Inverse Design of Composite Cure Cycles Using a Machine Learning Surrogate

J.M. Ashfiqur Rahman[1], Arash Ramian[1], Tian Zhao[1], and Rani Elhajjar[1]

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

Abstract. Cure cycle selection for thick composite laminates is difficult because the oven profile controls heat transfer, cure kinetics, exotherm, and stress development during manufacturing. Finite element cure simulation can predict these coupled responses, but direct use in cure cycle search is slow because each new oven profile needs a separate transient analysis. This makes repeated simulation and trial and error design expensive when many cure schedules must be tested. To improve efficiency, this paper uses a feed forward surrogate model inside an inverse search loop to screen promising cure cycles before full validation simulation. The inverse search changes ramp rates, dwell temperatures, dwell times, cooling rate, and total cycle duration. Each candidate profile is passed through the surrogate, and the predicted temperature and degree-of-cure histories are evaluated using cure-completion, thermal-response, cycle-time, feasibility, and training-data-proximity criteria. The objective favors high final degree of cure and shorter cycle time while penalizing non-physical cure behavior, thermal-limit violations, unrealistic temperature shapes, large distance from the reference profile or training-data region, and elevated stress as a secondary diagnostic. This provides a practical way to screen many cure schedules before validation. The method is tested on thermoset laminate cases with different thicknesses, layups, and cure qualities. The inverse search identifies alternative oven profiles that reach the desired cure level and preserve the main thermal and cure trends. The numerical validation results show that the forward cure surrogate can serve as a fast predictor and a design screening tool for inverse cure-cycle optimization of a fixed laminate.

Keywords: Composite curing, cure cycle, inverse design, machine learning surrogate, thermoset composites

1 Introduction

Cure cycle design is an important part of thermoset composite manufacturing. The final part quality depends not only on the material and layup, but also on the heating and cooling history used during cure. This is more important for thicker laminates, especially when the through-thickness thermal resistance is large enough for internal temperature gradients and exothermic heating to develop. In this dataset, the thickest cases corresponded to the 32-ply and 48-ply laminates, which showed stronger thermal-gradient and exotherm effects than the thinner cases. The outer surface follows the oven temperature more closely, while the inside of the laminate can heat and cool at a different rate. This can create thermal gradients, uneven cure, and stress buildup during processing.

Fiber reinforced polymer composites are used in aerospace, automotive, wind energy, marine, and other structural applications because they can give high strength and stiffness with low weight [1, 2]. In thermoset composites, the resin cure process controls the development of the polymer network. During cure, the resin reaction produces heat. This heat can raise the part temperature above the oven temperature, especially in thicker laminates. If the cure cycle is not well selected, the part can have incomplete cure, high thermal gradients, or large residual stress [3–5].

Finite element cure simulation is a useful tool for studying this process. It can predict the temperature field, degree of cure, and stress response during a given cure cycle. It can also include the effect of laminate thickness, layup, heat transfer, and material cure kinetics. For this reason, finite element analysis is widely used for cure simulation and process study [6–10]. However, this approach is slow when many cure cycles must be tested. Each new oven profile usually needs a new transient simulation. This makes trial and error cure cycle design expensive.

Machine learning surrogate models can reduce this cost. A surrogate model can learn from a set of high fidelity simulations and then predict new cases much faster than a full finite element model. This idea has been used in composite processing and cure modeling in different forms, including neural networks, data driven cure models, and physics informed methods [11–13]. The surrogate does not remove the need for physics based simulation. Instead, it gives a fast screening tool. Engineers can use it to test many possible cases before running a smaller number of detailed simulations.

In our earlier work, we developed a finite element generated dataset and a forward neural network surrogate for composite cure prediction [14]. That model predicts three response histories from cure cycle and laminate descriptors: degree of cure, temperature, and maximum principal stress. The input vector contains process features such as ramp rates, dwell temperatures, dwell times, cooling rate, total cycle time, and number of holds. It also contains laminate features such as number of plies, thickness, heat of reaction, symmetry, and layup angle statistics. The forward model was trained using ACCS (Ansys Composite Cure Simulation) simulation data and was evaluated on unseen cases.

The earlier work focused on the forward problem. In that setting, the cure cycle is known, and the model predicts the response. The present paper uses the same forward surrogate for a different task. Here, we study the inverse problem. The goal is to find a new oven profile that gives a desired cure response for a selected laminate. The laminate and material features are kept fixed. The search changes the oven profile parameters, such as ramp rate, dwell temperature, dwell time, cooling rate, and total cycle duration.

This inverse problem is not unique. More than one oven profile can lead to a similar final degree of cure. For this reason, the search should not use final cure alone. It should also check the temperature history, cure evolution, cycle time, and whether the candidate profile satisfies predefined process constraints. In this study, these constraints include final degree of cure, cure monotonicity, peak part temperature, exothermic overshoot above the scheduled oven temperature, temperature-curve smoothness, post-peak reheating, and distance from the surrogate training-data region. Stress is also predicted by the surrogate. In this paper, the predicted maximum principal stress history is used as a secondary diagnostic output rather than as the main success measure because its prediction accuracy is lower than that of temperature and degree of cure. A high predicted stress level may indicate greater potential for cure-induced residual stress and related part-quality concerns. However, stress is not used as a standalone acceptance criterion in this study, and final candidate assessment remains dependent on ACCS validation.

