

Autonomous Agentic AI for Adaptive Cure Optimization and Defect Prevention in Thermoset Polymer Composite Manufacturing

Manish Kumar Jha*

Abstract

Thermoset polymer composites occupy a central position in modern structural manufacturing, from aircraft fuselages to wind-turbine blades. Despite progress in resin chemistry and fiber architecture, the “cure process” that transforms compliant preforms into load-bearing structures remains difficult to manage. Manufacturers encounter ‘voids’, “interlaminar delaminations”, and “spring-back distortion” when curing complex or thick-section parts. The cause is not ignorance of the relevant physics, but rather that ‘temperature’, ‘chemistry’, ‘rheology’, and ‘mechanics’ couple in ways that no fixed cure schedule can anticipate across the full range of production variability. This paper presents the “Autonomous Agentic Artificial Intelligence System for Adaptive Cure Optimization (AAAI-ACO)”, a framework designed to close this gap. Rather than prescribing a cure cycle in advance, the system watches the evolving process state through embedded sensors - “fiber Bragg gratings”, “multi-frequency dielectric probes”, ‘thermocouples’, and “ultrasonic transducers” - fuses those measurements in real time using an Extended Kalman Filter, and feeds the result into a continuously running digital twin built around a physics-informed neural network (PINN). The PINN re-produces the Kamal-Sourour cure kinetics and the coupled energy equation at roughly 200 times the speed of a full finite-element model, making it fast to project process trajectories forward by 30 minutes and evaluate candidate control actions before committing to them. Decisions are taken by a hierarchy of four software agents: spatially distributed sensing agents for local data quality and anomaly detection; a defect prediction agent running gradient-boosting and LSTM models for probabilistic void, delamination, and residual stress risk estimation; a coordination agent that queries a Proximal Policy Optimization (PPO) deep reinforcement learning policy; and a supervisor agent that enforces hard safety limits throughout. The PPO agent was pre-trained across an initial 50,000 simulated cure cycles before full training extended to 250,000 episodes with varied resin batch properties and geometric conditions, building a policy that generalizes rather than memorizes. Conceptual validation against Hexcel 8552/AS4 laminate cure data from the published literature projects reductions of approximately 47% in void content, 37% in spring-back distortion, and 20% in energy consumption relative to the manufacturer’s standard cure cycle, alongside an 18.6% reduction in cycle time. These are simulation-derived projections, and the author acknowledges that hardware-in-the-loop testing will be required to confirm framework performance in a real autoclave. The results are internally consistent, benchmarked against comparable studies, and grounded in established control theory and materials science. This work provides a coherent conceptual foundation for researchers working toward autonomous cure management in safety-critical composite production.

*Author for Correspondence

Manish Kumar Jha

PhD, Department of Computer Science, Patliputra University,
Patna, Bihar, India

Received Date: March 20, 2026

Accepted Date: March 27, 2026

Published Date: April 10, 2026

Citation: Manish Kumar Jha. Autonomous Agentic AI for Adaptive Cure Optimization and Defect Prevention in Thermoset Polymer Composite Manufacturing. Journal of Polymer & Composites. 2026; 14(Special Issue 2): S301–S320p.

Keywords: Thermoset composite curing, autonomous agentic AI, reinforcement learning, digital twin, defect prevention, adaptive process control, multi-agent systems, industry 4.0

INTRODUCTION

Background and Industrial Context

Walk through any large aerospace fabrication facility and one will certainly encounter an autoclave - usually more than one. These

pressurised ovens are where carbon-fiber-reinforced epoxy laminates complete their chemical transformation, cross-linking under controlled heat and pressure into the stiff, lightweight panels and spars that keep aircraft airborne. The appeal of thermoset polymer matrix composites (PMCs) is not difficult to understand: their strength and stiffness are difficult to match with metals, their fatigue behaviour is generally favourable, and their anisotropy can be exploited by designers in ways that isotropic alloys cannot accommodate [1, 2].

The global composites market was valued at approximately USD 96 billion (approximately ₹7.97 lakh crore, at the 2023 average exchange rate of 83 INR per USD) in 2023 and is expected to surpass USD 157 billion (approximately ₹13.03 lakh crore) by 2030, with thermoset-based systems - predominantly carbon fiber-reinforced epoxy, bismaleimide, and cyanate ester composites - constituting more than 65% of structural applications [3].

Yet the ‘cure’ process itself is a persistent source of manufacturing anxiety. The thermal history applied to a preform - its “ramp rates”, “dwell temperatures”, “dwell durations”, “applied pressure profile”, and “cool-down trajectory” - determines nearly everything about the finished part: its “degree of cure”, its “residual stress state”, its “dimensional accuracy”, and its “mechanical performance” [6]. The problem is that those parameters interact with part geometry, tooling thermal mass, fiber volume fraction, and resin batch-to-batch variability in ways that are genuinely non-linear and hard to anticipate in a fixed schedule [4, 7, 8]. Resin that behaves one way at laboratory scale may gel some-what earlier or later in production and an autoclave that heats a flat coupon uniformly may impose 15°C or more of spatial temperature variation across a frame.

The failure modes arising from imperfect cure management are well catalogued. Voids form when en-trapped ‘moisture’ or ‘volatiles’ cannot escape before gelation locks them in place, because pressure was applied too ‘early’ or too ‘late’ relative to the resin viscosity window [9]. Interlaminar delaminations can develop when through-thickness thermal gradients impose stresses that exceed the matrix fracture toughness before sufficient cross-linking has occurred [10]. Residual stresses accumulate from differential chemical shrinkage and thermal mis-match between ‘fiber’ and ‘matrix’, producing the spring-back that causes assembled structures to require shimming and rework [5, 11]. Incomplete cure leaves a depressed glass transition temperature and a material prone to creep under service loads [12]. Standard cure cycles - the Hexcel 8552/AS4 two-stage cycle at 180°C and 586 kPa is a representative example - are deliberately conservative precisely for these reasons, but conservatism comes at a cost: “longer cycle times”, “higher energy consumption”, and “no guarantee of defect-free parts when process perturbations occur” [13].

Why Static Control Is Not Sufficient

The composites manufacturing community recognised some time ago that improving ‘cure’ quality requires better process control, not simply ‘longer’ or ‘hotter’ cycles. Model predictive control (MPC) approaches have been explored for cure scheduling since at least the early 1990s and have shown genuine promise under planned laboratory conditions [16]. The fundamental limitation of conventional MPC, however, is structural rather than incidental: it relies on a “kinetics model” whose parameters must be characterised with considerable precision, and it cannot recover gracefully when the actual process diverges from what that model predicts [16]. Resin batch variability alone can shift gelation time by 20 to 30 minutes in a nominally identical material [7, 8], and a fixed-model controller has no mechanism to accommodate that degree of variation.

The Role of Industry 4.0 Technologies

The convergence of embedded sensing infrastructure, edge computing capability, and machine learning algorithms has opened a more promising path [14, 15, 18]. The goal is not simply to replace one fixed recipe with another fixed model, but to build a system that learns the state of the process it is controlling and adjusts its strategy accordingly. Autonomous agentic AI - characterised by perception-action loops, the ability to plan over uncertain futures, and the capacity to improve through

operational experience - embodies exactly this kind of adaptability [17]. Re-inforcement learning is well-suited to this context because it does not require an explicit system model; the agent discovers effective control strategies through simulated experience, observing the consequences of its actions and refining its policy accordingly [19]. That property makes RL-based controllers genuinely robust to the parametric uncertainties that defeat conventional MPC [30, 31].

Research Gap and Contributions

A reading of the existing literature reveals that prior work has addressed individual components of the cure management problem, but no study has assembled them into a single coherent architecture. Isolated RL agents have been tested in simplified cure simulations [20, 35]. Digital twins have been built for process monitoring without autonomous decision-making capability [21, 57]. Machine learning has been applied to defect prediction without feeding those predictions back into a closed control loop [22]. To the best of the author's knowledge, no published framework has brought all these components together into a unified architecture and described, in operationally sufficient detail, how they interact during an actual cure cycle.

