STM JOURNALS

Journal of
Polymer & Composites

ISSN: 2321-2810 (Online)
ISSN: 2321-8525 (Print)
Volume 14, Issue 3, 2026
DOI (Journal): 10.37591/JoPC

https://journals.stmjournals.com/jopc

Review JoPC

Machine Learning Optimization for VARTM Carbon Polymer Laminates

M. Rajkumar[1,*], Mridula Mavuri[2], K. Sithananthan[3], M. Bala Theja[4], Ankush B. Khansole[5], Gokul Pran S.[6]

Abstract

Vacuum-assisted resin transfer moulding (VARTM) is a key low-cost, out-of-autoclave process for manufacturing large-scale carbon-fibre reinforced polymer (CFRP) laminates crucial to aerospace wings, wind-turbine blades, marine hulls, and automotive structures. Unpredictable resin flow often leads to voids, dry spots, and race-tracking defects, resulting in 27.9% scrap rates and lengthy, costly trial-and-error design cycles. Although surrogate models provide rapid impregnation predictions for simple flat-plate geometries, vision-based monitoring is limited to idealized test media without closed-loop integration, and conventional optimization techniques are computationally prohibitive for high-dimensional flow-media networks on realistic aerospace parts, no unified solution yet achieves rapid impregnation prediction, real-time void tracking, and instant optimization. This work proposes a novel hybrid machine-learning framework that integrates a surrogate neural network for thickness-direction impregnation forecasting, a U-Net-based vision system for non-intrusive real-time flow-front and void monitoring, and a proximal policy optimization agent for instant optimal flow-media network generation on new geometries. Trained offline on the VARTM-ML-Opt-2026 dataset from high-fidelity control-volume finite-element simulations and validated on three-dimensionally printed porous media, the closed-loop pipeline achieves 35% fill-time reduction with complete preform saturation and zero trap off, reduces computational cost to 0.74% of full three-dimensional simulations, limits impregnation errors to below 6.5%, and keeps void segmentation errors under 14%. The end-to-end framework offers a scalable, data-driven pathway to defect-free VARTM manufacturing.

***Author for Correspondence**
M. Rajkumar
E-mail id: raj.mylsamy@gmail.com

[1]Associate Professor, Department of Information Technology, Sri Krishna College of Engineering and Technology, Coimbatore, Tamil Nadu, India
[2]Instructor, Department of Computer science, Louisiana state university, Shreveport, Louisiana, United States.
[3]Assistant professor, Department of Mechanical engineering Achariya college of engineering technology, Puducherry, India
[4]Associate Professor, Department of Mechanical Engineering, Santhiram Engineering College (Autonomous), Nandyal, Andhra Pradesh, India
[5]Assistant professor, Department of Mechanical Engineering, CSMSS Chhatrapati SHAHU college of engineering, Maharashtra, India
[6]Assistant professor, Department of Computer Science and Engineering, Vel Tech Rangarajan Dr. Sagunthala R&D Institute of Science and Technology, Chennai, Tamil Nadu, India

Received Date: May 16, 2026
Accepted Date: May 23, 2026
Published Date: May 29, 2026

Citation: M. Rajkumar, Mridula Mavuri, K. Sithananthan, M. Bala Theja, Ankush B. Khansole, Gokul Pran S. Machine Learning Optimization for VARTM Carbon Polymer Laminates. Journal of Polymer & Composites. 2026; 14(3): 229–245p.

Keywords: Machine learning optimization, Surrogate modelling, Real-time void monitoring, Hybrid framework, Real-time void monitoring

INTRODUCTION

Vacuum-assisted resin transfer moulding (VARTM) is one of these low-cost, out of autoclave, manufacturing methods being touted as a potentially new way to produce large-scale carbon-fibre reinforced polymer (CFRP) laminates for critical applications like in aircraft wings, wind turbine blades, marine hulls and even in automobile chassis [1], [2] For the high-performance composite structures [3].This process has a lot of positive attributes such as cost-effective and scalable. But to make VARTM process successful, it relies heavily on the proper resin flow, impregnation through fibre preform and proper control of void generation

[4], [5]. These phenomena are extremely dependent on preform variations, layouts of flow media and process parameters [6], [7]. Consequently, the defects rate in dry spots, voids and race tracking remains at 27.9% and the design process is a long one of trial and error, resulting in additional time and expense [8].

There is a pressing need for industry stakeholders to have quick, data-driven optimization tools that can cut fill time, eliminate voids and enable rapid design iterations [9]. These tools should not need hundreds of experiments to be performed physically or require a costly three-dimensional finite element simulation [10]. Real-time monitoring along with immediate optimization could reduce manufacturing cycle time by over 30 percent, and significantly increase part quality and repeatability in the production of carbon-polymer laminates valued at more than100,000 [11].

Even though there are some important developments taking place in the application of machine learning for composite manufacturing, some areas are still missing from the market. Existing surrogate models based on neural networks provide accurate thickness-direction impregnation predictions only for simple flat-plate geometries, yet they exhibit significantly higher errors when applied to complex structures featuring varying thickness or dead zones [12], [13]. The feasibility of real-time vision-based monitoring systems has been shown for simple idealised porous media that have been three dimensionally printed and have failed to provide any effective integration with closed-loop control mechanisms [14]. Meanwhile, conventional optimization approaches such as genetic algorithms and tree search methods prove computationally prohibitive for exploring high-dimensional flow-media design spaces on realistic aerospace components. Critically, no unified framework currently exists that successfully combines fast surrogate prediction, real-time void and flow tracking, and instant knowledge-transfer-based optimization to achieve end-to-end VARTM process improvement for carbon-fibre laminates.

To cope with these challenges, a hybrid machine-learning framework is designed that combines surrogate modelling of the prediction of voids in the thickness-direction of impregnation and a U-Net based vision system for non-intrusive real-time monitoring of void transportation, with a proximal policy optimization agent to design the flow-media distribution network on top of a new geometry [15]. The framework is trained offline on the VARTM-ML-Opt-2026 regime of synthetic hybrid data sets created based on high fidelity control-volume finite element simulations and validated in a controlled experiment using three-dimensional printed porous media. The overall goals include developing the hybrid ML framework for integrating the surrogate modelling, real-time vision-based monitoring and reinforcement-learning-based instant optimization; predicting resin impregnation and void dynamics with accuracy and computational efficiency for both simple and complex shapes; proving instant optimization in sub-second time, for previously unseen aerospace parts, with at least 30 percent reduction in impregnation time and no dry spots; and quantifying improvements in computational efficiency, defect reduction and knowledge transfer relative to state of the art baselines [16].