The method used in this paper places the trained forward surrogate inside a candidate search loop. A reference simulation is first selected. New oven profiles are then generated near the reference profile. Each candidate profile is passed through the surrogate. Metrics extracted from the predicted temperature and degree-of-cure histories are evaluated using the scoring function and hard feasibility checks. The candidates are ranked using an objective function that rewards cure completion and reasonable curve shape, while penalizing long cycle time, high temperature, nonphysical cure behavior, and large distance from the training data region.

The aim is not to replace ACCS validation. The aim is to reduce the number of full simulations needed during cure cycle design. The inverse search can test many oven profiles quickly and select a small number of useful candidates. These candidates can then be checked with ACCS. In this way, the final decision remains tied to physics based simulation, but the search process becomes faster.

The main contributions of this paper are as follows. First, the finite element dataset and forward surrogate from the earlier work are included as the base model for inverse design. Second, a bounded local inverse search, referred to here as a trust-region search, is developed to generate candidate cure cycles for a fixed laminate. Third, the candidates are ranked using temperature, degree of cure, cycle time, feasibility, and stress diagnostics. Fourth, selected inverse designed profiles are sent back to ACCS for validation. The results show that the surrogate can be used not only as a fast predictor, but also as a practical screening tool for cure cycle design.

2 Finite Element Cure Model and Dataset

The inverse search in this paper uses the forward surrogate model from our earlier work [14]. That surrogate was trained using finite element cure simulations. For this reason, the finite element model and the dataset are part of the inverse design framework. This section gives the main details of the cure model, simulation setup, dataset, and extracted outputs.

2.1 Cure Kinetics Model

The resin cure state is described using the degree of cure, α. The value of α changes from 0 to 1. A value near 0 means the resin is mostly uncured. A value near 1 means the resin is close to fully cured. The cure reaction also releases heat. This heat can raise the part temperature during the cure cycle.

The cure kinetics were based on differential scanning calorimetry data and an autocatalytic cure kinetics model. The heat flow data can be related to the reaction rate as [15, 16]

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

where H_T is the total heat of reaction and dH/dt is the measured heat flow rate.

The degree of cure can also be written from the heat released during the reaction:

$$\alpha = \frac{1}{H_T}\int_0^t \frac{dH}{dt}\,dt. \tag{2}$$

The cure reaction rate was modeled using an autocatalytic relation:

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

where m and n are cure kinetics exponents. The rate constant $K(T)$ depends on temperature and follows the Arrhenius form:

$$K(T) = A\exp\left(-\frac{E}{RT}\right), \tag{4}$$

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

The material system used in this work is a Toray T700GC/epoxy unidirectional prepreg. The kinetic parameters used in the ACCS simulations are $A = 1.3 \times 10^6$ s^{-1}, $E = 67000$ J/mol, $m = 0.4$, and $n = 1.2$. These parameters define the cure kinetics model used to generate the simulation data. Figure 1(a) shows the agreement between the experimental DSC data and the cure kinetics model at 100°C, 125°C, and 150°C.

2.2 Finite Element Simulation Setup

The simulation data were generated using ANSYS Workbench with ANSYS Composite Cure Simulation (ACCS) and Composite PrePost (ACP). ACCS was used to model the coupled thermal and cure response during the cure cycle [17]. The same simulation workflow was used in the earlier forward surrogate study [14].

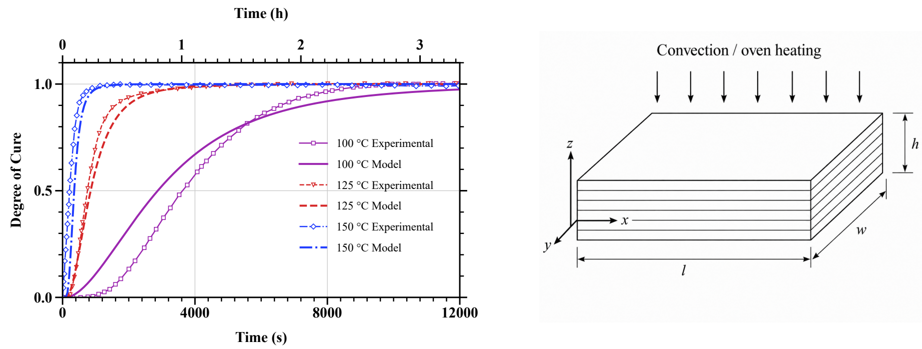

(a)

(b)

Fig. 1. Cure model validation and laminate geometry used for dataset generation. (a) Experimental and model degree of cure at 100°C, 125°C, and 150°C. (b) Flat laminate model used in the ACCS cure simulations.

The model geometry was a flat composite laminate made from the T700GC/epoxy prepreg system. The in-plane dimensions were fixed for all ACCS simulations, with a length of 125 mm and a width of 24 mm. These in-plane dimensions were not treated as design variables. The laminate thickness changed through the number of plies and the single-ply thickness. The laminate was modeled using three-dimensional coupled-field solid elements. The mesh used an element size of about 1.0 mm. The full model had about 7000 elements and 7300 nodes. The initial temperature was set to 20°C. Convection was applied on the free surfaces to model heat exchange with the surrounding environment. Figure 1(b) shows the flat laminate geometry used in the ACCS simulations.

