution histories. A shared-task neural network surrogate is trained to learn the mapping from processing conditions and laminate characteristics to the cure response. To better capture rapid thermal transitions and exothermic behavior, the simulation histories are represented on a shared warped time grid with increased resolution in highly nonlinear regions. The trained model reproduces finite element predictions of temperature, degree of cure, and maximum principal stress across the explored design space with low computational cost. This forward surrogate provides a fast prediction tool for comparing cure-cycle responses, identifying conditions associated with incomplete cure or large thermal gradients, and supporting future physics-informed process design studies.

Keywords: Composite Structures · Curing Process · Simulation · Time-Temperature-Cure · Machine Learning

1 Introduction

Fiber-reinforced polymer (FRP) composites are now widely used in high-performance structures in aerospace, automotive, wind energy, and marine applications, where strict weight reduction targets must be satisfied together with high stiffness, strength, and durability requirements [1, 2]. In these systems, manufacturing is not only a downstream operation but also an important design variable, and the curing of thermoset matrices has a key influence on the final part quality. During cure, an exothermic reaction drives the cross-linking of the polymer network, while heat generation and heat conduction inside the laminate and between the part and the tooling produce non-uniform temperature fields and significant thermal gradients [3]. These spatial and temporal variations in degree of cure and modulus evolution, combined with cure shrinkage and tool–part interaction, lead to the development of residual stresses. At the structural scale, these coupled thermo-chemical–mechanical phenomena appear as process-induced deformations, such as spring-in, warpage, and distortion. These effects make dimensional control more difficult and can reduce the long-term mechanical performance of composite components [4, 5].

Despite these challenges, finite element analysis (FEA) remains the most widely used method for predicting the coupled thermo-chemical–mechanical response during curing [6–8]. It can resolve spatially varying temperature, degree of cure, and stress fields under realistic boundary conditions [9]. However, the governing heat equation is tightly coupled with non-linear cure kinetics, requiring fine temporal resolution and detailed through-thickness discretization to capture steep thermal gradients and rapid exothermic events [7]. These models are indispensable for detailed analysis, but their computational cost becomes prohibitive when the design space contains many laminate configurations and off-nominal cure profiles. Process window exploration, cure cycle tuning, and iterative design cycles may require hundreds or thousands of model evaluations, which are difficult to carry out within realistic limits of time and computing resources. This computational bottleneck also limits the use of physics-based cure models in higher-level frameworks for uncertainty quantification, robust design, or real-time process monitoring and control [10].

To overcome this limitation, we adopt a surrogate modeling strategy, in which a data-driven model is trained to emulate the thermo-chemical-mechanical response at a fraction of the cost of full finite element analysis. Recent work has applied artificial intelligence and machine learning techniques to composite processing, including artificial neural networks for cure kinetics approximation, Gaussian process models for process window identification, and deep learning architectures for online quality monitoring [11–13]. The central idea is to use the high-fidelity cure model as a generator of training data, systematically sampling the space of process parameters and laminate configurations to build a rich synthetic dataset of spatial–temporal temperature, degree-of-cure, and stress histories. This dataset represents the mapping function from cure cycle parameters and structural descriptors to the local thermal, kinetic, and stress response. A feed-forward neural network surrogate is then trained on this dataset to learn

that mapping and provide rapid inference of temperature, degree of cure, and maximum principal stress for new cure schedules and laminate designs.

The novelty of this work is that it considers generalization, physical fidelity, and computational efficiency at the same time. Unlike many previous surrogate modeling approaches that are restricted to fixed geometries or nominal conditions, the proposed model generalizes across variable laminate thicknesses and complex, multi-stage cure cycles. In addition, a specialized shared time-sampling strategy is used to capture sharp, non-linear exothermic peaks, so that the data-driven model keeps high physical fidelity even during critical reaction phases where standard approximations usually fail. By including maximum principal stress as an additional predicted response, the model also connects the thermo-chemical cure history to stress development during processing. In an industrial manufacturing context, this surrogate is intended to support cure-cycle and layup decisions that currently require repeated finite element runs. Once trained, the model returns full DoC, temperature, and maximum-principal-stress histories from a compact cure-cycle and laminate descriptor, so it can be used by process engineers for cure-cycle screening, process-window exploration, and early flagging of conditions associated with incomplete cure, large through-thickness thermal gradients, or elevated residual stress before autoclave trials. The dataset is built around an industrially representative T700GC/epoxy prepreg system and spans manufacturer-recommended, accelerated, conservative low-gradient, and step-accelerated cure cycles together with intentionally incomplete-cure cases, so that the surrogate is exposed to both nominal and off-nominal production scenarios. The present work establishes this capability for flat thick laminates space; extension to more complex part geometries, additional material systems, and coupling with optimization and uncertainty-quantification methods is the next step toward direct industrial deployment.