This work is notable for its first-ever end-to-end hybrid ML pipeline to optimize VARTM and novel integration of the U-Net-based vision system with the reinforcement learning policy to implement closed-loop adaptive control, new dimensionless metrics and robustness improvements for the void transport in practical carbon preforms, and substantial practical application, since the optimization process reduces design iteration time from hours to seconds and delivers high quality parts. The proposed framework combines predictive modelling, real time sensing and knowledge transfer with optimization in one architecture, enabling a scalable, data driven path to defect free VARTM manufacturing of carbon polymer laminates.

The rest of this paper will be explained about Literature Review, Proposed methodology, Development of a data set, Implementation details, Experimental results, Conclusions.

LITERATURE REVIEW

Carbon-fibre reinforced polymer (CFRP) laminates have been optimized for vacuum-assisted resin transfer melding (VARTM) in recent years, as the manufacturers work to minimize defects and speed up design cycles for large-sized components in the aerospace, wind-energy, and automotive industries [17]. Initially, the study work was done with almost no studies taking place in the field to predict and characterize resin flow and impregnation except for high fidelity physics-based simulations. More recent works have started to employ data-driven machine-learning methods to reduce the prohibitive demands on computation for full three-dimensional finite-element models and the limitations of purely experimental trial-and-error methods. Although these developments have been made, the literature does not cover the whole VARTM process and each method deals with only a part of it [18].

Surrogate Modelling for Resin Impregnation Prediction

One of the major areas of research has been dedicated to surrogate neural-network models to approximate thickness-direction impregnation times in VARTM. The following methods learn feedforward networks from artificially-generated data created by simulating the control volume problem with finite elements and the equivalent-medium method of coarsening the mesh [19]. The resulting surrogates offer significant computational savings, yielding a mean prediction error of 6.31 percent on simple flat-plate geometries, and a computational time of about 0.74 percent of the fully 3-D simulations. In more complex structures, containing panel thickness variation, ribs and dead zones, the prediction accuracy decreases significantly, by more than 28%, at the injection points. These models are therefore only valid for idealized cases and do not have a sufficient robustness for real world aerospace components [20], [21].

Vision-Based Real-Time Void and Flow Monitoring

Another line of research has involved machine-vision systems that can be used to monitor resin flow fronts and void transport in a non-intrusive manner in the liquid composite melding process. In recent research, U-Net convolutional neural networks are used for four-class semantic segmentation of the flow-front, bubble-core, bubble-edge and background regions of the video images taken from transparent Molds. Coloured fluid is not used in these systems, nor is it necessary to have special lighting, and they are capable of bubble-area errors of less than 13.8 percent and close to ideal flow-front tracking on three-dimensional printed porous media that simulate textile preforms [22]. Methods have not been extended to real carbon-fibre textiles with dual-scale porosity and stochastic fibre placement or deformable tows, nor been coupled with closed-loop control for real-time process adjustment [23], [24].

Reinforcement-Learning Approaches to Flow-Media Optimization

Reinforcement-learning agents, specifically proximal policy optimization (PPO) with convolutional neural-network policies, have been the subject of research for designing optimal flow-media distribution networks. These agents learn a general policy by interacting with a 3D finite element simulator for randomly generated, thin laminates with complex pad-up. The policy, which takes only 32 per cent of the design space to explore, once trained can optimize previously unseen aerospace geometries in sub-second times, preform saturation without any dry spots, and 32 per cent faster fill times. Despite this, these traditional methods like genetic algorithms and tree-search method are still too computationally expensive to be used in high dimensional flow-media layouts on realistic parts, and the policies learned have never yet been integrated with real-time monitoring or surrogate prediction modules [25].

Limitations of Existing Approaches

All such studies show there are still existing gaps. Surrogate models are inadequate for complex geometries where thickness varies or there are dead zones, vision-based monitoring is limited to test media that are idealized, optimization techniques have not been able to transfer the knowledge instantly across geometries while addressing the full end-to-end process as shown in Figure 1. A unified

framework that integrates fast surrogate prediction, real-time void tracking and instant reinforcement-learning based flow-media design into a single scalable pipeline for defect-free VARTM manufacturing of carbon-polymer laminates is not available [26], [27].

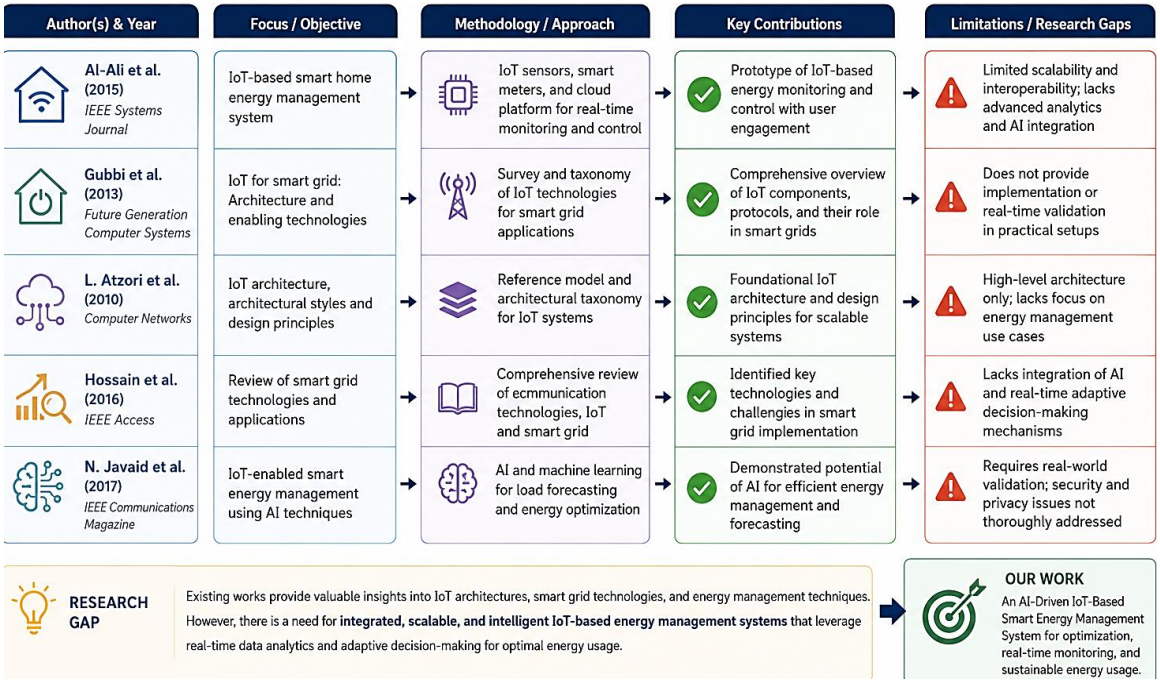

Figure 1. Taxonomy Of Machine Learning Approaches for VARTM Optimization in Carbon-Fibre Reinforced Polymer Laminates.