During the simulation, the cure reaction was coupled with the temperature field. When the resin cures, it releases heat. This heat changes the local temperature. The local temperature then changes the cure rate. This coupling is important because thick laminates can develop internal temperature peaks. These peaks may be higher than the oven temperature.

The simulation outputs used in this paper are temperature, degree of cure, and maximum principal stress. Temperature and degree of cure describe the thermo chemical response. Maximum principal stress gives a mechanical measure of stress buildup during the cure process.

2.3 Dataset Generation

The finite element model was used as a data generator. A set of cure simulations was created by changing the laminate, material, and cure cycle parameters. The dataset contains 1016 three dimensional simulation cases. Each case has a different combination of laminate geometry, cure profile, cure quality, and material heat of reaction.

The listed thickness values correspond to the single-ply thickness, not the total laminate thickness. The total laminate thickness is therefore determined by the product of the number of plies and the single-ply thickness. A stacking sequence denotes the ordered set of ply orientations through the laminate thickness. The allowed ply orientations were $-45°$, $0°$, $+45°$, and $90°$; no other ply angles were used. The ply orientations were selected from this discrete set, so no continuous angle increment was used. Symmetric layups were generated by mirroring the stacking sequence about the laminate mid-plane, while asymmetric layups did not satisfy this mirror condition. Representative examples are the symmetric layup $[0/90/{+45}/{-45}]_s$ and the asymmetric layups $[0/90/{+45}/{-45}]$ and $[0/+45/90/-45]$.

The laminate variables include the number of plies, ply thickness, and stacking sequence. Both symmetric and asymmetric layups were included. The cure cycle variables include ramp rate, dwell temperature, dwell time, number of holds, cooling rate, and total cycle time. The dataset also includes both well cured and incomplete cure cases. These incomplete cure cases are useful because the inverse search must be able to recognize profiles that do not reach a good final cure level.

The heat of reaction was also varied to include material variation. The main variables used for dataset generation are summarized in Table 1. Four main cure profile families were used. The manufacturer recommended profile was used as the baseline process. The accelerated profile used faster heating. The conservative low gradient profile used slower heating and cooling to reduce thermal gradients. The step accelerated profile used a multi stage heating path with an intermediate dwell before the final cure temperature. Figure 2 shows these cure profile families.

The simulations were performed using ANSYS/ACCS 2025 on a standard engineering workstation with an Intel Core i7 CPU, 64 GB RAM, and Windows operating system. Each complete case consisted of two sequential analyses: a thermal-chemical cure analysis followed by a mechanical analysis using the thermal and cure results. The average runtime for one complete case was approximately 1 hour, including both analyses. Runtime increased with laminate thickness and number of plies. For the thicker 32-ply and 48-ply cases, the runtime was typically about 2.5 hours or more per case. Overall, the 1016-case dataset required on the order of 1000+ solver hours, depending on laminate thickness and cure-cycle complexity.

2.4 Extracted Simulation Outputs

For each simulation case, time histories were extracted from the ACCS result files. The main outputs used in this work are degree of cure, temperature, and maximum principal stress. These outputs were selected because they describe the chemical, thermal, and mechanical parts of the cure response.

The temperature, degree-of-cure, and maximum-principal-stress histories were extracted using a maximum-history approach. At each saved time step, the output value corresponds to the maximum value over the laminate. Thus, the reported temperature and degree of cure are not spatial averages or fixed center-

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

Category	Variable	Range or values
Geometry	Laminate length	125 mm
Geometry	Laminate width	24 mm
Geometry	Number of plies	4, 8, 16, 32, 48
Geometry	Single-ply thickness	0.1 mm–0.25 mm
Geometry	Ply orientations	$-45°$, $0°$, $+45°$, $90°$
Geometry	Stacking sequence	Symmetric and asymmetric layups
Process	Cure profile strategy	Manufacturer recommended, accelerated, low-gradient, and step-accelerated
Process	Isothermal holds	One-hold and two-hold cycles
Process	Cure completion	Well-cured and incomplete-cure cases
Material	Total heat of reaction	$2.0–5.4 \times 10^5$ J/kg

point values. Similarly, the stress output corresponds to the maximum value of the maximum principal stress over the laminate at each saved time step. The stress values are reported in MPa.

The degree of cure shows how far the resin reaction has progressed and is the main measure used to decide whether the part has reached the desired cure level. The temperature history shows the thermal response of the laminate during heating, dwell, exotherm, and cooling. The maximum principal stress history provides a mechanical indicator of tensile stress buildup during the cure process. Elevated maximum principal stress can indicate increased risk of residual-stress related distortion or damage. In this paper, however, stress is used as a diagnostic output rather than as a standalone acceptance criterion.

3 Preprocessing and Forward Surrogate Model

The inverse search uses the trained forward surrogate as a fast evaluator. Before the surrogate can predict a response, each simulation case must be converted into a fixed input output format. Figure 3 summarizes this full workflow. The upper part shows the preprocessing steps used to build the input features and output targets. The lower part shows the shared task head surrogate model used to predict degree of cure, temperature, and stress histories.

3.1 Input Features

Each simulation case was represented by a 25-feature input vector. The inputs were grouped into process and timing descriptors, cure-cycle structure, laminate/material descriptors, and layup descriptors.