2 Theoretical Framework and Numerical Modeling

2.1 Governing Equations

To accurately model the cure kinetics, Differential Scanning Calorimetry (DSC) experiments were carried out following the methodology described in [14]. The kinetic model expresses the reaction rate as a function of temperature and the current degree of cure. Small uncured prepreg specimens (approximately 5 mg) were hermetically sealed in aluminum pans and exposed to both dynamic and isothermal temperature profiles. For the isothermal tests, the samples were heated rapidly to the target cure temperatures and then held at these levels to quantify the evolution of the reaction. The resulting heat flow data were used to calibrate the model parameters, so that the numerical framework can accurately predict the resin cure state under realistic manufacturing conditions. The transformation of data relied on the principle that reaction kinetics correlate directly with the measured heat flow rate, which can be expressed as [15]:

$$\frac{d\alpha}{dt} = \frac{1}{H_T}\frac{dH}{dt} \tag{1}$$

and

$$\alpha = \frac{1}{H_T} \int_0^1 \left(\frac{dH}{dt} \right) dt \tag{2}$$

The cure kinetics are described by an autocatalytic model of the form [15]:

$$\frac{d\alpha}{dt} = K\alpha^m (1-\alpha)^n \tag{3}$$

where m and n are reaction exponents capturing the autocatalytic contribution and depletion of reactive species, respectively, and $K(T)$ is the temperature-dependent rate constant. The rate constant follows an Arrhenius law:

$$K = Ae^{-\frac{E}{RT}} \tag{4}$$

where A is the pre-exponential factor, E is the activation energy, R is the universal gas constant, and T is the absolute temperature.

In this work, the matrix corresponds to a Toray T700GC/epoxy unidirectional prepreg system, calibrated previously using isothermal DSC tests. The fitted kinetic parameters are: $A = 1.3 \times 10^6$ s^{-1}, $E = 67,000$ Jmol^{-1}, $m = 0.4$, $n = 1.2$. These cure-kinetics relations are calibrated from experimental data and provide the thermo-chemical foundation for the finite element simulations. The resulting thermal and degree-of-cure fields are subsequently coupled within the ACCS finite element framework to obtain the stress response used as an additional machine-learning target in this study.

2.2 Finite Element Simulation

Finite element simulations were performed using ANSYS Workbench, using the ANSYS Composite Cure Simulation (ACCS) and Composite PrePost (ACP) modules [16] to generate high-fidelity training data for the machine learning model. The model geometry corresponds to a thick composite plate manufactured from Toray T700GC/epoxy prepreg (Figure 1). To capture the through-thickness thermal gradients and the evolution of the degree of cure with sufficient accuracy, the laminate was discretized using three-dimensional coupled-field solid elements with a minimum edge length of 1.0 mm, giving a mesh of approximately 7,000 elements and 7,300 nodes. The simulation was initialized at 20°C to represent ambient laboratory conditions. Convection boundary conditions were applied to all free surfaces to simulate heat transfer with the surrounding environment, while the internal heat generation from the exothermic curing reaction was fully coupled to the thermal field to ensure realistic feedback between the reaction rate and the temperature evolution.

To characterize the system's thermal response under the applied boundary conditions, a transient thermal plot was generated (Figure 2) to track the evolution of the glass transition temperature throughout the cure cycle. The curve shows a rapid increase during the primary cure stage, followed by a plateau as the material approaches full conversion. Capturing this progression is critical, because the evolution of the glass transition temperature governs the transition between viscoelastic material states and is a key driver for the development of residual stresses within the laminate.

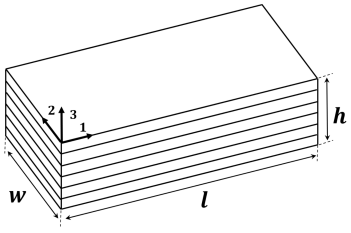

Fig. 1: Schematic of this study's laminated composite

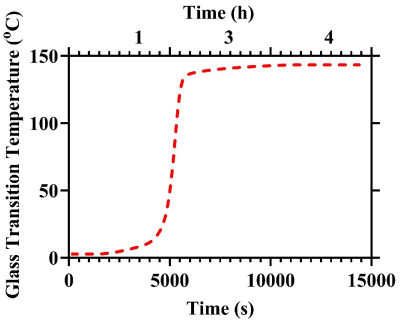

Fig. 2: Evolution of glass transition temperature with time during the cure cycle