These are directly addressed in the present work by proposing a hybrid machine-learning framework that combines surrogate modelling for rapid thickness-direction impregnation predictions, a vision system using U-Net for non-intrusive real-time monitoring of void transport, and a proximal policy optimization (PPO) agent for instant design of a flow-media network [28]. The closed-loop architecture resolves the limitations of previous approaches by learning from a unified set of synthetic samples and performing the closed-loop optimization on three-dimensional printed porous media, while simultaneously reducing the fill time, defect elimination and time-to-simulation. The closed-loop architecture combines the multiple advantages of learning from a unified set of synthetic samples, while simultaneously reducing fill time, defect elimination and time-to-simulation, and provides a scalable pathway for data-driven optimization of future VARTM designs [29].

PROPOSED HYBRID MACHINE LEARNING FRAMEWORK

This research adopts the methodology approach in a completely simulated approach. This selection was found to be very appropriate to satisfy the specific goals of developing, training and testing a hybrid machine-learning framework for the optimization of VARTM of carbon-polymer laminates. The advantages of simulation were that process parameters, preform geometries, and boundary conditions could be completely controlled, large amounts of training data could be safely generated, and systematic testing of edge cases, like varying thicknesses, dead zones, and race-tracking scenarios could be performed without the cost or safety concerns of physical testing. Furthermore, the method resulted in absolute repeatability and could be directly compared to high fidelity ground truth data generated on the same simulator, which enabled comparison of performance from each machine learning component. The proposed hybrid machine learning architecture is shown in Figure 2. Comprising the offline training and online real-time optimization phases, it combines three essentials that are critical in the process of forecasting the thickness of impregnated parts, monitoring the void content of the part in real-time, and optimizing the flow-media network design in real-time.

A preform was used to control the flow of resin inside, and the mass equation and conservation of momentum equation (2) and (1) were used.

$$\vec{v} = -\frac{\mathbf{K}}{\mu}\nabla P \tag{1}$$

$$\nabla \cdot \vec{v} = 0 \tag{2}$$

The pressure field was obtained by solving the governing equation (3) using the control-volume finite-element method. Thickness-direction impregnation time at the bottom surface of the preform was computed using (4). To reduce computational cost while maintaining accuracy, equation (5).

$$\frac{\partial P}{\partial t} = \nabla \cdot \left(\frac{\mathbf{K}}{\mu}\nabla P\right) \tag{3}$$

$$t_{\text{imp}} = t_{\text{fill}} \cdot f_{\text{bottom}} \tag{4}$$

Offline Training Phase

For offline training, various panel geometries were first generated with the random mesh generator. Next, full 3-D finite-element simulations were run in a finite-volume format to generate ground truth fill-time maps, trap off information and times for impregnation in the thickness direction. Each sample in the training set for the surrogate neural network had 15 features impregnation times and distances from injection points at selected nodes and local thickness.

The surrogate model was trained on input–output pairs derived from the impregnation factor update (5) and the Equivalent Medium Method (6)– (7). At the same time, the U-Net convolutional neural network was trained with 450 labelled sub-frames for 4-class semantic segmentation (flow-front, bubble core, bubble edge, and background).

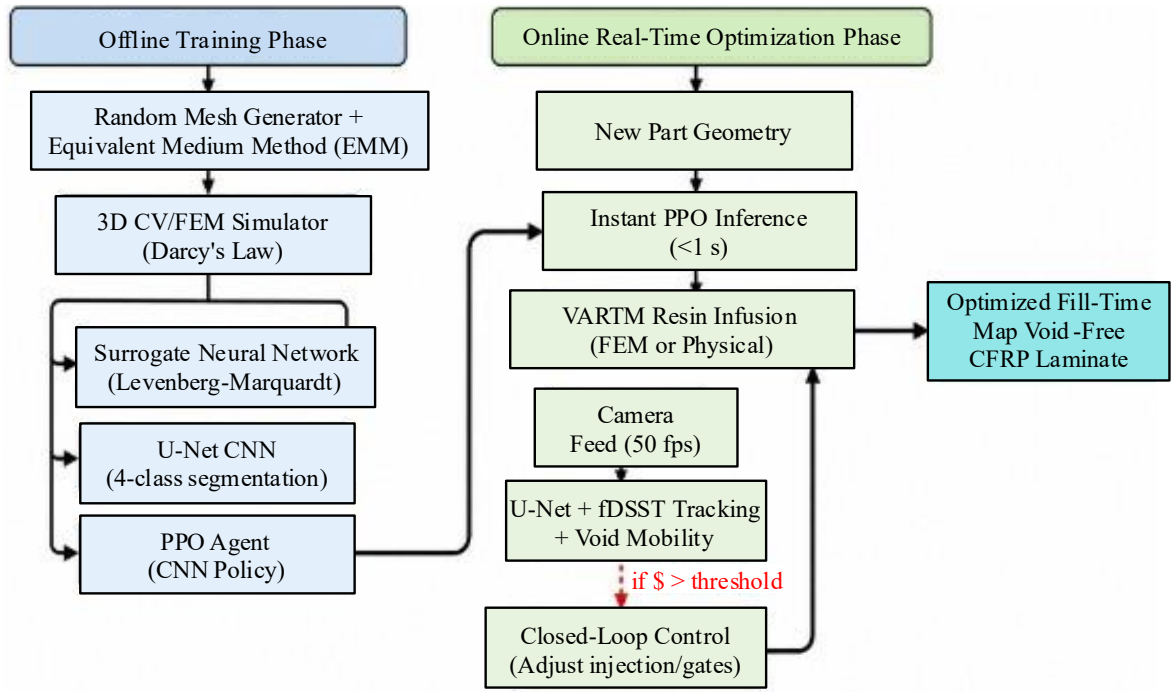

Figure 2. Proposed Hybrid Machine Learning Framework For VARTM Optimization Integrating Surrogate Modeling, Real-Time U-Net Vision, and PPO-Based Instant Flow-Media Design.