This paper addresses that gap through five specific contributions. First, a complete hierarchical architecture for autonomous cure management is described at a level of detail sufficient for practical implementation. Second, the Kamal-Sourour cure kinetics [41] are embedded within a physics-informed digital twin that runs fast enough for real-time state estimation and predictive simulation. Third, a PPO-based reinforcement learning formulation is developed with a multi-objective reward function covering all four principal defect modes [68]. Fourth, a multi-agent coordination protocol is proposed for the spatially heterogeneous conditions characteristic of industrial autoclave environments [36, 40]. Fifth, a benchmarked performance analysis demonstrates the projected defect reduction and efficiency improvements achievable relative to standard cure cycles [13, 29]. The framework is conceptual at this stage rather than experimentally validated, and the paper is explicit about that limitation throughout. The intention is to establish a foundation on which future experimental work can build.

LITERATURE REVIEW

Artificial Intelligence in Manufacturing: A Selective Overview

The application of machine learning to manufacturing quality control has expanded substantially over the past decade, driven partly by the availability of low-cost sensing hardware and partly by significant advances in algorithm capability [23, 24]. Convolutional neural networks trained on ultrasonic C-scan images have achieved defect detection accuracy comparable to experienced human inspectors in non-destructive evaluation of composite structures [25]. Long short-term memory (LSTM) architectures have been applied to time-series process data for anomaly detection in injection moulding and extrusion, flagging developing faults well before they propagate into finished-part defects [26]. Bayesian optimisation has found a productive niche in manufacturing parameter space exploration, where its sample efficiency is particularly valuable given the cost of physical experiments [27].

Within the composites domain, the work of Fernandez et al. [28] demonstrated that neural networks trained on dielectric spectroscopy data could predict degree of cure with $R^2 > 0.95$ - a result useful for monitoring but not extended to closed-loop control. Struzziero et al. [29] employed surrogate model-assisted optimisation for thick-section aerospace laminates and reported cycle time reductions of approximately 23% with maintained mechanical property targets, demonstrating that even offline optimisation methods can yield significant practical improvements.

Reinforcement Learning for Process Control

Reinforcement learning has attracted growing interest from the process engineering community [30, 31]. Its core appeal is that an RL agent can discover effective control strategies through simulated experience without requiring an explicit mathematical plant model - a significant advantage in

manufacturing contexts where building accurate process models is expensive and time-consuming [19]. Algorithms such as Proximal Policy Optimisation (PPO) and Soft Actor-Critic (SAC) have emerged as reliable workhorses for continuous-action manufacturing control problems, offering stable training and competitive sample efficiency [32].

Several applications are relevant here. Kim et al. [33] applied PPO to injection moulding thermal profile control and reduced warpage defects by 31% relative to standard cycles. Bortot et al. [34] combined Gaussian process dynamics models with RL in a model-based framework for extrusion screw speed optimisation, improving melt homogeneity scores by 18%. The most directly pertinent prior study is that of Gonzalez et al. [35], who demonstrated Q-learning-based dwell temperature adaptation in a one-dimensional cure simulation, reducing peak exotherm by 12°C. That result is encouraging, but the framework was limited to a single controllable parameter, a highly simplified one-dimensional kinetics model, and had no connection to real-time sensor data or multi-physics defect prediction - a gap that the present work directly addresses.

Multi-Agent Systems in Industrial Settings

Industrial autoclaves used for composite curing present a spatially distributed control problem that single-sensor, single-controller architectures cannot adequately manage [36, 37]. Chambers may be 6 to 8 metres in diameter and up to 20 metres long, and temperature, pressure, and resin state vary across that volume in ways that require distributed observation and localised response. Multi-agent systems (MAS) provide a natural computational structure for this situation, with each agent responsible for a spatial zone and coordinating with neighbours rather than deferring to a central controller for every decision [36].

Foundational work by Leitao and Restivo [38] demonstrated that holonic agent-based manufacturing systems outperform centralised schedulers in dynamic, multi-machine environments. Trentesaux et al. [39] formalised the PROSA reference architecture for hierarchical agent-based manufacturing control, a conceptual lineage that directly influenced the agent hierarchy proposed in this paper. More recently, Park et al. [40] deployed a multi-agent reinforcement learning system in an aerospace autoclave context, where spatially distributed temperature agents coordinated to reduce inter-zone thermal gradients by 27% - a result that specifically validates the multi-agent approach for autoclave cure management.

Curing Process Modelling: Kinetics and Consolidation

Cure process control requires a mathematical representation of the curing reaction. The autocatalytic phenomenological model of Kamal and Sourour [41] has remained the dominant framework for decades, owing to its physical interpretability, compact parameterisation, and demonstrated validity across epoxy, bismaleimide, cyanate ester, and polyimide resin systems. Johnston et al. [42] extended this model to incorporate coupled heat transfer and degree-of-cure field equations, providing the basis for three-dimensional process simulations now implemented in tools such as the RAVEN software platform [43].

Parallel developments in consolidation modelling are equally important for defect prediction. Loos and Springer [44] established the resin flow and pressure framework linking applied autoclave pressure to void dynamics, while Gutowski et al. [45] contributed permeability-based fibre bed compaction models needed to predict void evolution under consolidation pressure. Mesogitis et al. [46] demonstrated through stochastic cure simulation that uncertainty in resin viscosity characterisation is the dominant source of scatter in predicted residual stress fields - a finding that underlines the importance of tight material characterisation and motivates the real-time state estimation approach adopted in the present framework.

Defect Formation: What the Literature Shows

‘Void’ formation is fundamentally a pressure balance problem: nucleation occurs when the partial pressure of dissolved volatiles exceeds local resin pressure during the low-viscosity cure window [9,

48]. Centea and Hubert [49] quantified this experimentally for out-of-autoclave prepreps, finding that 'void' content could range from below 0.5% to above 3% depending on cure temperature ramp rate and vacuum level - a variation that translated directly into 15-25% differences in interlaminar shear strength. For practitioners, this result establishes pressure application timing relative to resin viscosity evolution as probably the most consequential controllable variable for void management.

Delamination during cure is driven by a different mechanism: through-thickness thermal gradients that generate transient interlaminar stresses in the still-curing laminate [10, 50]. In thick sections (above approximately 25 mm), internal exotherms from the cure reaction can exceed 20°C above the tool surface temperature [50], imposing stresses that can reach the matrix fracture toughness before the material has developed sufficient toughness to resist them. Residual stresses and spring-back distortion arise from differential chemical shrinkage and thermal mismatch between fibre and matrix [51], with chemical shrinkage alone contributing between 1.5 and 5% volumetric strain depending on the resin system [52] - a wide enough range that batch-to-batch variability has a measurable effect on dimensional accuracy [47].

Digital Twins in Composites Manufacturing

Digital twin technology - a virtual model of a physical system updated continuously from sensor data and capable of forward prediction - has become a key enabler for autonomous manufacturing management [53, 54]. Grieves and Vickers [55] formalised the conceptual framework, and Tao et al. [56] proposed the five-dimensional model encompassing physical entity, virtual entity, services, data, and connections that has been widely adopted as a reference architecture.

The practical challenge in composites manufacturing is computational speed: full finite-element cure simulations are typically orders of magnitude slower than real time, rendering them unsuitable for direct use in a closed-loop control system [43]. Recent developments address this constraint. Schapira et al. [57] demonstrated a digital twin combining FBG strain measurements with FE cure simulation for degree-of-cure field estimation, achieving 4.2% mean absolute error. Raju et al. [58] showed that physics-informed neural networks (PINNs) can emulate full FEM cure simulation results three orders of magnitude faster while maintaining temperature prediction accuracy within 5°C of embedded thermocouple measurements - a result that makes the PINN-based digital twin technically credible for real-time use, though accuracy must be validated carefully for each new application.

Summary of Research Gaps

The literature reviewed above supports three conclusions that directly motivate the present work. First, AI methods have been applied to composite cure problems almost exclusively in isolation: reinforcement learning for optimisation [20, 35], machine learning for monitoring [22, 28], digital twins for state estimation [21, 57] - rarely in combination and never within a hierarchical autonomous architecture addressing all principal defect modes simultaneously. Second, the multi-agent approach best suited to spatially distributed autoclave environments [36, 40] has received only limited attention in the composites manufacturing context specifically. Third, scalability from simplified laboratory-scale studies to production-scale curing of complex structural components has not been seriously addressed by any existing AI cure management proposal [29, 35]. The AAAI-ACO framework proposed in subsequent sections addresses each of these gaps directly.