2.3 Dataset Generation

To populate a representative design space for Machine Learning, a comprehensive parametric study was carried out using the calibrated ACCS model as a data generator. The geometric space includes laminates with different numbers of layers and ply thicknesses ranging from 0.1 mm to 0.25 mm, as well as diverse stacking sequences. Both symmetric and unsymmetric layups are considered to capture the effect of through-thickness stiffness imbalance and process-induced warpage. On the process side, the cure cycle is systematically varied by changing the main heating ramp rate and the structure of the isothermal holds, including both single-hold and double-hold profiles derived from industrially relevant cure schedules. Cure quality and material variability are explicitly included by combining nominal, well-cured cycles with deliberately under-cured scenarios, and the total heat of reaction is perturbed within a prescribed range to reflect stochastic variations in material properties. This parametric study produces a dataset with 1016 unique simulation cases, each representing a distinct combination of geometry, cure cycle, cure quality, and material characteristics. Table 1 shows the variables which were changed for generating the dataset.

To ensure the dataset encompasses a comprehensive range of thermal histories, the cure cycles were categorized into four distinct processing regimes, as illustrated in Figure 3. The Manufacturer Recommended profile functions as the standard industrial baseline, providing a balanced compromise between processing efficiency and part quality. The Accelerated Cure regime investigates the upper limits of manufacturing throughput by utilizing a single, rapid heating ramp to the final cure temperature, reflecting aggressive production schedules. The Conservative Low-Gradient profile uses reduced heating and cooling rates to minimize through-thickness thermal gradients, representing conditions for thick laminates where mitigation of residual stress is the main constraint. Finally, the Step-Accelerated cycle applies a multi-stage thermal profile with a rapid initial ramp to an intermediate dwell (around 80°C) before the final cure, testing the model's ability to capture viscosity evolution and gelation under non-linear,

Table 1: Summary of geometric, process, and material parameters used for dataset generation

Category	Variable	Range / Discrete Values
Geometry	Number of Plies	4, 8, 16, 32, 48
	Physical Thickness of Each Ply	0.1 mm – 0.25 mm
	Stacking Sequence	Symmetric & Unsymmetric
Process	Cure Profile Strategy	Manufacturer, Accelerated, Low-Gradient, Step-Accelerated
	Isothermal Holds	Single-hold vs. Double-hold
	Cure Completion	Nominal (Well-Cured) vs. Incomplete (Defect Scenarios)
Material	Heat of Reaction	$2.0 - 5.4 \times 10^5$ J/kg

stepped thermal excursions.

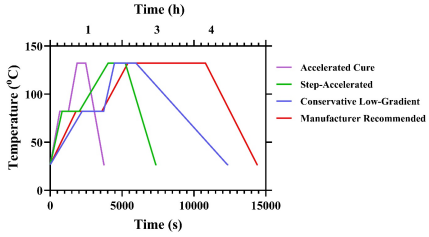

Fig. 3: Representative cure cycle temperature profiles from the generated simulation dataset, spanning the Manufacturer Recommended, Accelerated, Conservative Low-Gradient, and Step-Accelerated processing families

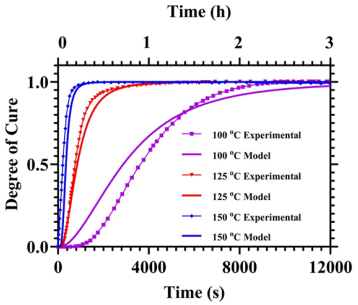

Fig. 4: Comparison between the degree of cure from experimental data and numerical simulations at 100°C, 125°C, and 150°C

3 Machine Learning Algorithm

3.1 Preprocessing

Figure 5 shows the preprocessing pipeline used to convert the raw FEA simulation data—generated from the validated cure-kinetics model (Figure 4)—into a fixed representation for learning. Each sample contains two parts: metadata that describes the laminate and cure schedule, and time-dependent simulation outputs for degree of cure, temperature, and stress. From these data, we extract a 25-dimensional input vector. These features summarize the cure schedule and laminate in a compact form. They include quantities such as start temperature, minimum and maximum schedule temperature, total cycle time, ramp rates, hold temperatures, hold durations, number of plies, number of holds, thickness, total

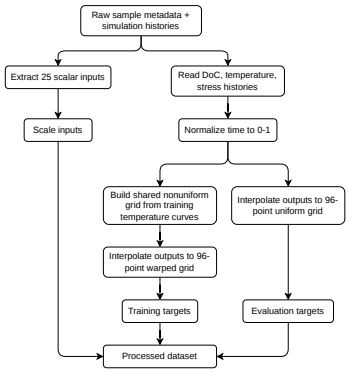

Fig. 5: Preprocessing pipeline used to convert raw simulation data into fixed-length model inputs and targets.