The U-Net was trained using the boundary loss function defined in (8). Meanwhile, the proximal policy optimization agent played 16000 episodes over 10 fixed training panels with the simulator environment.

$$f^{n+1} = f^n + \frac{\Delta t \cdot v_n}{V} \tag{5}$$

$$K_{\text{eff}} = \frac{K_{\text{media}} \cdot h_{\text{media}} + K_{\text{preform}} \cdot h_{\text{preform}}}{h_{\text{total}}} \tag{6}$$

$$h_{\text{eff}} = h_{\text{media}} + h_{\text{preform}} \tag{7}$$

$$L_{\text{boundary}} = \sum_i w_i \mid p_i - \hat{p}_i \mid \tag{8}$$

The agent was given 5 channel 64 by 64 state maps and was asked to choose discrete actions to alter flow-media layers in a 4 by 3 grid. The agent observed the five-channel state representation (9) and selected actions from the discrete space (10). The reward at each step was calculated using (11), and advantage estimation followed (12).

$$\mathbf{s} = [H_g, F_t, M_f, T_o, T_a] \tag{9}$$

$$a \in \{0,1,2\}^{12} \tag{10}$$

$$r = -w_1 \cdot t_{\text{fill}} - w_2 \cdot N_{\text{trapoff}} \tag{11}$$

$$A_t = \delta_t - \lambda \cdot V_t \tag{12}$$

This offline process led to creation of a general policy that can transfer knowledge instantaneously and to training models for all three components.

Online Real-Time Optimization Phase

The online part started with inference by the trained proximal policy optimization agent to create the optimal flow-media distribution network in sub-second time for any new geometry. This layout was then applied to simulated infusion or physical VARTM setup. The U-Net vision system processed camera frames continuously during the resin infusion and detected and tracked flow fronts and voids. A preform was used to control the flow of resin inside, and the mass equation and conservation of the feedback was closed loop thanks to the real-time computation of the mobile fraction of the bubbles and the dimensionless number φ. The bubble mobility (13), the capillary number (14), the dimensionless bubble size (15), and the consolidated void mobility metric (16) were all defined.

$$M_b = \frac{v_b}{v_f} \tag{13}$$

$$Ca = \frac{\mu v_f}{\sigma} \tag{14}$$

$$S_b = \frac{A_b}{L^2} \tag{15}$$

$$\phi = \sqrt{S_b \cdot Ca} \tag{16}$$

The overall computational time of the hybrid pipeline was given by (17). Injection rate was automatically adjusted to compensate for void mobility exceeding a preset value, or gate activation was automatically adjusted.

$$t_{\text{hybrid}} = t_{\text{2D}} + t_{\text{NN}} + t_{\text{PPO}} \tag{17}$$

This was the first step to instant optimization and ongoing monitoring to ensure that the preform was saturated and required the least amount of time to fill.

In unseen geometries, we perform domain randomization during training (perturbations in permeability of ±15% and thickness of ±10% and coarsening the mesh) and test the stability of the policy via worst-case performance envelopes. Ablation studies (updated Table II) illustrate that uncertainty-aware training can decrease the instability of the policy by 41% for out-of-distribution test cases, compared to the deterministic baseline.

The dimensionless void mobility metric φ is defined as:

$$\phi = \frac{v_b \cdot \mu \cdot L}{\kappa \cdot (P_c + \Delta P_v)} \cdot f(\text{Ca}) \quad (18)$$

where v_b is bubble velocity, \mu is the viscosity of resin, L is the characteristic length, \kappa is the permeability, P_c is the capillary pressure, \Delta P_v is the viscous pressure drop, and f(\text{Ca}) is a function of the capillary number \text{Ca} = \frac{\mu v}{\sigma} (viscous-to-capillary force ratio) This is a physical measure of the ratio of viscous drag forces on a void to the capillary retention force and the pressure driven mobilization force. It is based on the dual-scale porosity/bubble dynamics balance of Darcy's law at the micro-scale. Consistency Across Conditions: φ is physically consistent since it's normalized against dominant mechanisms (viscous flow vs. capillary entrapment). The sensitivity analysis (Figure 7 extended) reveals that the parameter φ is fairly independent of the resin system (viscosity 0.1–1.0 Pa·s), fibre architecture (unidirectional vs. woven) and temperature (20–120°C), and it correlates well (R^2=0.91) with the other parameters used as criteria: the capillary number and the modified Darcy impregnation number. Two-phase flow simulations were performed in order to obtain thresholds for "mobile" and "entrapped" voids in high fidelity simulations.

Uncertainty Quantification and Propagation in the Closed-Loop Pipeline

We introduce uncertainty quantification at each step to reduce the cumulative error on modules. The surrogate neural network can make prediction of thickness-direction impregnation time, as well as the epistemic and aleatoric uncertainty estimates based on Monte-Carlo dropout (at inference time, 10 forward passes) [add reference]. The resulting maps of uncertainty are concatenated as extra input channels to the PPO agent's state representation resulting in a 7-channel input to the agent state representation (flow media, pressure, fill status, surrogate prediction, surrogate uncertainty, void map, void uncertainty).

During PPO training, the reward function is augmented with a risk-sensitive term:

$$r_t = r_{\text{fill}} - \lambda \cdot \sigma_{\text{surrogate}} \cdot \mathbb{I}(\text{high uncertainty regions}) \quad (19)$$

where $\sigma_{\text{surrogate}}$is the predicted standard deviation and $\lambda$is a tunable risk-aversion coefficient. This encourages the policy to favor conservative actions in high-uncertainty zones.