THEORETICAL FOUNDATIONS

Cure Kinetics: The Kamal-Sourour Model and Its Extensions

The reaction rate of a curing thermoset resin is described using the autocatalytic model proposed by Kamal and Sourour in 1973 [41]. Despite its age, this model remains the workhorse of industrial cure process simulation because it captures the essential feature of autocatalytic crosslinking reactions - an initial acceleration followed by diffusion-limited deceleration - with only a handful of parameters that can be determined from differential scanning calorimetry (DSC) experiments. The reaction rate is written:

$$d\alpha/dt = (k_1 + k_2 \cdot \alpha^m) \cdot (1 - \alpha)^n \quad (1)$$

where α is the degree of cure (ranging from 0, unreacted, to 1, fully cross-linked), k_1 and k_2 are temperature-dependent rate constants, and m and n are empirical reaction order parameters fitted to DSC data. The rate constants obey Arrhenius kinetics:

$$k_i = A_i \cdot \exp(-E_{ai}/RT), \quad i = 1, 2 \quad (2)$$

with A_i the pre-exponential frequency factor, E_{ai} the activation energy for the i -th reaction pathway, R the universal gas constant (8.314 J/mol·K), and T the absolute temperature in Kelvin. In practice, real cure processes deviate from Equation (1) at high conversion because molecular mobility decreases as the polymer network forms. A diffusion factor $D(\alpha, T)$ is incorporated to capture this:

$$d\alpha/dt = D(\alpha, T) \cdot (k_1 + k_2 \cdot \alpha^m) \cdot (1 - \alpha)^n \quad (3)$$

where $D(\alpha, T) = [1 + \exp(C(\alpha - \alpha_c))]^{-1}$, with C being a diffusion constant and α_c the critical degree of cure at the onset of diffusion control [59]. The governing equation for the temperature field within the curing laminate couples this reaction rate to heat conduction:

$$\rho \cdot C_p \cdot \partial T / \partial t = \nabla \cdot (k_{th} \nabla T) + \rho_r \cdot H_r \cdot d\alpha/dt \quad \dots (4)$$

where ρ is the composite density, C_p the specific heat, k_{th} the thermal conductivity tensor, ρ_r the resin density, and H_r the total heat of reaction. This system of coupled partial differential equations, when solved over the part volume with boundary conditions reflecting autoclave air temperature and tooling thermal contact, provides the thermo-chemical state evolution that the digital twin must reproduce in real time.

Controllable Process Parameters

From a control engineering standpoint, the cure cycle defines the input trajectory applied to the autoclave system. For a standard two-stage epoxy cure, the relevant parameters are the first-stage ramp rate $(dT/dt)_1$, the first-stage dwell temperature T_{d1} and duration τ_{d1} , the second-stage ramp rate $(dT/dt)_2$, the second-stage dwell temperature T_{d2} and duration τ_{d2} , the pressure profile $P(t)$, and the cool-down rate $(dT/dt)_{cool}$. Collecting these into a vector:

$$\theta = \{(dT/dt)_1, T_{d1}, \tau_{d1}, (dT/dt)_2, T_{d2}, \tau_{d2}, P(t), (dT/dt)_{cool}\} \quad (5)$$

The objective is to find the trajectory $\theta^*(t)$ that minimises a composite defect risk function while maintaining the final degree of cure above an acceptable threshold (typically $\alpha_{final} \geq 0.95$) and respecting actuator limits and safety boundaries. This is the optimisation problem that the reinforcement learning agent is trained to solve.

Principal Defect Mechanisms

Void Formation

Voids nucleate when the partial pressure of dissolved volatiles inside the resin exceeds the pressure resisting bubble growth. The condition for void growth is:

$$P_{void} = P_{atm} + (2\gamma/r) > P_{applied} + P_{resin} \quad \dots (6)$$

where P_{void} is the internal void pressure, γ the resin surface tension, r the bubble radius, $P_{applied}$ the autoclave consolidation pressure, and P_{resin} the resin hydrostatic pressure accounting for flow resistance in the fiber bed [9, 49]. Void suppression requires the applied pressure to exceed the volatile partial pressure throughout the low-viscosity window (the period when $\eta < \text{approximately } 10 \text{ Pa.s}$). Timing matters enormously: pressure applied after gelation accomplishes nothing, while pressure applied too early can displace resin before the fiber bed is properly consolidated.

Delamination

Through-thickness temperature gradients generate transient interlaminar stresses in anisotropic laminates. A useful first approximation for the interlaminar shear stress at a critical interface is:

$$\tau_{\{xz\}} \approx Q_{\{55\}} \cdot \gamma_{\{xz\}} + (E_{\{11\}} \cdot \alpha_{\{11\}} - E_{\{22\}} \cdot \alpha_{\{22\}}) \cdot \Delta T/h \quad (7)$$

where $Q_{\{55\}}$ is the laminate shear stiffness, $\gamma_{\{xz\}}$ the interlaminar shear strain, $\alpha_{\{11\}}$ and $\alpha_{\{22\}}$ the in-plane and transverse thermal expansion coefficients, ΔT the through-thickness temperature gradient, and h the interlaminar spacing [10, 50]. Delamination occurs when $\tau_{\{xz\}}$ exceeds the mode II fracture toughness threshold. The implication for cure control is that heating ramp rates must be moderated in thick sections to keep thermal gradients below the critical level, even when a faster ramp would otherwise be preferred for productivity.

Residual Stresses and Cure Completion

Cure-induced residual stresses accumulate incrementally as the resin transitions from a viscous liquid to a glassy solid. The incremental residual stress model is:

$$\Delta \sigma_r = C(T, \alpha) \cdot [\Delta \epsilon_{\{mech\}} - \Delta \epsilon_{\{th\}} - \Delta \epsilon_{\{ch\}}] \quad (8)$$

where $C(T, \alpha)$ is the evolving stiffness tensor and the three strain increments represent mechanical, thermal, and chemical contributions respectively [51, 52]. The key control lever for residual stress management is the cool-down profile: a slower cool-down allows more stress relaxation before the resin vitrifies fully, reducing the locked-in stress state and consequent spring-back. Incomplete cure ($\alpha_{\{final\}}$ below the target threshold) compounds this problem by leaving a material with reduced crosslink density and a glass transition temperature too close to service conditions.

The Autonomous Agent as a Markov Decision Process

An autonomous agent is, in formal terms, a computational entity that observes a state s , selects an action a according to some policy $\pi(a|s)$, and receives a scalar reward r evaluating how well the action served the agent's objectives [60]. The cure management problem maps naturally onto a Markov Decision Process:

$$MDP = (S, A, P, R, \gamma) \quad (9)$$

where S is the process state space (temperature fields, degree-of-cure distribution, resin viscosity, defect risk estimates), A is the action space (adjustments to cure cycle parameters), $P(s'|s, a)$ is the state transition probability modelled by the digital twin, $R(s, a)$ is the multi-objective reward function, and γ is the discount factor balancing immediate against long-term rewards [61]. The optimal policy $\pi^*(a|s)$ - the one that maximises expected cumulative discounted reward - constitutes the optimal cure management strategy. Finding this policy is the learning problem that the PPO algorithm solves through experience in the simulated cure environment.

PROPOSED FRAMEWORK: THE AAAI-ACO ARCHITECTURE

Overall Structure

The AAAI-ACO system is organised into five layers that together implement a continuous perception-reasoning-action loop operating at multiple timescales. Raw sensor data arrives at up to 100 Hz. The digital twin integrates those measurements and updates the process state estimate at 1 Hz. The reinforcement learning agent evaluates the current state and selects adjustments to cure cycle parameters every 2 minutes. The supervisor agent monitors all activity in real time and overrides any action that would violate a safety constraint. Figure 1 presents a schematic overview of the AAAI-ACO layered architecture and the principal data flow pathways between system components. Table 1 summarises the five layers, their principal components, and the data they exchange.


Physical Sensing Layer

Instrumentation choices matter considerably in any real-time control system, and composite cure is no exception. The sensing layer adopted here combines four measurement modalities that are complementary in what they observe and at what spatial resolution. Embedded Type-K thermocouples at the tool-part interface and mid-plane locations provide temperature-time histories with $\pm 0.5^\circ\text{C}$ accuracy - relatively coarse in spatial resolution but reliable and cost-effective. Fiber Bragg Grating

Table 1. AAAI-ACO system layer summary.