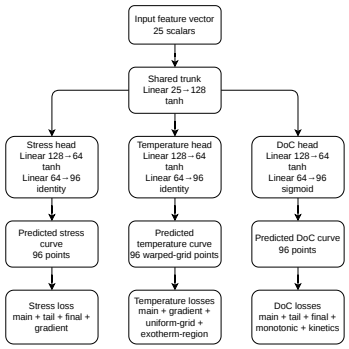

Fig. 6: Shared-task-head forward surrogate used to predict 96-point DoC, temperature, and stress histories from 25 scalar inputs

heat of reaction, symmetry, and angle-based laminate descriptors. The output histories are then read from the simulation files and mapped to normalized time over the interval [0,1]. This step removes differences in absolute cycle length and allows all cases to be represented on a common time basis.

To improve the representation of rapid thermal events, we do not rely only on uniform sampling. Instead, a shared nonuniform time grid is built from the training temperature histories. This grid places more resolution in regions where the temperature response changes more sharply, such as high-slope regions, high-curvature regions, and exothermic portions of the cure. The resulting warped grid is used to generate 96-point training targets. In parallel, a 96-point uniform grid is also retained for evaluation. As a result, the model is trained with a representation that is more informative near important thermal events, while evaluation remains easy to interpret on a standard grid. The final processed dataset therefore contains scaled 25-dimensional inputs together with fixed-length DoC, temperature, and stress histories for training and assessment.

The choice of 96 output points was made to balance temporal resolution and model size. A smaller output dimension may under-represent rapid heating, cooling, and exothermic regions, while a much larger output dimension increases the number of predicted quantities without necessarily improving interpretability. The warped 96-point grid therefore provides a compact representation that concentrates sampling density in thermally active regions, while the corresponding 96-point uniform grid provides a consistent basis for reporting errors across all simulations.

3.2 Forward model

Figure 6 presents the forward surrogate model. The network takes a 25-dimensional input vector and passes it through a shared trunk. This trunk

consists of a linear layer from 25 to 128 followed by a tanh activation. The shared representation is then sent to three separate task-specific heads. Each head has a linear layer from 128 to 64, a tanh activation, and a final linear layer from 64 to 96. The DoC head uses a sigmoid output so that the predicted cure remains within a physical range. The temperature and stress heads use identity outputs. This architecture produces three full response histories: a 96-point DoC curve, a 96-point temperature curve on the warped grid, and a 96-point maximum-principal-stress curve. The shared trunk allows the model to learn common structure across the outputs, while the separate heads preserve task-specific behavior.

The training objective is designed to improve both accuracy and physical consistency. The main loss includes weighted terms for DoC, temperature, and maximum principal stress prediction. Additional auxiliary terms are used to shape the output behavior more carefully. On the DoC side, the loss includes terms for the full curve, the tail region, the final cure value, monotonicity, and a mild kinetics-based regularization. On the temperature side, the loss includes the main temperature term together with gradient, uniform-grid, and exotherm-region penalties. On the stress side, the loss includes the main stress term together with tail-region, final-value, and gradient penalties. This design allows the model to learn a unified thermo-chemical-mechanical response while giving additional attention to physically important regions of the predicted histories.

3.3 Evaluation

The forward model was evaluated using both cross-validation and a held-out test set. Ten-fold cross-validation was used on the training split to measure consistency across the available samples, while a separate held-out test set was used to assess generalization on unseen cases. In the implementation, each fold is trained on the scaled 25-dimensional input features and evaluated on full response histories for degree of cure, temperature, and maximum principal stress. After cross-validation, a final model is trained on the full training split and tested on the held-out set. The evaluation pipeline also saves prediction files and case-wise comparison plots for later inspection.

We report standard regression metrics for the predicted histories, including MSE, RMSE, MAE, and R^2. These metrics are computed for DoC, temperature, and maximum principal stress. Temperature and stress errors are reported after conversion back to native physical units, while DoC is evaluated in its nondimensional form. This gives a direct measure of numerical fit while keeping the results interpretable in engineering terms.

In addition to these global metrics, we use a benchmark suite designed for cure modeling. The benchmark includes temperature-gradient error, temperature-curvature error, and DoC-gradient error. It also evaluates specific temperature regions that are important in practice, including the early-ramp region, the exotherm region, and a window around the peak temperature. Peak temperature value and peak timing are measured separately, and the DoC response is further assessed using final-cure error and the predicted time to reach key DoC thresholds of 0.50, 0.70, 0.80, and 0.90. These metrics are included because low overall error

does not guarantee that the model captures the thermal and cure behavior that matters for process design.