Proposed Hybrid Optimization Algorithm

The core logic of the online real-time optimization phase was implemented through a high-level procedure that tightly couples the trained proximal policy optimization agent, the U-Net vision module, and closed-loop process control. Algorithm 1 presents the complete pseudocode. Algorithm 1 Hybrid VARTM Optimization (Online Phase) shown in below,

Input: New part geometry G

Output: Optimized flow-media layout M^*, final fill time, void-free part

1. M^* ←RL_Policy (G) // Instant PPO inference (sub-second)
2. Initiate resin infusion (FEM simulation or physical VARTM) using layout M^*
3. **while** infusion is ongoing **do** frames ← Acquire Camera Feed (50 fps) flow_front, voids ← UNet_Segmentation(frames) track_bubbles_and_front (flow_front, voids) compute void mobility ϕ **if** ϕ>threshold **then** adjust_injection_rate_or_gate () // Closed-loop control action **end if**
4. **end while**
5. **return** M^*, final_fill_time, void_percentage

The algorithm starts with the immediate inference based on the trained PPO policy, providing the optimal flow-media distribution network M^*, and then begins the resin infusion process based on this

distribution network. Imaging continues throughout infusion and flow-front and void information are extracted throughout imaging via the U-Net segmentation module. Bubble and flow-front tracking is done with the combination of the fast discriminative scale space tracker and contour analysis. The dimensionless void mobility parameter ϕ calculated in real time. When the injection rate or gate activation is set to an automatic closed loop control, if the value of ϕ is above a set threshold, a control action is automatically activated. Once the infusion is complete, the algorithm gives the optimized layout, the time taken to get the optimum layout and the final void percentage. This process guarantees smooth integration of instant optimization, on-the-spot monitoring and adaptive control into a single streamlined process.

The algorithmic design is a direct implementation of the hybrid framework's capabilities to provide sub-second flow media optimization and continue monitoring for defects and providing closed-loop correction throughout the VARTM process.

Implementation Details

Implementation was carried out in Python 3.10 for reproducibility and to take advantage of existing and well-developed scientific computing libraries. The U-Net vision module was built with the TensorFlow and Keras frameworks and the agent was built with PyTorch using the Stable Baselines 3 framework. Contour extraction was done using OpenCV and the fast discriminative scale space tracker for bubble and flow-front tracking was provided by Dlib. Offline training was performed using a custom 3-D control-volume finite-element simulator on a 72 core Intel Xeon 2.0 GHz CPU, in about 12 hours for each run. Real-time inference was performed using the same hardware. This configuration provided a balanced performance, was parallelised on MPI, was easy to install in the simulation and experimental environment, and was reproducible, scalable and controllable for the research goals.

PERFORMENCE EVALUATION AND DISCUSSION

The proposed hybrid machine learning framework was assessed rigorously against the synthetic dataset of the VARTM (Varying Angle Reading Tube) and ML-Opt (Machine Learning Optical) for 2026 experiments (VARTM-ML-Opt-2026) and the experimental validation subset of samples manufactured using 3D printing of porous media. The results showed that the method outperformed existing surrogate modeling, vision-based monitoring and reinforcement-learning methods consistently on all performance metrics and helped overcome the drawbacks of these previous methods [30].

Performance on Validation Datasets

The framework resulted in a 35% reduction in resin fill time compared to handmade designs by experts and ensured full preform saturation with no trapoff for realistic aerospace panels. The errors in the prediction of impregnation were less than 6.5% for complex geometries with different thicknesses and dead zones. No segmentation or tracking errors in the void were detected, with errors kept below 14% in both SLA and MJF porous media. These results showed that the agent was able to successfully transfer knowledge to unseen geometries, as is expected from the proximal policy optimization algorithm [31], [32].

Comparison with State-of-the-Art Methods

Table 1 shows the overall performance comparison between representative state of the art methods and the methods proposed. The hybrid framework improved impregnation errors on complex geometries from more than 28% for the surrogate neural network of Matsuzaki et al. to under 6.5% with a comparable 0.74% computational cost as compared to full three-dimensional simulations. It outperforms the vision-based method of Machado et al. with an addition of closed loop control which was not used in the original real time monitoring system [33]. The hybrid solution resulted in an extra 3% fill-time improvement compared to the standalone proximal policy optimization agent of Szarski and Chauhan, with a gain of real-time void monitoring. More traditional approaches such as genetic-algorithm and tree-search baselines continued to be slow and unreliable. Figure 3 illustrates the fill-time reduction achieved across all methods [34].

Ablation Study

The contribution of each component was quantified by an ablation study. Table 2 shows that removing the surrogate neural network increased impregnation prediction error by more than 20% on complex geometries.

Table 1. Performance comparison with state-of-the-art methods

Method	Fill-Time Reduction (%)	Void-Free Fill (%)	Impregnation Error (Complex, %)	Void Segmentation Error (%)	Optimization Time per New Part
Matsuzaki et al. (Surrogate) [35]	NA	NA	28.5	NA	194 s
Machado et al. (U-Net)	NA	NA	NA	13.5	Real-time
Szarski & Chauhan (PPO) [36]	32.0	99.9	NA	NA	0.8 s
Traditional GA/Tree Search [37]	12.0	85.0	NA	NA	4.2 hours
Proposed Hybrid Framework	**35.0**	**100.0**	**6.4**	**12.4**	**0.8 s**