System Layer	Key components	Primary function	Data interface
Physical Sensing	Thermocouples (Type-K), FBG sensors, dielectric probes, ultrasonic transducers, IR cameras	Multi-modal real-time measurement of process state	Raw sensor streams at up to 100 Hz
Digital Twin	PINN surrogate, Extended Kalman Filter, FEM reference model, state estimator	Real-time state reconstruction; 30-min predictive simulation horizon	Fused state estimates at 1 Hz
Autonomous Agents	Sensing agents (per zone), Defect Prediction Agent, Coordination Agent, Supervisor Agent	Distributed perception, defect risk quantification, decision coordination, safety enforcement	Shared blackboard; agent messages
Optimisation Engine	PPO policy network, multi-objective reward calculator, replay buffer	Cure policy learning and action generation	Action vectors; reward signals
Control Interface	Zone PID controllers, MPC layer, actuator drivers, safety interlocks	Setpoint execution and actuator management	Setpoints dispatched to autoclave hardware

Achieving optimal cure cycle control

**Figure 1.** A schematic overview of the AAAI-ACO layered architecture.

(FBG) sensors embedded within the laminate track strain evolution throughout cure, capturing the onset of resin shrinkage and the development of residual stress in real time at discrete laminate locations [62]. Multi-frequency dielectric sensors measure the ionic viscosity of the resin through the Debye model, providing the closest available real-time measurement of flow state and degree of cure [63]. High-frequency ultrasonic transducers mounted on the tooling surface monitor acoustic impedance changes associated with consolidation and void content evolution [64].

None of these sensor types alone provides a complete picture. Thermocouples see temperature but not cure state. FBGs see strain but not void content. Dielectric sensors see bulk resin properties but are sensitive to sensor placement and contact quality. The fusion step - implemented via an Extended Kalman Filter (EKF) - combines all four modalities with the physics-based prediction from the digital twin to produce a statistically optimal estimate of the full process state vector [65]. This is where the heterogeneity of the sensing system becomes a strength rather than a complication.

Digital Twin Layer

The digital twin serves two distinct functions: “state estimation” (what is happening in the part right now?) and ‘prediction’ (what will happen if the agent takes a particular action?). Both require a model that is both accurate and fast. Full three-dimensional finite-element cure simulations satisfy the first requirement but emphatically not the second - they typically run 50 to 500 times slower than real time for complex part geometries. Physics-informed neural networks offer a way out of this bind.

The PINN adopted here is trained to emulate the thermo-chemical system defined by Equations (1) through (4), with the cure kinetics physics embedded as soft constraints in the loss function:

$$L = L_{\text{data}} + \lambda_{\text{phys}} \cdot L_{\text{physics}} + \lambda_{\text{BC}} \cdot L_{\text{BC}} \quad (10)$$

where L_{data} is the mean squared error between PINN predictions and sensor measurements, L_{physics} is the PDE residual evaluated at collocation points inside the part volume, and L_{BC} is the boundary condition residual [66]. The physics terms prevent the network from fitting sensor measurements in physically implausible ways - a failure mode that purely data-driven models are prone to when measurements are sparse, as they inevitably are in a real composite laminate. Training the PINN for a specific resin system and part family takes 24 to 72 hours on GPU hardware, after which inference runs in well under a second - fast enough to project 30 minutes of cure process dynamics in real time.

The digital twin operates in two alternating modes: an assimilation mode in which EKF-fused sensor measurements update the state fields at 1 Hz, and a predictive mode in which candidate control actions proposed by the RL policy are evaluated by rolling the PINN forward in time. This look-ahead capability transforms the agent from a reactive controller into a genuinely anticipatory one.

Autonomous Agent Layer

Four agent types handle different aspects of the monitoring and decision problem, coordinated through shared blackboard architecture. This structure was chosen over a flat peer-to-peer scheme because the tasks naturally fall into a hierarchy: “local sensing”, “zone-level defect risk estimation”, “global coordination”, and “safety enforcement” are genuinely different problems operating at different spatial and temporal scales.

Sensing Agents (SA)

One sensing agent is instantiated per sensor cluster zone in the autoclave. Its responsibilities are “local data acquisition”, “outlier detection using a pre-trained SVM-based anomaly model”, and “computation of a zone-level degree-of-cure estimate using a Gaussian Process model conditioned on local thermocouple and dielectric measurements”. Zone-level state summaries are reported to the Coordination Agent every second.

Defect Prediction Agent (DPA)

The DPA is arguably the most novel element of the architecture. It maintains three parallel prediction models: a gradient boosting classifier trained on over 15,000 simulated cure scenarios to estimate the probability of void content exceeding acceptance thresholds given the current resin viscosity, temperature gradient, and applied pressure; a physics-guided LSTM network that evaluates evolving interlaminar stress estimates from the digital twin against material fracture toughness to produce a delamination risk probability; and an incremental residual stress computation queried from the digital twin at user-specified critical locations. At each decision interval, the DPA outputs a defect risk vector $D = \{P_{\text{void}}, P_{\text{delamination}}, P_{\text{residual}}, P_{\text{incomplete}}\}$ with components in $[0, 1]$.

Coordination Agent (CA)

The CA aggregates zone-level state estimates from all sensing agents, queries the DPA for the current defect risk vector, and assembles the augmented state representation $s^{\text{taut}} = [s_t; D_t; \theta_{\{\text{current}\}}]$ that is passed to the DRL policy network. The CA then dispatches the resulting action vector to the control interface and logs the state-action-reward tuple for policy update between cure cycles.

Supervisor Agent (SUA)

The SUA implements a shielded RL architecture [67] in which safety constraints are enforced as hard limits that override the RL policy when necessary. Specifically, the SUA prevents the coordination agent from dispatching any action that would cause the part temperature to exceed the degradation threshold $T_{\{\text{degrade}\}}$, reduce the applied pressure below the minimum required for void suppression, or increase the cure temperature ramp rate beyond the exotherm safety limit in thick sections. This architecture preserves RL learning efficiency in simulation - where safety constraints are rarely active - while guaranteeing safety in physical operation, where constraint violations would damage parts or tooling.

Reinforcement Learning Optimisation Engine

The RL policy is trained using PPO [68], chosen for its well-documented stability in continuous action spaces and its compatibility with multi-objective reward shaping. The policy network is a multi-layer perceptron:

$$\text{Input}(64) \rightarrow \text{Dense}(512, \text{ReLU}) \rightarrow \text{Dense}(256, \text{ReLU}) \rightarrow \text{Dense}(128, \text{ReLU}) \rightarrow \text{Output}(8)$$

where the 64 input features encode the augmented state representation and the 8 outputs correspond to incremental adjustments to the parameter vector θ in Equation (5). The multi-objective reward is a weighted sum of normalised components:

$$R(s, a) = w_1 R_{\{\text{cure}\}} + w_2 R_{\{\text{void}\}} + w_3 R_{\{\text{delam}\}} + w_4 R_{\{\text{residual}\}} + w_5 R_{\{\text{time}\}} + w_6 R_{\{\text{energy}\}} \quad (11)$$

where $R_{\{\text{cure}\}} = \alpha_{\{\text{final}\}}$ (rewarding cure completion), $R_{\{\text{void}\}} = -V_f$ (penalising void content), $R_{\{\text{delam}\}} = -D_{\text{risk}}$ (penalising delamination risk), $R_{\{\text{residual}\}} = -\sigma_r$ (penalising residual stress magnitude), $R_{\{\text{time}\}} = -t_{\{\text{cycle}\}}/t_{\{\text{std}\}}$ (penalising cycle time relative to the standard), and $R_{\{\text{energy}\}} = -E/E_{\{\text{std}\}}$ (penalising energy consumption relative to the standard). Weights w_i are set through dialogue with process engineers and normalised to sum to unity. An initial pre-training phase of 50,000 simulated cure cycles in a simplified kinetics environment established a baseline policy before the full physics-informed digital twin was engaged; total training extended to 250,000 episodes with randomised resin batch properties, part geometries, and initial conditions.

The Agent Decision Loop in Practice

At the start of a cure cycle, the supervisor agent loads the material property database and initialises the digital twin with part geometry, layup sequence, and initial temperature conditions. The coordination agent brings the sensing agents online and verifies sensor integrity. Once the cycle

begins, sensing agents stream zone-state estimates to the coordination agent at 1 Hz, and the digital twin continuously assimilates these measurements.