The evaluation is performed on a common 96-point uniform normalized-time representation. Temperature is trained on the shared warped grid, which places more resolution in thermally active regions, and the predicted temperature history is then mapped back to the uniform grid for evaluation. This allows the model to benefit from a more informative training representation while keeping the reported performance easy to interpret.

Overall, this evaluation setup is intended to test more than curve fitting alone. A useful forward model for composite cure prediction must capture the timing of cure progression, the evolution of temperature during early heating and exotherm, the maximum-principal-stress response, and the final cure state with sufficient accuracy to support downstream engineering analysis. For that reason, the reported results combine standard regression metrics with cure-specific benchmarks and case-wise visual comparisons.

4 Results and Discussion

4.1 Model Calibration

The calibrated cure-kinetics model was validated against DSC-derived degree-of-cure data at 100°C, 125°C, and 150°C, as shown in Figure 4. The comparison shows close agreement between the experimental measurements and the numerical cure model over the temperature range used for parameter fitting. In each isothermal case, the degree of cure increases rapidly during the main reaction stage and then approaches a plateau as the reaction nears completion.

This agreement supports the use of the calibrated autocatalytic kinetics model in the subsequent ACCS simulations. Since the finite element dataset is generated from this kinetics model, the calibration step provides the physical basis for the synthetic cure histories used to train the machine-learning surrogate. The calibration result is therefore used to validate the thermo-chemical input model, while the surrogate-model accuracy is evaluated separately using cross-validation and held-out test results in the following sections.

4.2 Surrogate Model Performance

The following table summarizes the prediction accuracy across all temperature profiles, showing the best-case, worst-case, and average RMSE (Table 2) for both degree of cure and temperature predictions. The results demonstrate that the neural network achieves consistently low prediction errors across all profiles. Temperature RMSE values are presented in normalized scale [0,1]. DoC predictions show sub-1% RMSE in most cases, while temperature predictions maintain excellent accuracy even for aggressive heating profiles with rapid temperature changes.

4.3 Architecture Comparison

To justify the selected surrogate architecture, we compared the proposed shared-task-head MLP with three sequence-decoder alternatives: GRU, LSTM, and

Table 2: Prediction accuracy summary across all temperature profiles

Profile	Variations	Best DoC RMSE	Best Temp RMSE	Worst DoC RMSE	Worst Temp RMSE	Avg DoC RMSE	Avg Temp RMSE
Accelerated Cure Cycle	Var1: 2-Hold / Fully Cured	0.0049	0.291	0.0372	3.619	0.0113	0.793
	Var2: 2-Hold / Incomplete Cure	0.0036	0.295	0.0279	3.488	0.0056	0.748
	Var3: 1-Hold / Fully Cured	0.0098	0.579	0.1633	23.518	0.0407	5.967
	Var4: 1-Hold / Incomplete Cure	0.0081	0.368	0.0084	2.331	0.0093	1.141
Fast Heat Fast Cool Cycle	2-Hold / Fully Cured	0.0067	0.365	0.0127	1.558	0.0081	0.718
Manufacturer Recommended Cycle	Var1: 2-Hold / Fully Cured	0.0085	0.310	0.0119	2.036	0.0105	0.765
	Var2: 2-Hold / Incomplete Cure	0.0066	0.318	0.0094	1.179	0.0089	0.655
	Var3: 1-Hold / Fully Cured	0.0092	0.576	0.0116	2.379	0.0107	1.137
	Var4: 1-Hold / Incomplete Cure	0.0078	0.347	0.0187	8.322	0.0095	1.238
Conservative Low-Gradient Cycle	Var1: 2-Hold / Fully Cured	0.0077	0.431	0.0605	25.161	0.0194	3.095
	Var2: 2-Hold / Incomplete Cure	0.0037	0.294	0.0365	2.770	0.0085	0.922
	Var3: 1-Hold / Fully Cured	0.0052	0.459	0.0277	17.043	0.0116	2.135
	Var4: 1-Hold / Incomplete Cure	0.0074	0.347	0.0116	2.019	0.0092	0.796
Step-Accelerated Cycle	Var1: 2-Hold / Fully Cured	0.0079	0.216	0.2346	23.529	0.0252	3.161
	Var2: 2-Hold / Incomplete Cure	0.0049	0.299	0.5277	30.225	0.0277	2.270
	Var3: 1-Hold / Fully Cured	0.0097	0.487	0.1953	15.663	0.0264	2.486
	Var4: 1-Hold / Incomplete Cure	0.0043	0.337	0.0133	2.007	0.0076	0.782

Temporal CNN. All models used the same 25-dimensional input vector, 96-point output representation, train/test split, and evaluation metrics. Since the model input is a case-level descriptor vector rather than a temporal input sequence, the recurrent and convolutional models were implemented as decoders that map each case descriptor to the full response histories. The comparison is summarized in Table 3. The GRU decoder gives a slightly lower cross-validation temperature RMSE, but the proposed model gives the best held-out DoC, temperature, and stress RMSE. It also requires much less training time than the sequence-decoder models. Therefore, the shared-task-head MLP was retained as the main surrogate model.