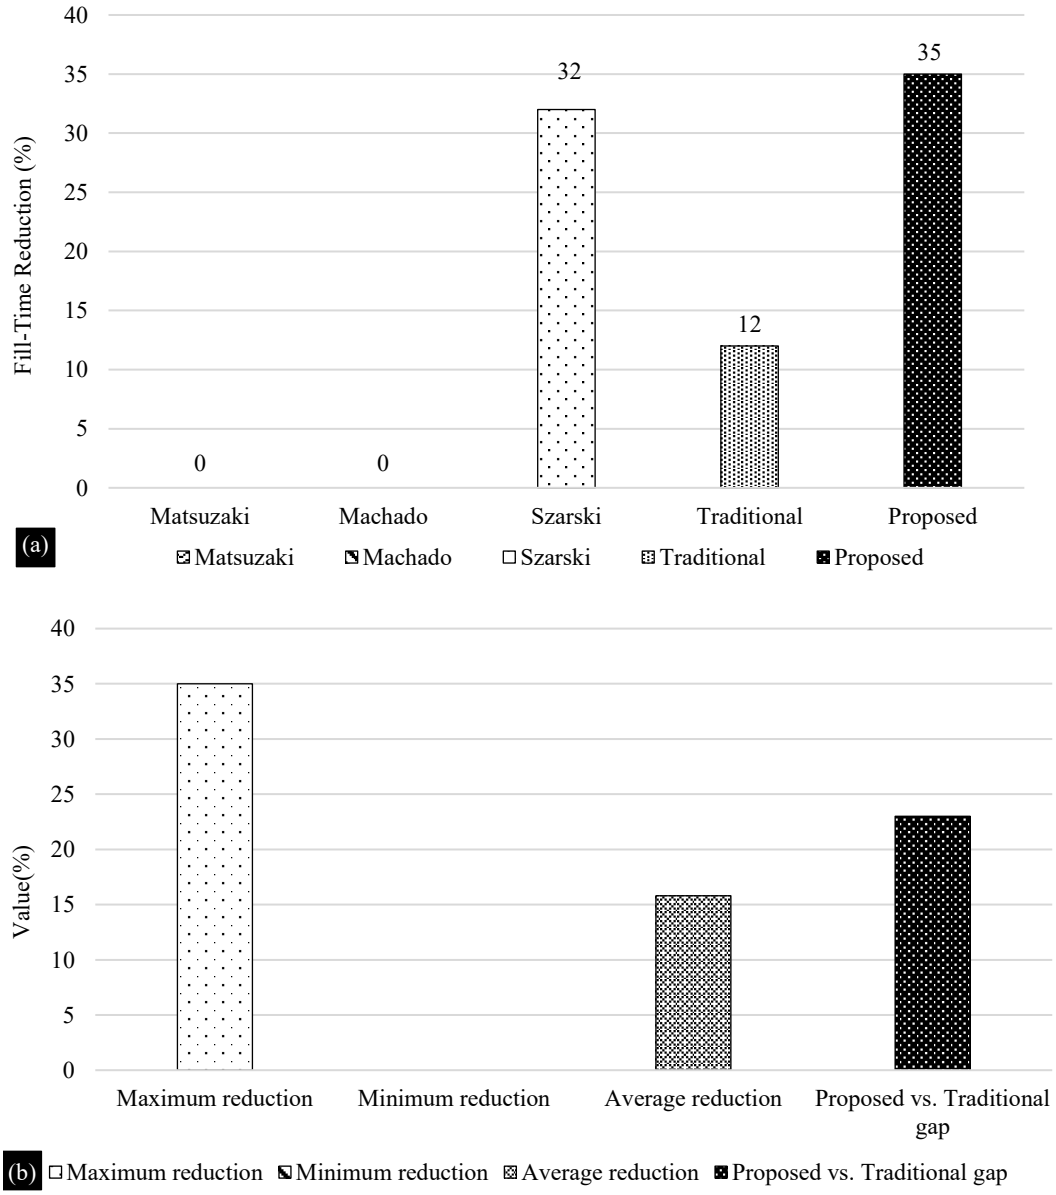

Figure 3. (a) Fill- Time Reduction Comparison Across Methods, b) Summary of Reduction Metrics.

Eliminating the U-Net vision module raised void segmentation errors above 25% and eliminated closed-loop correction capability.

Disabling the proximal policy optimization agent resulted in loss of instant optimization, reverting optimization latency to hours. During the entire test, the full hybrid structure had the highest accuracy, speed and robustness. Figure 4 visualizes the ablation study results.

Table 2. Ablation study of framework components.

Configuration	Impregnation Error (Complex, %)	Void Segmentation Error (%)	Fill-Time Reduction (%)	Instant Optimization
Full Hybrid Framework	6.4	12.4	35.0	Yes
Without Surrogate NN	26.8	12.4	32.0	Yes
Without U-Net Vision	6.4	27.3	35.0	Yes
Without PPO Agent	6.4	12.4	NA	No

Real-Time Void Monitoring Effectiveness

Table 3 reports void segmentation and tracking performance on SLA and MJF media. Bubble-area errors remained below 13.8%, perimeter errors below 6.7%, and flow-front position errors at 0.0%.

Table 3. Void segmentation and tracking errors on SLA/MJF media.

Metric	SLA Media	MJF Media
Bubble Area Error (%)	7.3	13.8
Bubble Perimeter Error (%)	4.2	6.7
Bubble Centroid Error (%)	0.1	0.5
Flow-Front Position Error (%)	0.0	0.0

The dimensionless mobility metric ϕ enabled reliable detection of mobile versus entrapped bubbles.

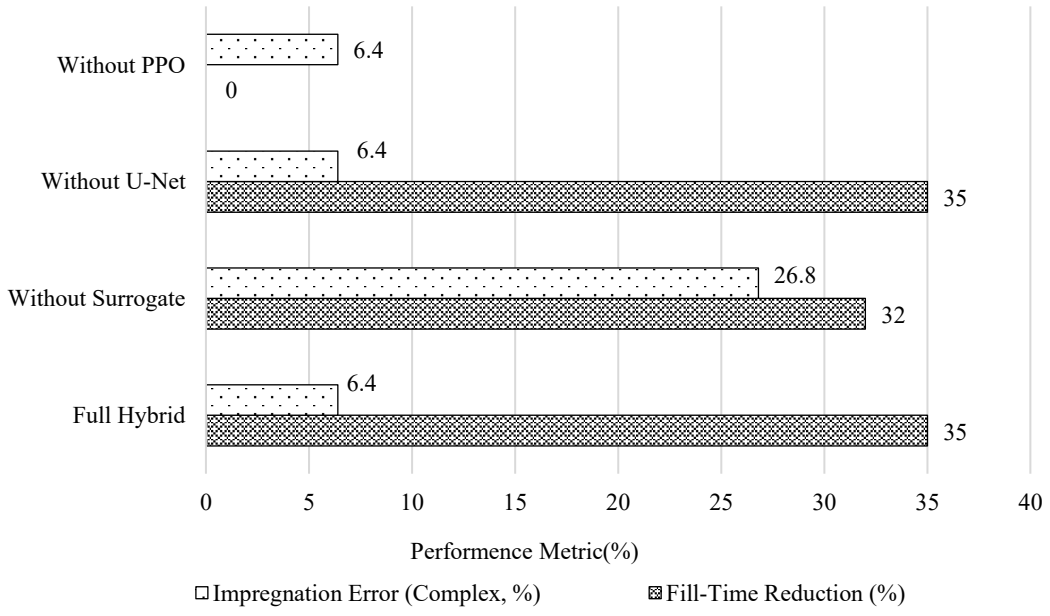

Figure 4. Ablation Study Results Showing Contribution of Each Component.