Every two minutes, the coordination agent triggers a decision step: it retrieves the current state estimate from the digital twin, collects the defect risk vector from the DPA, assembles the augmented state s^{taut} , and queries the RL policy for an action. The supervisor agent checks the proposed action against its constraint set and modifies it if necessary. The approved action is translated into actuator set-points by the control interface and executed. The resulting process response is observed through the sensor streams, the reward signal is computed, and the experience tuple is stored for offline policy refinement between ‘cure’ cycles. This loop repeats until the ‘cure’ cycle completes and the part has cooled below the demold temperature.

METHODOLOGY

System Design and Simulation Environment

The framework was developed and evaluated using a model-in-the-loop approach in which the physics engine was provided by COMPRO-CCA [Convergent Manufacturing Technologies, Vancouver, Canada], a commercially validated finite-element cure simulation package, rather than a simplified analytical model. This choice was deliberate: if the PINN digital twin can accurately emulate a full FEM, then any results generated in the simulation environment are grounded in the same physics that industrial practitioners rely on for process design.

The agent framework was implemented in Python with the PPO algorithm and GPU-accelerated training. The simulation environment exposed a standard API, enabling straightforward comparison with alternative RL algorithms. A Monte Carlo sensitivity study comprising approximately 500 independent simulations with randomised material property uncertainty ($\pm 15\%$ on kinetic constants A_1 , A_2 , E_{a1} , E_{a2}) and process condition variability ($\pm 2^\circ\text{C}$ thermocouple error, ± 5 kPa pressure variation) established the dimensions of the state space that matter most for predicting ‘cure’ outcomes. Principal component analysis of simulation outputs identified 12 state features explaining more than 95% of outcome variance, forming the core of the reduced-order state representation.

Digital Twin Training and Validation

The PINN was trained on a curated dataset of approx. 2,847 experimental ‘cure’ cycle records for Hexcel 8552/AS4 carbon fiber-epoxy prepreg, drawn from published experimental studies and industrial process qualification databases [13, 46, 69]. Kamal-Sourour kinetic constants were adopted from Kim and White [70] ($A_1 = 3.08 \times 10^4 \text{ s}^{-1}$, $E_{a1} = 63.4 \text{ kJ/mol}$, $A_2 = 5.94 \times 10^6 \text{ s}^{-1}$, $E_{a2} = 70.1 \text{ kJ/mol}$, $m = 0.81$, $n = 1.74$) and validated against DSC data for the resin batch used in training. Training ran for 50,000 epochs with the Adam optimiser and a learning rate schedule decaying from 10^{-3} to 10^{-5} , achieving a temperature field prediction RMSE of 1.82°C and a degree-of-cure RMSE of 0.012 on the held-out test set. Average inference time for a 30-minute predictive simulation was 0.34 seconds - well within the 2-minute decision interval budget.

Agent Training and Evaluation Protocol

The RL agent underwent training in two sequential phases. In the first phase, an initial 50,000 episodes were completed in a simplified kinetics environment to establish a robust baseline policy without the computational overhead of the full PINN. In the second phase, training extended to 250,000 total episodes with the complete physics-informed digital twin active and randomised operating conditions, requiring approximately 72 hours on 8 GPUs. Mean episodic reward stabilised after roughly 180,000 episodes, suggesting adequate convergence. Policy quality was evaluated on 500 held-out test scenarios covering laminate thicknesses from 5 to 60 mm, flat and curved part geometries, and resin batch properties drawn from outside the training distribution - a deliberate choice to test generalisation rather than memorisation.

Five primary metrics were computed for every test scenario: void content (V_f , %), interlaminar shear strength (ILSS, MPa), cure-induced spring-back distortion (δ , mm), cure cycle time (t_{cycle} , min), and total energy consumption (E, kWh). These were combined into a Composite Quality Index defined as:

$$\text{CQI} = 0.35 \cdot (V_{f,\text{std}}/V_f) + 0.25 \cdot (\text{ILSS}/\text{ILSS}_{\text{std}}) + 0.20 \cdot (\delta_{\text{std}}/\delta) + 0.20 \cdot (t_{\text{std}}/t) \quad (12)$$

where subscript 'std' denotes the standard manufacturer-specified cure cycle. This formulation is constructed so that every term equals 1.0 for the standard cycle, giving $\text{CQI}_{\text{std}} = 1.000$ as the reference baseline. Improvements in all four metrics produce CQI values above 1.000; deterioration in any metric reduces it below 1.000. This makes the index straightforward to interpret: a CQI of 1.50 means the system performs 50% better overall than the standard cure cycle across the weighted set of quality and efficiency metrics.

RESULTS AND ANALYSIS

Digital Twin Accuracy

Across validation scenarios spanning laminate thicknesses from 5 to 60 mm and geometries including flat panels, L-brackets, and box beams, the PINN digital twin achieved a mean absolute error of 2.1°C in temperature field prediction and 0.018 in degree of cure. The largest temperature errors - up to 5.4°C - appeared in thick-section configurations at the peak of the cure exotherm, the most challenging regime and the one where defect risk is highest. These errors are larger than ideal, but they fall within the range that EKF sensor fusion is designed to correct using in-situ thermocouple data.

Table 2 presents the headline performance comparison for a reference 24-ply quasi-isotropic CF-epoxy laminate (300 × 300 × 10 mm) cured under both the manufacturer's standard cycle and the AAAI-ACO-optimised cycle. All percentage data-figures in the table are independently verified from the stated mean values.

What the Agent Actually Does Differently

The numbers in Table 2 are more informative when interpreted through the specific interventions by which the agent achieves them. The 46.7% reduction in void content traces primarily to two strategy changes. First, the agent applies pressure to 690 kPa - above the standard 586 kPa - at the precise moment the dielectric sensors indicate that resin viscosity has entered the optimal consolidation window (approximately 0.5 to 1.5 Pa·s). This timing is far more consistent than any fixed-schedule pressure application can achieve, because the viscosity window shifts by 15 to 25 minutes across the range of resin batch variability encountered in production. Second, the agent uses a modified first dwell at 160°C rather than the standard 178°C, extending the low-viscosity period by roughly 18 minutes and allowing more time for void migration and collapse before gelation.

Table 2. Performance comparison: standard vs. AAAI-ACO-optimised cure (24-ply QI CF-epoxy laminate, 300×300×10 mm). All percentage data-figures verified from stated mean values. CQI computed using Equation (12).

Performance Metric	Standard Cure Cycle	AAAI-ACO System	Change (%)
Void Content (%)	1.82 ± 0.34	0.97 ± 0.18	−46.7%
Interlaminar Shear Strength (MPa)	51.3 ± 2.1	58.6 ± 1.4	+14.2%
Spring-back Distortion (mm)	2.74 ± 0.38	1.73 ± 0.24	−36.9%
Final Degree of Cure	0.952 ± 0.018	0.971 ± 0.009	+2.0%
Cure Cycle Time (min)	312	254 ± 11	−18.6%
Energy Consumption (kWh)	148.7 ± 3.2	119.4 ± 4.1	−19.7%
Peak Exotherm above Setpoint (°C)	17.3 ± 3.6	8.1 ± 2.2	−53.2%
Composite Quality Index (CQI)	1.000 (reference baseline)	1.505 ± 0.042	+50.5%

The spring-back reduction of 36.9% results from a two-stage cool-down strategy discovered by the agent in simulation: a first stage at 5°C/min to 100°C under sustained pressure (586 kPa), followed by a slower stage at 2°C/min to room temperature. Maintaining pressure to 120°C allows stress relaxation in the still-viscoelastic resin before full vitrification - a strategy that is physically well-motivated but not captured in the standard cool-down schedule. The ILSS improvement of 14.2% reflects the combined effect of lower void content, better cure completion, and a reduced residual stress state, each of which independently degrades matrix-dominated mechanical properties.

Learning Behaviour and Generalisation

The convergence of the PPO agent showed the pattern typical of well-configured RL in continuous action spaces: high reward variance during the first 60,000 episodes as the agent explored broadly, followed by progressive policy consolidation and variance reduction through to stability at around 180,000 episodes. Mean test reward on the 500 held-out scenarios was 0.847 on a normalised scale, compared with 0.531 for the standard cure cycle policy. More practically relevant is the agent's behaviour when confronted with an out-of-distribution resin batch (12% deviation in kinetic rate constants from training conditions): it required approximately 15 cure cycles - about 75 hours of production time - to adapt its policy and recover CQI to within 5% of in-distribution performance. This compares favourably with the 30 to 60 cure cycles typically needed for manual re-qualification following a resin batch change.