Table 3: Architecture comparison using the same dataset, preprocessing workflow, train–test split, and evaluation metrics. DoC is nondimensional, temperature RMSE is reported in °C, stress RMSE is reported in native simulation units, and training time is reported in seconds.

Architecture	CV DoC RMSE	CV Temp. RMSE	CV Stress RMSE	Test DoC RMSE	Test Temp. RMSE	Test Stress RMSE	Training Time (s)
Proposed model	**0.0356**	3.215	**10.376**	**0.0295**	**3.203**	**9.177**	**55.8**
GRU decoder	0.0405	**3.199**	12.703	0.0333	3.416	11.434	823.6
LSTM decoder	0.0436	3.392	12.463	0.0360	4.024	11.260	376.2
Temporal CNN decoder	0.0441	3.612	14.497	0.0347	3.714	13.687	844.0

4.4 Sensitivity to Simulation Parameters

To further examine how the trained surrogate responds to the input simulation parameters, a permutation-based sensitivity analysis was performed on the

held-out test set. First, the baseline prediction errors were computed using the original test inputs. The baseline RMSE values were 0.0283 for DoC, 3.287 °C for temperature, and 9.073 for maximum principal stress in native simulation units. Then, each input feature was randomly shuffled across the held-out test cases while all other features were kept unchanged. The model was evaluated again, and the increase in RMSE was recorded. A large increase in RMSE means that the model prediction is sensitive to that parameter. This analysis does not imply direct physical causality, but it identifies which simulation descriptors are most influential for the trained surrogate.

Table 4 lists the most influential input parameters ranked by a composite sensitivity score. The score is the average positive percentage increase in RMSE across DoC, temperature, and stress predictions. The results show that the model is most sensitive to physically meaningful cure-cycle and laminate descriptors. The maximum cure temperature, simulation duration, number of holds, hold temperature, hold time, and cooling rate strongly affect the thermal and cure predictions. In contrast, layup descriptors such as angle spread, laminate symmetry, and ply-angle fractions mainly affect the stress response. This behavior is consistent with the curing process, where thermal-cycle parameters control the cure reaction and temperature history, while stacking-sequence descriptors influence stress development.

Table 4: Permutation sensitivity analysis on the held-out test set. The sensitivity score is the mean positive percentage increase in DoC, temperature, and stress RMSE after shuffling each input feature across held-out test cases.

Rank	Parameter	Score (%)	Main affected output	Interpretation
1	`Tmax_C`	657.5	DoC, Temp., Stress	Maximum cure temperature controls cure rate, thermal peak, and exothermic response.
2	`sim_t_end`	318.2	DoC, Temp.	Total simulated cure duration affects cure completion and thermal history.
3	`angle_std`	309.3	Stress	Layup-angle spread affects the mechanical and stress response.
4	`holds`	256.2	DoC, Temp.	Number of dwell stages affects cure progression and temperature evolution.
5	`hold_temp2_C`	234.1	DoC, Temp., Stress	Second dwell temperature influences cure advancement and thermal response.
6	`symmetric`	229.2	Stress	Laminate symmetry affects stress buildup during curing.
7	`hold_time2_sec`	217.9	DoC, Temp.	Second dwell duration controls time available for cure completion.
8	`cool_rate_C_per_sec`	180.6	DoC, Temp.	Cooling rate affects the late-stage thermal response.
9	`frac_n45`	173.6	Stress	The fraction of -45° plies affects stress prediction.
10	`hold_temp1_C`	154.6	DoC, Temp.	First dwell temperature affects reaction progress during the early cure stage.

4.5 Simulation–Machine Learning Comparison

Figure 7 presents selected best-case and worst-case DoC and temperature predictions for the Manufacturer Recommended and Accelerated Cure Cycle Var1 profiles. Figure 8 shows corresponding predictions for the Fast Heat Fast Cool and Conservative Low-Gradient Cycle Var1 profiles. Each panel compares the neural-network prediction with the FEA simulation result, and these comparisons are used to assess how closely the surrogate reproduces the ACCS-generated response histories within the evaluated simulation dataset.