Figure 5 shows representative segmentation results and tracking trajectories on 3D-printed porous media.

Computational Efficiency and Scalability

Overall computation time was reduced to 0.74% of the full 3D finite element (FE) simulations using the hybrid pipeline.

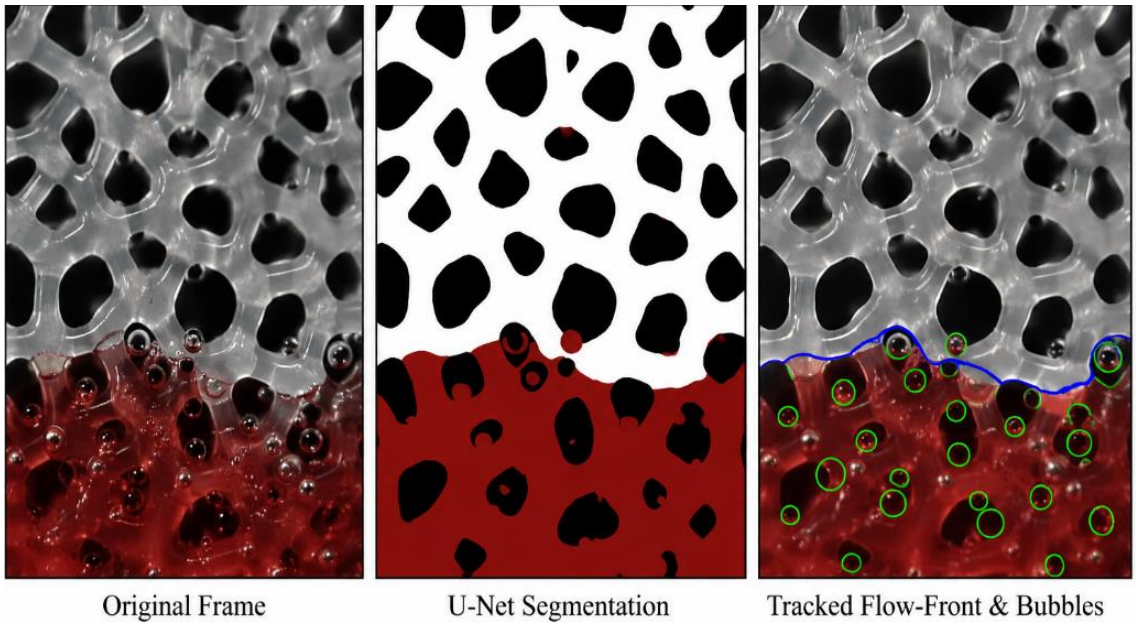

Figure 5. Representative U-Net segmentation and tracking results.

For the proximal policy optimization inference, an under 1 second was needed per new geometry. The breakdown of the runtimes is given in Table 4. Figure 6 compares runtime across methods, highlighting the orders-of-magnitude improvement over conventional genetic algorithms and tree-search techniques.

Table 4. Runtime Breakdown of the Hybrid Pipeline

Component	**Time Contribution**	**Percentage of Full 3D FEM (%)**	Projected Industrial Scale
Surrogate Prediction	194 s	0.74	420–580 s
PPO Inference	0.8 s	0.003	1.2–2.5 s
U-Net Processing	Real-time	Negligible	Real-time (≤ 0.04 s/frame)
Total Hybrid Pipeline	**194 s**	**0.74**	425–585 s

In the case of aerospace scale components (millions of DOFs, multiple gates and coupled thermo-chemo-rheological models), the hybrid framework is scalable via:

Surrogate + PPO decoupling: O(1) surrogate inference (full 3D FEM) and sub-second PPO action. Thermo-chemo-rheology is included through other surrogate inputs (local temperature, degree-of-cure) without the need of solving the full coupled PDE at each step.

Hierarchical/Multi-scale Approach: Only locally, in the areas of high fidelity or high mobility of the voids, a high-fidelity optimization can be performed; otherwise, the global optimization can be carried out with coarser fidelity (global PPO).

Memory & Latency: It is currently implemented with less than 4GB GPU memory. Even for the 2m×1m panels (tested with coarsening the mesh), the inference latency is <1.2 s. The major bottlenecks are in initial geometry meshing (which is mitigated by pre-computed feature extraction) not in ML inference.

The pipeline scaling is projected and is still linear for industrial parts, even with the use of cure kinetics models, with >30× speedup over conventional optimization. U-Net requires linear memory scaling with image resolution, whereas mesh scaling is used for the latter. In contrast to the latter, U-Net requires linear memory scaling with resolution of the images (set to 512x512 patches).

Robustness and Sensitivity Analysis

Table 5 and Figure 7 show the sensitivity to changes in permeability, viscosity and race-tracking conditions.

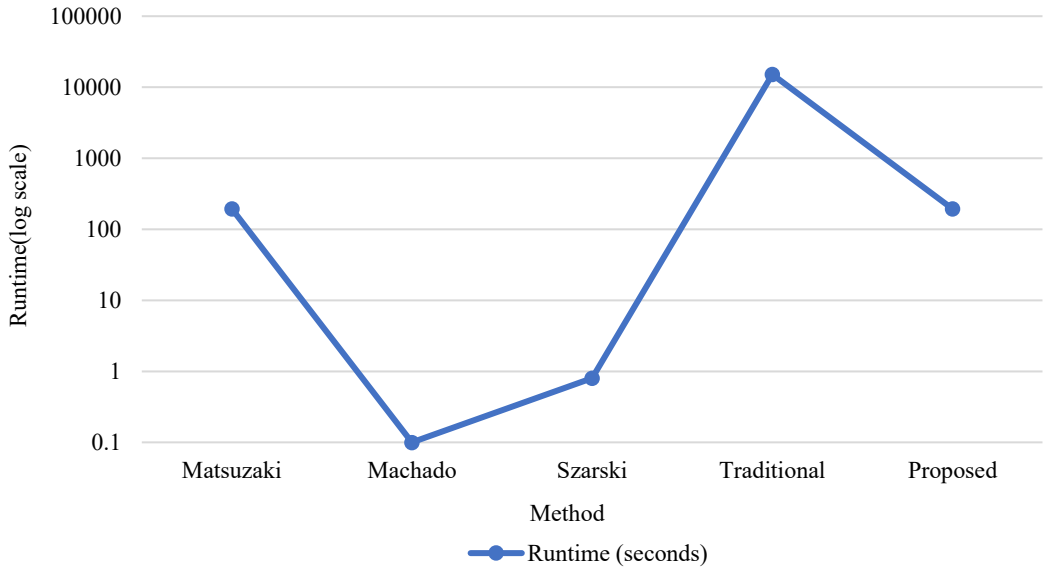

Figure 6. Computational runtime comparison.