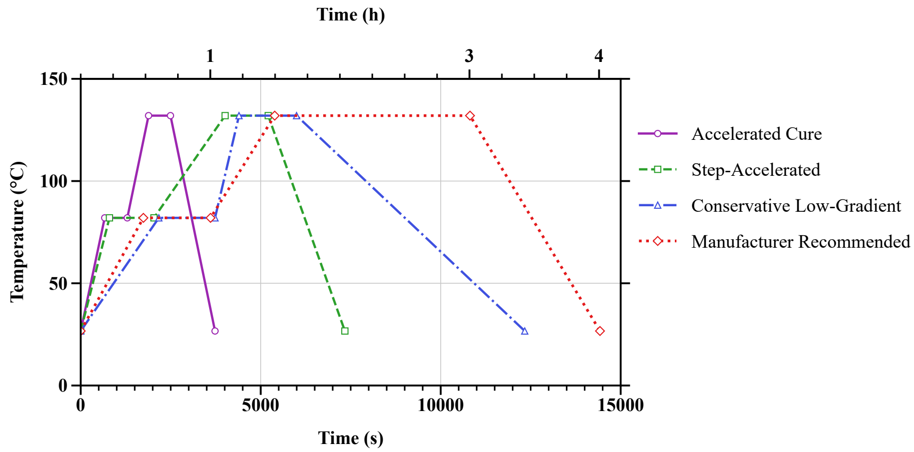

Fig. 2. Representative cure cycle temperature profiles used in the dataset. The profiles include manufacturer recommended, accelerated, conservative low gradient, and step accelerated cure cycles. Adapted from [14].

The 12 process and timing features were the initial temperature, minimum scheduled temperature, maximum scheduled temperature, scheduled cycle duration, first heating rate, first hold temperature, first hold duration, second heating rate, second hold temperature, second hold duration, cooling rate, and simulation-history duration.

The cure-cycle structure feature was the number of temperature holds. The four laminate and material features were the number of plies, total heat of reaction, ply thickness, and laminate symmetry.

The remaining eight layup-descriptor features were the fractions of $-45°$, $0°$, $45°$, $90°$, and other plies; the mean ply angle; the ply-angle standard deviation; and the difference between the fractions of $45°$ and $-45°$ plies.

All input features were min–max scaled using the training split before model training. The same scaling procedure was applied during inverse search.

3.2 Output Curve Representation

The ACCS simulations have different total cycle times. A direct point by point comparison is difficult because one cycle may be shorter than another. To make all cases consistent, time was normalized from 0 to 1. Each output history was represented using 96 points. The three output histories are degree of cure, temperature, and maximum principal stress. This fixed size representation allows the neural network to predict full response curves instead of only final values. A uniform normalized time grid was kept for evaluation and plotting. This makes the predicted and simulated histories easy to compare. For temperature training, a shared warped time grid was also used to give more attention to rapid thermal changes.

3.3 Shared Warped Time Grid

The temperature response does not change at the same rate over the full cure cycle. Some parts are smooth. Other parts change quickly, such as heating ramps, exothermic regions, and peak temperature regions. A uniform grid may use too many points in smooth regions and too few points in fast changing regions.

For this reason, a shared nonuniform time grid was built from the training temperature curves. The grid used temperature slope, curvature, and exotherm information. More points were placed where the temperature response changed more strongly. This helped the model learn the important thermal features without increasing the output size. The warped grid was used for temperature training. The predicted temperature history was then mapped back to the uniform grid for evaluation. Degree of cure and stress were evaluated on the same 96 point uniform grid.

The warped grid is used only for training the temperature-output head. It places more of the 96 temperature points in regions with rapid heating, cooling, curvature, or exothermic behavior, where prediction accuracy is especially important. The uniform normalized grid is retained for reporting, plotting, comparison, and inverse-search scoring because it gives common relative-time locations for temperature, degree of cure, and stress histories. After the temperature head predicts values on the warped grid, the predicted temperature curve is interpolated onto the uniform grid for these shared evaluation tasks.

3.4 Neural Network Architecture

The forward surrogate uses a shared task head neural network. The lower part of Fig. 3 shows the model architecture. The model takes the 25 feature input vector and passes it through a shared trunk. The shared trunk learns information that is useful for all three outputs.

After the shared trunk, the model uses three separate output heads. One head predicts degree of cure. One head predicts temperature. One head predicts maximum principal stress. Each head produces a 96 point history. The degree of cure head uses a sigmoid output so that the predicted cure value stays between 0 and 1. The temperature and stress heads use linear outputs. After prediction, temperature and stress are converted back to physical units using the target scalers.

3.5 Training Targets and Loss Terms

The model was trained to predict degree of cure, temperature, and maximum principal stress at the same time. The main loss included terms for all three outputs. Extra loss terms were also used to improve important parts of the response. For degree of cure, the loss included the full curve, the tail region, the final cure value, and a monotonicity term. The monotonicity term was used because degree of cure should not decrease during curing. A mild cure kinetics regularization term was also included.