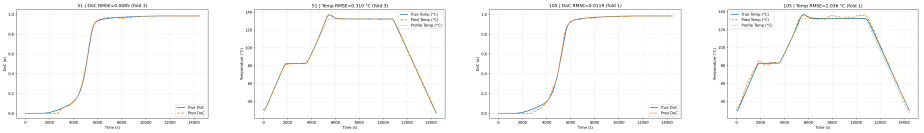

(a) DoC best case (b) Temp best case (c) DoC worst case (d) Temp worst case

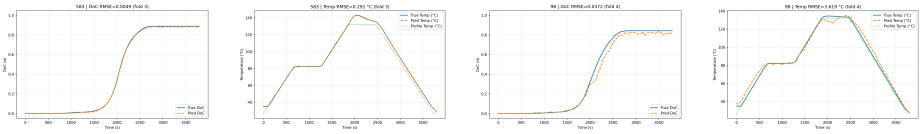

(e) DoC best case (f) Temp best case (g) DoC worst case (h) Temp worst case

Fig. 7: Manufacturer Recommended Cycle Var1 (top row, a–d) and Accelerated Cure Cycle Var1 (bottom row, e–h): (a,e) DoC prediction best case, (b,f) Temperature prediction best case, (c,g) DoC prediction worst case, (d,h) Temperature prediction worst case

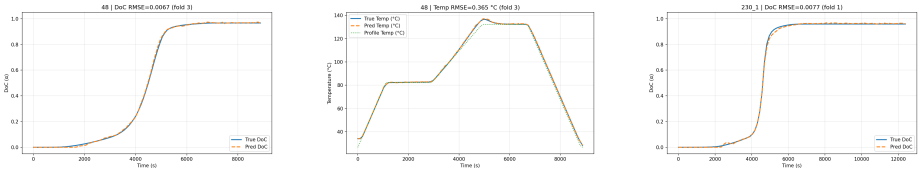

(a) DoC best case (b) Temp best case (c) DoC best case

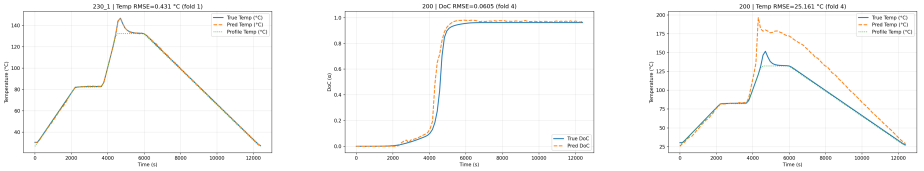

(d) Temp best case (e) DoC worst case (f) Temp worst case

Fig. 8: Fast Heat Fast Cool Cycle (a,b) and Conservative Low-Gradient Cycle Var1 (c–f): (a,c) DoC prediction best case, (b,d) Temperature prediction best case, (e) DoC prediction worst case, (f) Temperature prediction worst case

The selected cases show that the surrogate captures the main trends of the simulated cure response, including the rise in degree of cure, dwell-stage temperature behavior, and cooling-stage evolution. Some larger errors remain in the most difficult cases, especially near rapid thermal transitions and peak-temperature regions. This behavior is consistent with the profile-wise RMSE results in Table 2 and with the cure-specific evaluation metrics described in Section 3.3.

The 96-point representation is used in this work as a fixed-length output grid for learning complete cure histories. The shared warped grid places more sampling resolution in high-slope, high-curvature, and exothermic regions of the

temperature response, while the uniform 96-point grid is used for evaluation. This sampling choice is part of the present preprocessing pipeline and is not presented as a direct performance comparison with the 27-point representation used in previous work. Since the dataset, model inputs, architecture, and evaluation protocol differ from prior studies, the present results are interpreted only within the ACCS-generated dataset used in this paper.

5 Conclusion

The curing of composite structures is governed by strongly coupled thermo-chemical physics, described by the transient heat equation and an autocatalytic cure kinetics model (Eqs. (1)–(4)). Solving this system with FEA provides high-fidelity predictions, but the computational cost becomes prohibitive for large parametric studies or real-time applications. To address this limitation, we developed a data-driven Machine Learning model trained on a comprehensive synthetic dataset generated using ACCS, effectively using the validated FEA framework as a data factory for cure histories.

The combined dataset contains 1016 three-dimensional simulation cases covering a broad process and structural design space. The simulations include variations in heating ramps, dwell temperatures and times, number of holds, laminate thickness, number of plies, stacking sequence symmetry, total heat of reaction, and intentionally incomplete cure conditions. For each case, time-dependent histories of degree of cure, temperature, and maximum principal stress were extracted from the finite element results. These outputs allow the surrogate model to represent the chemical, thermal, and mechanical response of the laminate during curing.