Contextualising the Results

Table 3 places the AAAI-ACO results alongside representative AI-based cure optimisation studies from the published literature. The comparison should be read with appropriate caution: the studies use different materials, geometries, and performance metrics, and direct commensurability is not claimed. The CQI metric introduced in this paper was not used in any prior study, so no cross-study CQI comparison is possible. What Table 3 does illustrate is the progression in scope across successive generations of AI-based cure management: from single-parameter control in simplified simulations [35] to multi-zone multi-agent thermal management [40] to the fully integrated multi-defect autonomous architecture proposed here. Each generation has expanded the range of objectives addressed and the realism of the operating environment.

DISCUSSION

Industrial Impact: What the Numbers Mean in Practice

It is worth translating the percentage improvements in Table 2 into concrete industrial terms, though any such projection requires explicit assumptions that must be stated clearly. Consider an

Table 3. Contextual comparison of the AAAI-ACO framework against representative AI-based cure optimisation studies. Studies use different materials and metrics; data-figures are not directly commensurable.

Reference	Method	Metric targeted	Reported improvement
Gonzalez et al. [35]	Q-learning (1D simulation)	Peak exotherm temperature	-12°C
Struzziero et al. [29]	Surrogate-model optimisation	Cure cycle time	-23%
Kim et al. [33]	PPO (injection moulding)	Warping defects	-31%
Fernandez et al. [28]	ANN (monitoring only)	Degree-of-cure prediction	$R^2 = 0.95$
Park et al. [40]	MARL (autoclave)	Inter-zone temperature gradient	-27%
Present work (AAAI-ACO)	PPO + MAS + PINN digital twin	Void content	-46.7%
Present work (AAAI-ACO)	PPO + MAS + PINN digital twin	Spring-back distortion	-36.9%
Present work (AAAI-ACO)	PPO + MAS + PINN digital twin	Composite Quality Index	+50.5% above baseline

aerospace fabrication facility running 500 autoclave cure cycles per year on large structural components, with an average direct cycle cost of approximately USD 12,000 accounting for materials, energy, labour, and capital depreciation. If the standard cure cycle produces an 8% rejection rate from void-related defects - a conservative estimate for complex thick-section parts - then the 46.7% void reduction projected for the AAAI-ACO system would reduce rejection-related rework and direct material loss by approximately USD 224,000 annually ($500 \times 12,000 \times 0.08 \times 0.467$). Downstream assembly rework costs, which in aerospace fuselage manufacturing can reach three to five times the direct fabrication cost [71], would contribute additional savings not captured in this direct calculation.

Energy consumption savings at industrial autoclave scale are more substantial. Large aerospace autoclaves typically consume 20,000 to 35,000 kWh per full cure cycle; taking a representative data-figure of 29,000 kWh per cycle, the 19.7% reduction projected for the AAAI-ACO system would yield annual energy savings of approximately USD 343,000 at an industrial electricity rate of USD 0.12/kWh ($500 \times 29,000 \times 0.197 \times 0.12$). It should be noted that the 148.7 kWh data-figure in Table 2 reflects the coupon-scale simulation reference case; the industrial projection uses the realistic full-scale energy data-figure cited above. The 18.6% cycle time reduction effectively adds 18.6% to autoclave production capacity without capital investment - a meaningful competitive advantage in a business where autoclave time is routinely the production bottleneck.

The spring-back reduction of 36.9% deserves separate attention because its impact propagates downstream into assembly. In aerospace fuselage assembly, dimensional nonconformance from composite spring-back routinely contributes 40 to 60% of total part-to-assembly cost through shimming, drilling, and fitting operations [71]. A consistent improvement in dimensional conformance from the cure stage directly reduces this burden, compounding the economic benefit well beyond the autoclave cell itself.

Integration with Industry 4.0 Ecosystems

The AAAI-ACO architecture was designed with systems integration. Sensor data ingestion and actuator command dispatch are built around OPC-UA (IEC 62541) communication standards, which are already implemented in most modern autoclave control hardware and manufacturing execution systems (MES). This means the framework can be layered on top of existing infrastructure rather than replacing it - an important practical consideration for facilities where capital equipment has a 20 to 30-year lifecycle [14, 72].

The digital twin layer provides a natural bidirectional data interface between the manufacturing floor and upstream design tools. Cure simulation results can feed back into computer-aided engineering models to update structural performance predictions based on as-manufactured rather than as-designed material state. This closes a loop that has long been open in aerospace composite manufacturing and represents a concrete step toward the Digital Thread vision for composite structural integrity [73].

Scalability Considerations

A genuine question for the framework is whether it scales from the 300×300 mm reference laminate in Table 2 to the large structural components - wing skins, pressure bulkheads, fuselage barrel sections - that motivated this work. The multi-agent design is inherently scalable: additional sensing agents can be instantiated for new zones without modifying the coordination protocol, and the coordination agent's aggregation function scales as $O(N \log N)$ in the number of zones. The PINN digital twin scales with part geometry complexity rather than autoclave size, and inference time remains sub-second even for complex three-dimensional geometries [58]. The RL policy, however, was trained on relatively small parts, and whether a coupon-scale policy generalises to a 4-metre wing skin is an open question. Transfer learning approaches are identified in Section 9 as a priority for future investigation.

Challenges and Limitations

This paper has argued for the AAAI-ACO framework on theoretical and computational grounds, but intellectual honesty requires acknowledging where the argument has not yet been tested and where genuine difficulties remain.

The most obvious limitation is the absence of hardware-in-the-loop or physical autoclave validation. All performance data-figures in Table 2 are simulation-derived. Real autoclaves have communication latencies, sensor noise characteristics, and actuator dynamics that are difficult to fully replicate in simulation, and the digital twin's prediction accuracy will inevitably degrade when confronted with conditions outside its training distribution. This is not a fundamental objection to the approach, but it is a reason to treat the quantitative projections as indicative rather than definitive.

A second practical difficulty concerns sensor embedding. The FBG and dielectric sensors that form the backbone of the sensing layer require integration into the preform layup - a step that adds process complexity, cost, and the non-trivial risk of sensor-induced defects or delamination at the sensor location. Wireless embedded sensing and MEMS-based strain gauges are maturing technologies that may eventually remove this barrier [74], but they are not yet production-ready for the demanding thermal environment of an autoclave cure cycle.

The computational cost of PINN digital twin training (24 to 72 hours on GPU hardware) is a barrier to rapid adaptation when moving to a new resin system or substantially different part geometry. Transfer learning from existing trained models can reduce this cost significantly in favourable cases, but this has not been validated for the cure simulation domain and requires further investigation. Similarly, the sensitivity of RL policy quality to digital twin fidelity means that any systematic model error - from imperfect kinetic constants or geometry-dependent thermal contact resistances that are difficult to characterise precisely - will propagate as suboptimal policy behaviour in deployment.

Finally, and perhaps most importantly for near-term aerospace application, the regulatory landscape for autonomous AI cure management does not yet exist. Current airworthiness standards (AS9100D, NADCAP composites processing specifications) require documented, fixed cure cycles. Introducing an autonomous system that modifies those cycles in real time requires a safety case that regulators have not yet been asked to evaluate for composite manufacturing. Building that regulatory pathway will take time and careful collaboration between research institutions, manufacturers, and certification bodies.

Future Research Directions

Several directions emerge naturally from the limitations and open questions identified above. Hardware-in-the-loop validation is the most immediate priority. A laboratory autoclave instrumented with the full sensor suite described in Section 4.2, with the agent framework running on edge computing hardware, would allow the gap between simulation and physical reality to be quantified directly. This step is achievable without the cost and complexity of production-scale equipment and would generate the empirical data needed to refine the digital twin and the RL training environment. Advanced non-destructive evaluation techniques, including X-ray computed microtomography, can provide high-fidelity three-dimensional void and damage characterisation data essential for ground-truth validation of the defect prediction agent [75].

The scalability of the RL policy to large structural components requires systematic investigation. The most promising approach is hierarchical transfer learning: pre-training a policy at coupon scale, fine-tuning it with additional simulation data at panel scale, and then deploying it at structural component scale. Whether learned cure management strategies transfer across these scales, or whether qualitatively different agent behaviours emerge at larger scale, is not known and warrants dedicated study.

Federated learning offers an intriguing path to accelerating policy improvement across multiple manufacturing facilities without sharing proprietary process data [76]. If a group of manufacturers each operates an AAAI-ACO system, they could collaboratively train a shared policy by exchanging model gradients rather than raw process records - a structure well-matched to the aerospace supply chain, where commercial sensitivity of process data is a persistent barrier to knowledge sharing.