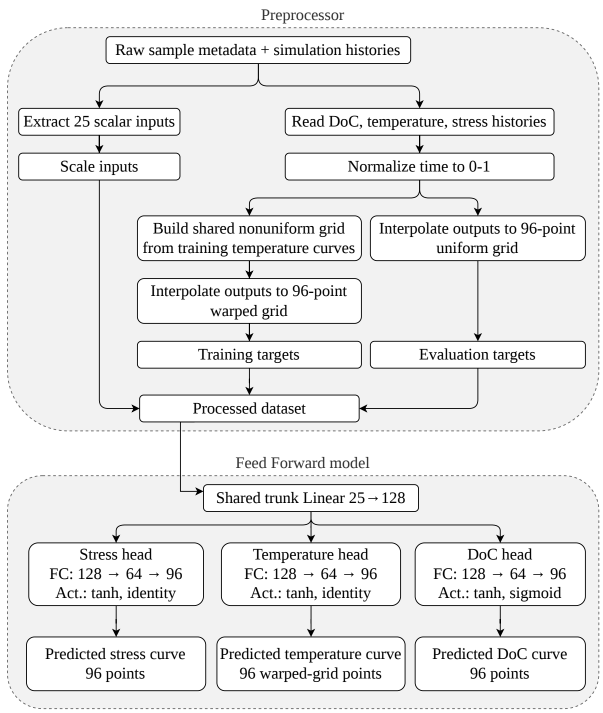

Fig. 3. Combined preprocessing and forward surrogate workflow. The raw simulation metadata and output histories are converted into fixed length input and target arrays. The processed 25-feature input vector is passed through the shared-task-head surrogate to predict degree-of-cure, temperature, and stress histories. The temperature head is trained on the shared warped grid and interpolated to the uniform grid for common evaluation and comparison with the degree-of-cure and stress histories. Adapted from [14]

For temperature, the loss included the main temperature error, a gradient term, a uniform grid term, and an exotherm region term. These terms helped the model capture heating rate, cooling rate, and rapid thermal changes. For stress, the loss included the main stress error, a tail region term, a final value term, and a gradient term. Residual stress was included because it is an important component impacting part quality. However, in the inverse search part of this paper, stress is treated as a diagnostic output. The main validation focuses on temperature and degree of cure.

3.6 Use of the Surrogate in the Inverse Search

For inverse design, the surrogate is used as a fast evaluator. A candidate oven profile is first converted into the same 25 feature input format. The input is

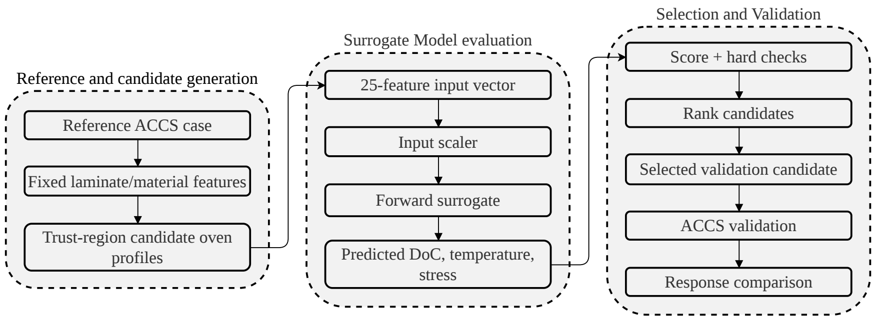

Fig. 4. Workflow of the surrogate assisted inverse cure cycle search. The laminate and material features are kept fixed from the reference case, while trial oven profiles are generated, evaluated by the trained forward surrogate, ranked using score and hard check terms, and selected for ACCS validation

scaled using the same scaler used during training. The surrogate then predicts the response histories. The inverse search mainly uses the predicted temperature and degree of cure curves to rank candidate oven profiles. The predicted stress curve is saved and checked, but it is not used as the main success measure in this paper. This setup allows many cure cycles to be tested quickly. Only the best candidates need to be checked later using ACCS. This is the main reason the forward surrogate is useful for inverse cure cycle design.

4 Inverse Cure Cycle Search

The forward surrogate predicts the cure response for a given oven profile and laminate. In this section, the same surrogate is used inside an inverse search. The goal is to find a new oven profile for a selected laminate case. The laminate and material features are kept fixed. Only the oven profile features are changed. The inverse search is not used as a final approval tool. It is used as a screening tool. It can test many possible oven profiles quickly and select a smaller number of profiles for ACCS validation.

4.1 Search Problem

The search starts from a reference simulation. The reference case has a known laminate, cure cycle, and ACCS response. For this reference case, the laminate and material features are kept unchanged. These include the number of plies, ply thickness, total heat of reaction, symmetry, and layup angle features. Figure 4 summarizes the inverse search workflow used in this study.

The search changes only the cure cycle features. These include ramp rates, dwell temperatures, dwell times, cooling rate, total cycle time, and number of holds. This setup asks a simple process question: for the same laminate, can another oven profile give a useful cure response? In this paper, a useful candidate should reach a high final degree of cure, keep a reasonable temperature response, and stay close to the region learned by the surrogate. When possible, the search also looks for a shorter cycle than the reference case. Stress is predicted and saved, but it is treated as a diagnostic output.

4.2 Candidate Profile Generation

In this study, a trust region means a bounded local search around an existing reference cure cycle. Instead of generating completely unrelated oven schedules, the algorithm changes the reference heating rates, hold temperatures, hold durations, cooling rate, and total cycle duration only within predefined limits. This approach keeps candidate schedules close to processing conditions represented in the training data and reduces the chance of selecting impractical or poorly supported surrogate predictions.