The proposed neural-network surrogate uses a compact 25-dimensional input vector and a shared-task architecture with separate prediction heads for degree of cure, temperature, and maximum principal stress. A shared warped time-grid representation is used to improve the resolution of thermally active regions, including rapid heating and exothermic portions of the cure cycle. The model is evaluated using cross-validation, a held-out test set, profile-wise error summaries, and cure-specific benchmark metrics. These evaluations show that the surrogate can reproduce the main finite element response trends across the explored design space while providing a much faster forward prediction tool.

The results demonstrate that the proposed framework is suitable for rapid comparison of cure-cycle responses and for identifying conditions associated with incomplete cure, large thermal gradients, or elevated stress response. Unlike approaches restricted to a fixed cure cycle or narrow parameter range, the present model generalizes across multiple cure-profile families, hold configurations, laminate thicknesses, and material heat-of-reaction values. Future work will extend the framework to more complex geometries, such as L-shaped or rib-stiffened components, and will investigate how the forward surrogate can be coupled with optimization or uncertainty-quantification methods for physics-informed cure-cycle design.

References

1. Harle, S.M.: Durability and long-term performance of fiber reinforced polymer (FRP) composites: A review. Structures (2024). Elsevier.
2. Kovács, G.: Optimization of a new composite multicellular plate structure in order to reduce weight. Polymers **14**(15), 3121 (2022).
3. Antonucci, V., Giordano, M., Hsiao, K. T., & Advani, S. G.: A methodology to reduce thermal gradients due to the exothermic reactions in composites processing. International Journal of Heat and Mass Transfer **45**(8), 1675–1684 (2002).
4. Muc, A., Romanowicz, P., Chwał, M.: Description of the resin curing process—formulation and optimization. Polymers **11**(1), 127 (2019).
5. Parlevliet, P.P., Bersee, H.E., Beukers, A.: Residual stresses in thermoplastic composites—A study of the literature—Part I: Formation of residual stresses. Composites Part A: Applied Science and Manufacturing **37**(11), 1847–1857 (2006).
6. Advani, S.G., Hsiao, K.-T.: Manufacturing techniques for polymer matrix composites (PMCs). Elsevier (2012).
7. Wu, C., Zhang, R., Zhao, P., Li, L., & Zhang, D.: Curing simulation and data-driven curing curve prediction of thermoset composites. Scientific Reports **14**(1), 31860 (2024).
8. Pantelelis, N., Vrouvakis, T., Spentzas, K.: Cure cycle design for composite materials using computer simulation and optimisation tools. Forschung im Ingenieurwesen **67**(6), 254–262 (2003).
9. Patil, A.S., Moheimani, R., Dalir, H.: Thermomechanical analysis of composite plates curing process using ANSYS composite cure simulation. Thermal Science and Engineering Progress **14**, 100419 (2019).
10. Zhang, Y., An, L., Zhao, C.: Predicting Curing Distortion in Composite Manufacturing—A Fast and Cost-Efficient Numerical Simulation Method. Polymers **16**(24), 3597 (2024).
11. Yang, B., Huang, H., Bi, F., Yin, L., Yang, Q., & Shen, H.: Analysis of cure kinetics of CFRP composites molding process using incremental thermochemical information aggregation networks. Composite Structures **331**, 117904 (2024).
12. Hui, X., Xu, Y., Zhang, W., & Zhang, W.: Cure process evaluation of CFRP composites via neural network: From cure kinetics to thermochemical coupling. Composite Structures **288**, 115341 (2022).
13. Malashin, I. P., Martysyuk, D., Nelyub, V., Borodulin, A., Gantimurov, A., & Tynchenko, V.: A review of physics-informed and data-driven approaches for manufacturing process optimization in polymer matrix composites. Advanced Manufacturing: Polymer & Composites Science **11**(1), 2547335 (2025).
14. Gkertzos, P., Kotzakolios, A., Katsidimas, I., Zaidi, S., Sanchez-Rodriguez, D., Costa, J., & Kostopoulos, V.: Comparative fitting methodology of cure kinetics models based on differential scanning calorimetry. Journal of Composite Materials **59**(1), 75–98 (2025).
15. Cure, T.: Determination of autocatalytic kinetic model parameters describing thermoset cure. J. Appl. Polym. Sci **51**, 761–764 (1994).
16. Kohnke, P.: ANSYS Theory Reference; Ansys Inc: Canonsburg, PA, USA (2020).