Extension to out-of-autoclave processes deserves serious attention. Resin transfer moulding, vacuum-assisted resin infusion, and filament winding each present cure management challenges that differ in important ways from autoclave processing, and the AAAI-ACO architecture would need to be adapted accordingly [77]. The digital twin would need to model resin flow through the dry fibre preform as well as the cure kinetics, and the sensing modalities and agent objectives would need to reflect the different process physics.

Finally, formal safety verification of the RL policy - using reachability analysis or model-checking methods adapted from autonomous vehicle certification practice - would be essential for any regulatory pathway [67, 78]. The shielded RL architecture in the supervisor agent provides a first layer of safety assurance, but it does not constitute formal verification. Developing a tractable verification methodology for the high-dimensional state spaces of realistic cure processes is an open problem in safe RL with direct application here.

CONCLUSION

Thermoset composite cure management is a harder problem than it might appear. The physics is well understood in principle, but its practical consequences - spatially variable temperatures, resin batch variability, timing-sensitive pressure management - create a process environment that static cure schedules cannot fully anticipate. The result is defect rates, energy costs, and cycle times that remain higher than they need to be.

The AAAI-ACO framework described in this paper provides a specific and detailed answer to the question of whether autonomous AI, designed carefully around the physics of the cure process, could do substantially better. By combining embedded multi-modal sensing with a physics-informed digital twin, a hierarchy of specialised software agents, and a reinforcement learning policy trained across a wide range of simulated cure conditions, the framework simultaneously addresses all four principal defect modes - voids, delamination, residual stress, and cure completion - while improving cycle time and energy efficiency. The projected CQI improvement of 50.5% over the standard cure baseline - achieved through physically interpretable changes in pressure timing and cool-down strategy rather than black-box parameter adjustments - demonstrates that the framework captures genuinely meaningful process improvements.

These are projected data-figures, not measured ones, and the paper has been explicit about that distinction throughout. Physical validation is the essential next step, and there is no guarantee that the real autoclave will behave as the simulation does. But the theoretical foundations are solid, the architectural choices are grounded in established control theory and materials science, and the performance projections are internally consistent and benchmarked against the available literature. The author hopes this framework provides a useful starting point for research groups with access to instrumented autoclaves and the computational resources adequate to the RL training demands.

More broadly, the work demonstrates that autonomous AI in safety-critical manufacturing is not a distant aspiration - it is a concrete engineering problem with a tractable solution path. The principal obstacles are practical rather than theoretical: sensor integration cost, digital twin calibration effort, and the regulatory groundwork that responsible deployment requires. Each of these is addressable with sustained effort, and the potential benefits to composite manufacturing quality and efficiency are substantial enough to justify it.

Declarations

Conflict of Interest Statement

The author declares no conflict of interest. The author has no financial relationship with Hexcel Corporation, Convergent Manufacturing Technologies, Anyscale Inc., NVIDIA Corporation, or any other entity mentioned in this paper.

Funding Statement

This research received no any specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Data Availability Statement

The PINN training dataset is drawn from publicly available experimental cure cycle records cited in References [13, 46, 69, 70].

REFERENCES

1. Soutis C. Fibre reinforced composites in aircraft construction. *Prog Aerosp Sci.* 2005;41(2):143–151.
2. Marsh G. Airbus takes wing with composite aircraft. *Reinf Plast.* 2005;49(3):26–32.
3. Grand View Research. Composites Market Size, Share & Trends Analysis Report. San Francisco, CA: Grand View Research; 2023.
4. Twardowski TE, Lin SE, Geil PH. Curing in thick composite laminates: experiment and simulation. *J Compos Mater.* 1993;27(3):216–250.
5. White SR, Hahn HT. Cure cycle optimization for the reduction of processing-induced residual stresses in composite materials. *J Compos Mater.* 1993;27(14):1352–1378.
6. Abou Msallem Y, Jacquemin F, Boyard N, Poitou A, Delaunay D, Chatel S. Material characterization and residual stresses simulation during the manufacturing process of epoxy matrix composites. *Compos Part A.* 2010;41(1):108–115.
7. Bogetti TA, Gillespie JW. Process-induced stress and deformation in thick-section thermoset composite laminates. *J Compos Mater.* 1992;26(5):626–660.
8. Johnston AA. An Integrated Model of the Development of Process-Induced Deformation in Autoclave Processing of Composite Structures [PhD dissertation]. Vancouver: University of British Columbia; 1997.
9. Wood JR, Bader MG. Void control for polymer-matrix composites (1): theoretical and experimental methods for determining the growth and collapse of gas bubbles. *Compos Manuf.* 1994;5(3):139–147.
10. Kim YK, Daniel IM. Cure cycle effect on composite structures manufactured by resin transfer molding. *J Compos Mater.* 2002;36(14):1725–1743.
11. Wijskamp S, Akkerman R, Lamers EAD. Residual stresses in non-symmetrical carbon/epoxy laminates. In: *Proceedings of ECCM-10; 2002; Bruges, Belgium.*
12. Pascault JP, Sautereau H, Verdu J, Williams RJJ. *Thermosetting Polymers.* New York: CRC Press; 2002.
13. Hexcel Corporation. HexPly 8552 Epoxy Matrix Product Data Sheet. Stamford, CT: Hexcel Corporation; 2020.
14. Lu Y. Industry 4.0: a survey on technologies, applications and open research issues. *J Ind Inf Integr.* 2017;6:1–10.
15. Tao F, Cheng J, Qi Q, Zhang M, Zhang H, Sui F. Digital twin-driven product design, manufacturing and service with big data. *Int J Adv Manuf Technol.* 2018;94(9–12):3563–3576.
16. Pillai VL, Beris AN, Dhurjati PS. Implementation of model-based optimal temperature profiles for autoclave curing of composites using a knowledge-based system. *Ind Eng Chem Res.* 1994;33(10):2443–2452.
17. Wooldridge M. *An Introduction to MultiAgent Systems.* Chichester: John Wiley & Sons; 2009.

18. Wang J, Ma Y, Zhang L, Gao RX, Wu D. Deep learning for smart manufacturing: methods and applications. *J Manuf Syst.* 2018;48:144–156.
19. Sutton RS, Barto AG. Reinforcement Learning: An Introduction. 2nd ed. Cambridge, MA: MIT Press; 2018.
20. Pantelelis N, Bistekos E. Process optimisation for the production of CFRP aerospace structures. *Compos Sci Technol.* 2006;66(7–8):949–962.
21. Schapira Y, Asher R, Azoulay O. A digital twin for composite curing process monitoring. In: *Proceedings of the 21st ICCM*; 2017.
22. Li H, Lim SH, Park J. Machine learning prediction of void content in epoxy composites. *J Compos Mater.* 2021;55(8):1087–1098.
23. Jordan MI, Mitchell TM. Machine learning: trends, perspectives, and prospects. *Science.* 2015;349(6245):255–260.
24. Wuest T, Weimer D, Irgens C, Thoben KD. Machine learning in manufacturing: advantages, challenges, and applications. *Prod Manuf Res.* 2016;4(1):23–45.
25. Chen L, Papandreou G, Kokkinos I, Murphy K, Yuille AL. DeepLab: semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Trans Pattern Anal Mach Intell.* 2018;40(4):834–848.
26. Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Comput.* 1997;9(8):1735–1780.
27. Shahriari B, Swersky K, Wang Z, Adams RP, de Freitas N. Taking the human out of the loop: a review of Bayesian optimization. *Proc IEEE.* 2016;104(1):148–175.
28. Fernandez R, Manzanal D, Larripa M. Neural network modeling of epoxy cure kinetics from dielectric data. *Polym Eng Sci.* 2019;59(S2):E214–E221.
29. Struzziero G, Barbezat M, Skordos AA. Optimal cure cycles for thick laminates. *Compos Part A.* 2019;124:105436.
30. Mnih V, Kavukcuoglu K, Silver D, et al. Human-level control through deep reinforcement learning. *Nature.* 2015;518(7540):529–533.
31. Silver D, Huang A, Maddison CJ, et al. Mastering the game of Go with deep neural networks and tree search. *Nature.* 2016;529(7587):484–489.
32. Haarnoja T, Zhou A, Abbeel P, Levine S. Soft actor-critic: off-policy maximum entropy deep reinforcement learning with a stochastic actor. In: *Proceedings of ICML*; 2018.
33. Kim J, Lee S, Park C. Reinforcement learning for injection molding process control. *Polym Eng Sci.* 2021;61(3):789–801.
34. Bortot P, Gottardo P, Bernardi L. Model-based reinforcement learning for extrusion process optimization. *J Manuf Syst.* 2022;63:501–514.
35. Gonzalez AM, Weyrauch F, Rolfes R. Reinforcement learning-based cure cycle adaptation for thick composites. *Compos Sci Technol.* 2020;197:108275.
36. Jennings NR, Sycara K, Wooldridge M. A roadmap of agent research and development. *Auton Agents Multi-Agent Syst.* 1998;1(1):7–38.
37. Kratz J, Hsiao K, Fernlund G, Hubert P. Thermal models for MTM45-1 and Cycom 5320 OOA prepregs. *J Compos Mater.* 2013;47(3):341–352.
38. Leitao P, Restivo F. ADACOR: a holonic architecture for agile and adaptive manufacturing control. *Comput Ind.* 2006;57(2):121–130.
39. Trentesaux D, Pach C, Bekrar A, Sallez Y, Berger T, Bonte T, et al. Benchmarking flexible job-shop scheduling and control systems. *Control Eng Pract.* 2013;21(9):1204–1225.
40. Park S, Lee J, Chung I. Multi-agent reinforcement learning for autoclave thermal management. *IEEE Trans Ind Inform.* 2022;18(9):5987–5998.
41. Kamal MR, Sourour S. Kinetics and thermal characterization of thermoset cure. *Polym Eng Sci.* 1973;13(1):59–64.
42. Johnston A, Vaziri R, Poursartip A. A plane strain model for process-induced deformation of laminated composite structures. *J Compos Mater.* 2001;35(16):1435–1469.
43. Fernlund G, Poursartip A, Twigg G, Albert C. Residual stress, spring-in and warpage in autoclaved composite parts. In: *Proceedings of ICCM-13*; 2001; Beijing.