In this work, a candidate means a trial oven temperature profile generated by the inverse search. Each candidate profile is a possible replacement cure cycle for the same laminate. The laminate and material features are not changed. Only the process variables are changed. Each candidate oven profile is written using the same process features used during surrogate training. The search can generate one hold or two hold profiles. The main variables are start temperature, hold temperature, hold time, heating rate, cooling rate, and total cycle time. The candidates are generated near the reference profile. This is done because a fully random global search can create profiles that are inside the numerical feature limits but are not realistic cure cycles. A global search can also move too far away from the training data. In that case, the surrogate prediction may not be reliable. For this reason, the search uses a trust region approach. The reference profile is changed within selected limits. The start temperature is changed only slightly. Hold temperatures can increase or decrease within a limited range. Hold times are sampled around the reference hold times. Heating and cooling rates are also sampled around the reference rates. The total cycle time is checked so that the generated profile remains reasonable. This approach keeps the search local and conservative. It reduces the chance of selecting a nonphysical oven profile.

4.3 Surrogate Prediction of Candidates

After a candidate profile is generated, it is converted into the 25 feature input vector. The process features come from the candidate profile. The laminate and material features come from the reference case. The input vector is then scaled using the same input scaler used during training. The scaled input is passed through the forward surrogate. The surrogate predicts degree of cure, temperature, and maximum principal stress histories. The degree of cure prediction is reported in the range from 0 (no cure) to 1. Temperature and stress are converted back to physical units after prediction. The inverse search mainly uses the predicted temperature and degree of cure curves. These outputs are used to

Table 2. Objective-function weights used in the inverse cure-cycle search.

Penalty term	Purpose	Weight
J_{cyc}	Cycle duration penalty	1.0
J_{α}	Final degree-of-cure deficit	5000
J_{mono}	Degree-of-cure nonmonotonicity	5000
J_{th}	Peak-temperature and schedule-temperature violations	50
J_{exo}	Exotherm above the scheduled oven temperature	50
J_{shape}	Temperature-shape violations, including jaggedness and post-peak reheating	100
J_{dist}	Distance from the nearest training-data sample	5
J_{prof}	Distance from the reference cure profile	0.10
J_{save}	Failure to meet the required cycle-time saving	50
J_{σ}	Maximum-stress diagnostic penalty	1
$J_{\sigma,f}$	Final-stress diagnostic penalty	1

rank the candidate profiles. The predicted stress curve is saved for checking, but stress is not the main validation measure in this paper.

4.4 Scoring Function

Each candidate receives an objective score, where a lower value indicates a more desirable candidate. The score combines cycle duration, cure completion, cure monotonicity, thermal behavior, temperature-shape behavior, proximity to the training-data region, similarity to the reference profile, and stress diagnostics. The scoring function is written as

$$\begin{aligned} J = {} & w_{\mathrm{cyc}} J_{\mathrm{cyc}} + w_{\alpha} J_{\alpha} + w_{\mathrm{mono}} J_{\mathrm{mono}} + w_{\mathrm{th}} J_{\mathrm{th}} \\ & + w_{\mathrm{exo}} J_{\mathrm{exo}} + w_{\mathrm{shape}} J_{\mathrm{shape}} + w_{\mathrm{dist}} J_{\mathrm{dist}} + w_{\mathrm{prof}} J_{\mathrm{prof}} \\ & + w_{\mathrm{save}} J_{\mathrm{save}} + w_{\sigma} J_{\sigma} + w_{\sigma,f} J_{\sigma,f} . \end{aligned} \tag{5}$$

Here, J_{cyc} favors shorter cycle duration, J_{α} penalizes insufficient final degree of cure, and J_{mono} penalizes nonphysical decreases in the predicted degree-of-cure history. J_{th} penalizes excessive peak temperature, scheduled temperature, or predicted peak temperature relative to the reference case. J_{exo} penalizes excessive exothermic overshoot above the scheduled oven temperature. J_{shape} penalizes jagged temperature histories and substantial post-peak reheating. J_{dist} penalizes candidates outside the local training-data region, while J_{prof} penalizes large deviation from the reference cure profile. J_{save} penalizes failure to achieve the required cycle-time saving. The terms J_{σ} and $J_{\sigma,f}$ are associated with maximum stress and final stress, respectively, and are used only as secondary diagnostic penalties.

Final degree of cure and degree-of-cure monotonicity were assigned the largest weights because incomplete cure and nonphysical cure evolution were considered unacceptable. Intermediate weights were assigned to thermal-limit and temperature-shape penalties because they indicate potential processing risk. Lower weights were assigned to stress terms because stress was treated as a secondary diagnostic output in this work.

The weighted objective is used to rank candidates, but it does not establish physical validity by itself. Strict feasibility checks are applied separately for final degree of cure, cure monotonicity, cycle-time reduction, thermal limits, temperature-shape behavior, training-data proximity, and stress response. Therefore, a low objective value alone does not qualify a candidate as a final cure cycle. ACCS validation remains necessary before final acceptance.

4.5 Hard Checks

The inverse search code also records hard checks. These checks are used to mark whether a candidate is strictly feasible. A candidate must pass all of the checks to be called strictly feasible. The main checks are final degree of cure, degree of cure monotonicity, cycle time, peak temperature, exothermic overshoot, temperature smoothness, post peak reheating, distance from the training data, and stress response. These checks are intentionally conservative. They help reject clearly bad candidates.

However, a strict feasibility flag can be too strong. A candidate may fail one hard check but still give a useful temperature and cure response after ACCS validation. This is especially possible when the failed check is related to stress or a conservative temperature shape limit. For this reason, this paper separates strict feasibility from validation usefulness. A strictly feasible candidate passes all automatic checks. A validation candidate is a profile selected by the search for ACCS testing. It may or may not pass every strict check.

4.6 Selection of Validation Candidates

For each reference case, the search saves the ranked candidate list, the selected oven profile, and the predicted response curves. If a strictly feasible candidate is found, the best feasible candidate is selected. If no candidate passes all strict checks, the best ranked rejected candidate is saved as a diagnostic validation candidate. These selected profiles are not treated as final cure cycles. They are sent back to ACCS for validation. The ACCS result is then compared with the surrogate prediction and with the original reference response. This workflow keeps the final decision tied to physics based simulation. The surrogate is used to reduce the search space, not to replace ACCS.

5 Results and Discussion

The inverse search was validated using selected cases from the processed simulation dataset. For each case, the original ACCS simulation was used as the reference case. The laminate and material features were kept fixed. The inverse search then generated a new oven profile for the same laminate.

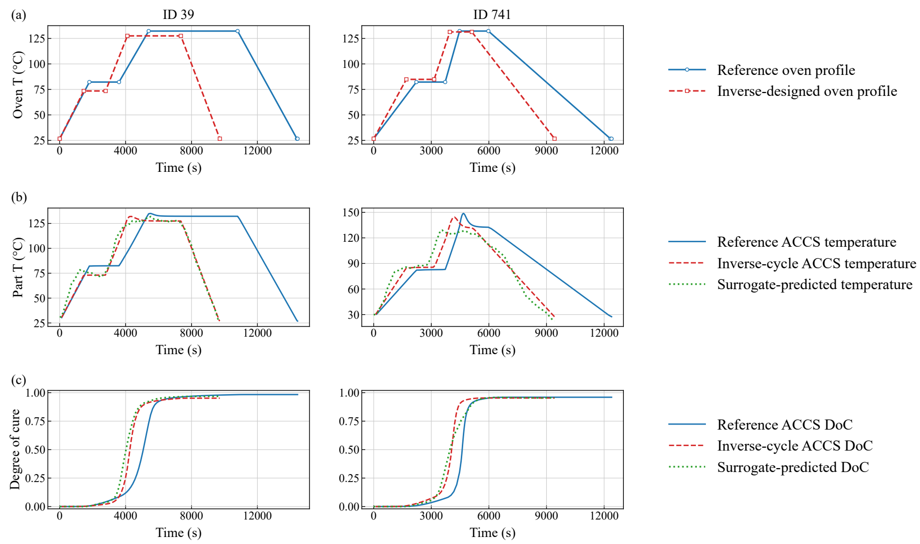

Fig. 5. Representative ACCS validation results for inverse designed cure cycles. Panel (a) compares the reference and inverse designed oven profiles. Panel (b) compares the reference ACCS temperature, inverse cycle ACCS temperature, and surrogate predicted temperature. Panel (c) compares the reference ACCS degree of cure, inverse cycle ACCS degree of cure, and surrogate predicted degree of cure

The selected oven profile was first evaluated by the surrogate. The same profile was then run again in ACCS. This gave two responses for the inverse designed profile: the surrogate prediction and the ACCS validation result. These responses were compared with the original reference response.

The goal of this section is not to show that the surrogate replaces ACCS. The goal is to check whether the inverse search can identify useful oven profiles that are worth testing in ACCS.

5.1 Validation Setup

Each validation case compares the reference cure cycle with the inverse designed cure cycle. The comparison includes the oven temperature profile, the ACCS part temperature response, and the degree of cure response.

Stress was checked separately, but it is not shown in the main validation figure. In this paper, stress is treated as a diagnostic output. The main validation focus is temperature and degree of cure.

Figure 5 shows two representative validation cases. The figure is divided into three rows. Panel (a) compares the reference and inverse designed oven profiles. Panel (b) compares the corresponding part temperature responses from ACCS and the surrogate. Panel (c) compares the degree of cure histories.

5.2 Oven Profile Change

The first row of Fig. 5 shows the oven profiles selected by the inverse search. The inverse designed profiles do not simply copy the reference profiles. The search changes the heating path, dwell time, and cooling duration while keeping the laminate fixed.

In both representative cases, the inverse designed oven profile is shorter than the reference profile. This is important because cycle time is one of the practical drivers for cure cycle design. However, a shorter profile is useful only if it still gives acceptable temperature and cure behavior. For this reason, the oven profile change must be judged together with the ACCS temperature and degree of cure responses.

5.3 Temperature Response

The temperature response is important because cure rate depends strongly on temperature. A useful inverse designed profile should avoid unrealistic temperature behavior and should not create an unsafe thermal response.

As shown in Fig. 5(b), the inverse cycle ACCS temperature follows the main trend predicted by the surrogate. The surrogate captures the main heating stage, the high temperature region, and the cooling stage. This means the surrogate gives useful guidance during the inverse search.

Some differences are still present. These differences are expected because the surrogate uses a compact input vector, while ACCS solves the full transient thermo chemical problem. The largest differences usually occur near rapid thermal changes and near the peak temperature region. These regions are difficult because the exothermic reaction can change the part temperature quickly.