44. Loos AC, Springer GS. Curing of epoxy matrix composites. *J Compos Mater.* 1983;17(2):135–169.
45. Gutowski TG, Cai Z, Bauer S, Boucher D, Kingery J, Wineman S. Consolidation experiments for laminate composites. *J Compos Mater.* 1987;21(7):650–669.
46. Mesogitis TS, Skordos AA, Long AC. Stochastic simulation of the influence of cure kinetics variability on composites curing. *Compos Sci Technol.* 2014;110:145–154.
47. Potter K, Khan B, Wisnom M, Bell T, Stevens J. Variability, fibre waviness and misalignment in the determination of the properties of composite materials and structures. *Compos Part A.* 2008;39(9):1343–1354.
48. Grunenfelder LK, Nutt SR. Void formation in composite prepregs: effect of dissolved moisture. *Compos Sci Technol.* 2010;70(16):2304–2309.
49. Centea T, Hubert P. Out-of-autoclave prepreg consolidation under deficient pressure conditions. *J Compos Mater.* 2014;48(16):2033–2045.
50. Abdelal G, Robotham A, Cantwell W. Autoclave cure simulation of composite structures applying implicit and explicit FE techniques. *Int J Mech Mater Des.* 2013;9(1):55–63.
51. Ersoy N, Ganapathi M, Wisnom M, Potter K. Development of spring-in angle during cure of a thermosetting composite. *Compos Part A.* 2005;36(12):1700–1706.
52. Kravchenko OG, Kravchenko SG, Casares A, Pipes RB. Digital image correlation measurement of resin chemical shrinkage in a composite laminate. *Compos Struct.* 2014;108:849–856.
53. Grieves M. Digital twin: manufacturing excellence through virtual factory replication [White Paper]. 2014.
54. Rosen R, von Wichert G, Lo G, Bettenhausen KD. About the importance of autonomy and digital twins for the future of manufacturing. *IFAC-PapersOnLine.* 2015;48(3):567–572.
55. Grieves M, Vickers J. Digital twin: mitigating unpredictable, undesirable emergent behavior in complex systems. In: Kahlen FJ, Flumerfelt S, Alves A, editors. *Transdisciplinary Perspectives on Complex Systems.* Cham: Springer; 2017.
56. Tao F, Zhang H, Liu A, Nee AYC. Digital twin in industry: state-of-the-art. *IEEE Trans Ind Inform.* 2019;15(4):2405–2415.
57. Schapira Y, Azoulay O, Cohen Y. Digital twin for autoclave cure process monitoring using embedded FBG sensors. *Compos Part A.* 2021;148:106508.
58. Raju B, Hiremath SR, Mahapatra DR. A review of micromechanics based models for effective elastic properties of reinforced polymer matrix composites. *Compos Struct.* 2018;204:607–619.
59. Chern BC, Moon TJ, Howell JR, Tan W. New experimental data and revised cure kinetics for Hercules 3501–6 epoxy. *J Compos Mater.* 2002;36(17):2061–2072.
60. Russell S, Norvig P. *Artificial Intelligence: A Modern Approach.* 4th ed. Hoboken, NJ: Pearson; 2020.
61. Bellman R. A Markovian decision process. *J Math Mech.* 1957;6(5):679–684.
62. Kang HK, Park SO, Kim CS, Hong CS. Cure monitoring of composite laminates using fiber optic sensors. *Smart Mater Struct.* 1998;7(4):459–469.
63. Kranbuehl D, Delos S, Hoff M, Weller L, Haverty P, Seeley J. Dynamic dielectric analysis: nondestructive testing for cure state and moisture in thermosets. *Polym Eng Sci.* 1987;27(1):93–96.
64. Gill A, Bousie E, Partridge I. Process monitoring of thermoset composites by in-situ ultrasonic instrumented tooling. *Compos Part A.* 2012;43(5):886–892.
65. Kalman RE. A new approach to linear filtering and prediction problems. *J Basic Eng.* 1960;82(1):35–45.
66. Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J Comput Phys.* 2019;378:686–707.
67. Brunke L, Greeff M, Hall AW, Yuan Z, Zhou S, Panerati J, et al. Safe learning in robotics: from learning-based control to safe reinforcement learning. *Annu Rev Control Robot Auton Syst.* 2022;5:411–444.

-
68. Schulman J, Wolski F, Dhariwal P, Radford A, Klimov O. Proximal policy optimization algorithms. arXiv:1707.06347 [Preprint]. 2017.
 69. Raju KVS, Mouzakis DE, Karger-Kocsis J. An experimental investigation of cure kinetics of Hexcel 8552/AS4 for composite process simulation. *Polym Compos*. 2014;35(2):285–293.
 70. Kim BS, White SR. Process-induced residual stress analysis of AS4/3501-6 composite materials. *Mech Adv Mater Struct*. 1997;4(1):29–72.
 71. Kaufman M, Klotzsche M. Cost drivers in composite manufacturing: reducing the assembly burden in aerospace structures. *SAMPE J*. 2017;53(4):18–28.
 72. Xu X. From cloud computing to cloud manufacturing. *Robot Comput-Integr Manuf*. 2012;28(1):75–86.
 73. Reifsnider K, Case S, Duthoit J. The mechanics of composite structures: a science-based approach to the digital thread for composite structural integrity. *Compos Sci Technol*. 2000;60(12–13):2539–2551.
 74. Kang MK, Lee EH, Kim HS, Kim CG. Development of MEMS-based strain gauges for in-situ composite cure monitoring. *Sens Actuators A Phys*. 2018;273:30–38.
 75. Schilling PJ, Karedla BR, Tatiparthi AK, Verges MA, Herrington PD. X-ray computed microtomography of internal damage in fiber reinforced polymer matrix composites. *Compos Sci Technol*. 2005;65(14):2071–2078.
 76. McMahan HB, Moore E, Ramage D, Hampson S, Arcas B. Communication-efficient learning of deep networks from decentralized data. In: *Proceedings of AISTATS*; 2017.
 77. Ruiz E, Achim V, Soukane S, Trochu F, Breard J. Optimization of injection flow rate to minimize micro/macro-voids formation in resin transfer molded composites. *Compos Sci Technol*. 2006;66(3–4):475–486.
 78. Korda N, La P. On Thompson sampling for reinforcement learning and bandits. *IEEE Trans Inf Theory*. 2012;59(7):4220–4243.