Overall, the temperature results show that the inverse search can find oven profiles that are reasonable enough for ACCS validation. The surrogate is not exact, but it helps reduce the number of full simulations needed during cure cycle screening.

5.4 Degree of Cure Response

The degree of cure is the main measure of cure completion. A useful inverse designed cycle should reach a high final degree of cure and should show a physically reasonable cure history.

As shown in Fig. 5(c), the inverse designed cycles reach a high final degree of cure in the selected validation cases. The surrogate also captures the main rise in degree of cure and gives a final cure level close to the inverse cycle ACCS result.

The timing of cure is also important. In both cases, the inverse designed profile changes when the cure reaction develops compared with the reference case. This is expected because the oven profile is shorter and the heating path is different. The important result is that the final cure level remains high after ACCS validation.

These results support the use of the inverse search as a screening method. The surrogate can identify candidate oven profiles that give useful cure behavior, but the final check still comes from ACCS.

5.5 Cycle Time Effect

The inverse search includes cycle time in the candidate ranking. Shorter cure cycles are useful in manufacturing, but cycle time cannot be minimized alone. A fast cycle is not useful if it gives incomplete cure, excessive temperature, or an unreliable prediction.

The representative cases in Fig. 5(a) show that the inverse designed profiles can reduce the total cycle duration compared with the reference profiles. The ACCS results in Fig. 5(b) and Fig. 5(c) show that these shorter profiles can still give reasonable temperature and cure responses.

This result shows why the inverse search must balance several terms. The search does not select the fastest possible profile. It selects a profile that gives a shorter cycle while still satisfying cure and thermal response requirements.

5.6 Strict Feasibility and Validation Usefulness

The strict feasibility flag was useful for rejecting clearly poor candidates. However, it was also conservative. Some candidate profiles may fail one hard check but still give useful temperature and degree of cure behavior after ACCS validation.

This is important for interpreting the inverse search result. A strict feasibility flag should not be treated as the only decision rule. A candidate that fails one conservative check may still be worth testing in ACCS. This is especially true when the failed check is related to stress or to a conservative temperature shape limit.

For this reason, this paper separates strict feasibility from validation usefulness. A strictly feasible candidate passes all automatic checks. A validation candidate is a profile selected by the search for ACCS testing. It may or may not pass every strict check.

5.7 Stress as a Diagnostic Output

Stress was predicted by the surrogate, but it was not used as the main success measure in this work. The stress response is harder to predict because it depends on temperature gradients, cure shrinkage, layup, evolving resin state, and stiffness development.

For this reason, stress was used only as a diagnostic output. A candidate with high predicted stress can be marked for further review. However, the main validation in this paper is based on temperature and degree of cure.

Future work should improve the stress model before stress is used as a hard acceptance condition. This may require more training data, better stress features, or a separate model focused on stress prediction.

5.8 Overall Discussion

The validation results show that the forward surrogate can support inverse cure cycle screening for a fixed laminate. The method is most reliable when generated profiles stay near the training data region. This is why the trust region search is important. The current method is still a screening method, not a final approval method. The selected candidate should still be checked using ACCS or experiment before being accepted as a final cure cycle.

6 Conclusion

This paper presented an inverse cure cycle search for thermoset composite laminates using a trained forward surrogate model. The surrogate predicts temperature, degree of cure, and maximum principal stress histories from cure cycle and laminate features. In the inverse search, the laminate and material features are kept fixed, while the oven profile is changed.

The method uses the surrogate as a fast evaluator inside a candidate search loop. Each trial oven profile is ranked using degree of cure, temperature response, cycle time, feasibility checks, and stress diagnostics. Selected profiles are then checked again using ACCS validation simulations.

The results show that the inverse search can identify alternative oven profiles that reduce cycle duration while preserving the main temperature and degree of cure trends of the reference case. This shows that a forward cure surrogate can be used not only for fast prediction, but also as a screening tool for cure cycle design.

The method should not be treated as a replacement for finite element validation. The surrogate reduces the search space and helps identify promising candidates, but the final cure cycle should still be checked using ACCS or experiment. Stress was included as a predicted output, but it was used only as a diagnostic response in this work because the stress prediction was less reliable than temperature and degree of cure.

Future work will improve the stress model, test larger validation sets, and extend the method to more complex composite geometries and experimental cure data. These steps are needed before the method can be used for practical cure cycle optimization in real manufacturing cases.

References

1. Harle, S.M.: Durability and long-term performance of fiber reinforced polymer (frp) composites: A review. Structures (2024)
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. Ramian, A., Rahman, J.M.A., Zhao, T., Elhajjar, R.: Modeling dataset development for machine learning prediction of the thermo-chemical curing process in advanced composite structures. In: Proceedings of the 27th International Conference on Engineering Applications of Neural Networks, EANN 2026. Communications in Computer and Information Science, vol. 3024. Springer, Chania, Crete, Greece (2026), proceedings Part I
15. 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)
16. Cure, T.: Determination of autocatalytic kinetic model parameters describing thermoset cure. Journal of Applied Polymer Science 51, 761–764 (1994)
17. Kohnke, P.: ANSYS Theory Reference. Ansys Inc., Canonsburg, PA, USA (2020)