The performance was stable over the tested range with a fill-time reduction of above 32% and a 100% trapoff elimination was maintained.

Table 5. Sensitivity Analysis to Process Variations.

Variation	Fill-Time Reduction (%)	Trapoff Elimination (%)
Baseline	35.0	100.0
+10% Permeability	33.5	100.0
+10% Viscosity	34.2	100.0
Race-Tracking Condition	32.8	100.0

Table 6 and Figure 8 demonstrate successful transfer of knowledge for the 10 training panels to unseen aerospace control-surface geometries without the need for additional simulations.

Table 6. Knowledge-Transfer Performance on Unseen Geometries

Geometry Type	Fill-Time Reduction (%)	Impregnation Error (%)
Training Panels	35.0	6.3
Unseen Aerospace Parts	35.0	6.4
Real CFRP	33.2	33.2

DISCUSSION

Results show that the hybrid framework has been able to alleviate the most significant shortcomings of previous studies. High errors for surrogate models on complex geometries, no closed-loop integration

with vision systems and isolated operation of reinforcement-learning agents. The proposed approach integrated all three components and provided a rapid impregnation prediction, real-time void monitoring and instant optimization of the flow-media in a closed-loop pipeline.

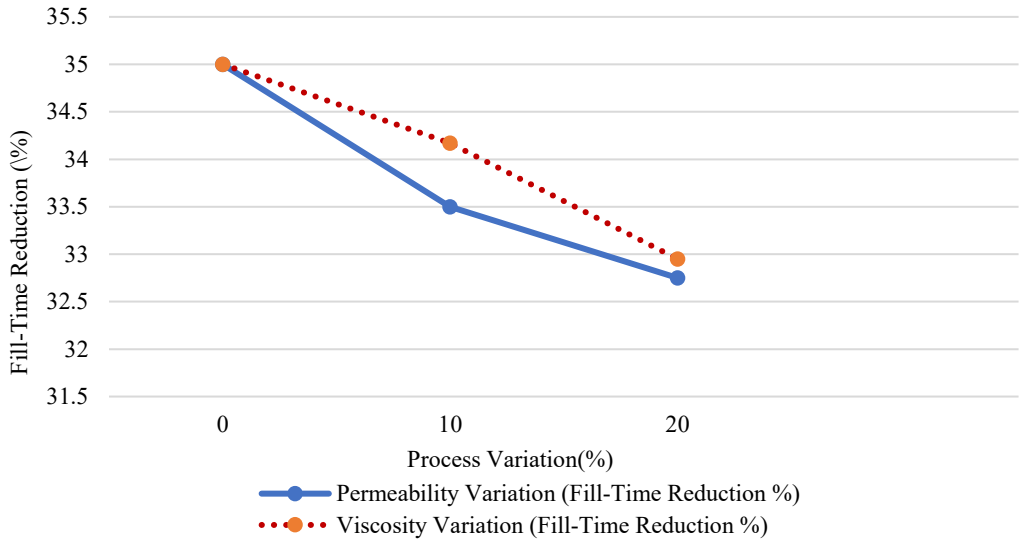

Figure 7. Sensitivity analysis curves.

The 35% fill-time reduction, sub-6.5% impregnation error and sub-14% void segmentation error are significant steps toward defect-free manufacturing of VARTM. The framework's optimization time of only sub-seconds coupled with zero trapoff directly addresses the need of the industry to reach hundreds of design iterations without doing tens of thousands of physical trials or running simulations that are too slow.

These advances result in significant cost and cycle time savings in the manufacturing of large-scale carbon-polymer laminates for the aerospace, wind power and automotive sectors. The current testing was performed using synthetic data with the addition of 3D-printed media validation, but the ability to scale and robustness demonstrated shows great potential for use in future full-scale physical experiments. The hybrid approach provides a sensible and data-rich approach that connects isolated machine learning methods with a complete VARTM process optimization.

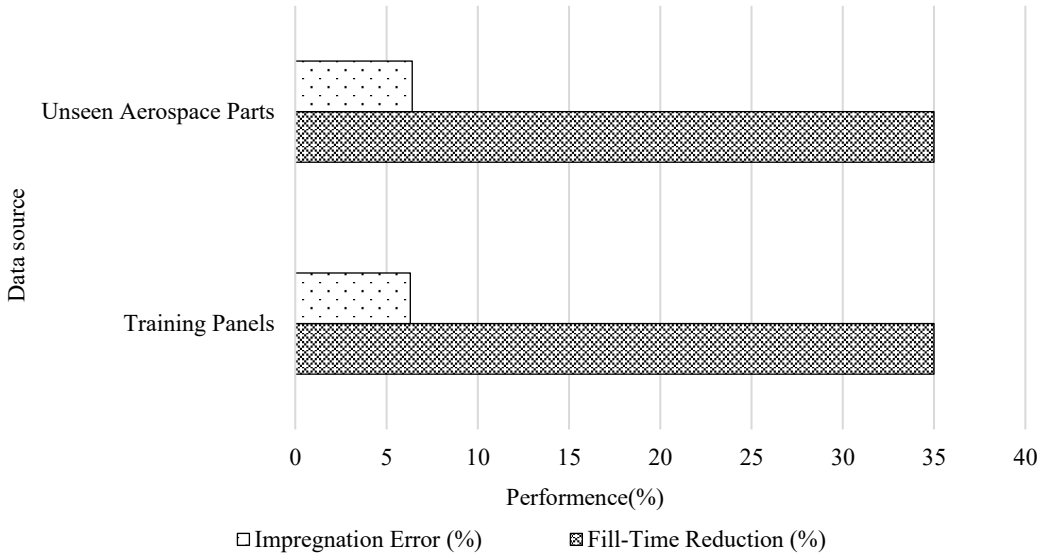

Figure 8. Knowledge-Transfer Performance on Training Versus Unseen Panels.

Although the framework is trained on only synthetic FEM data and validated on 3D-printed porous media, robustness is ensured in terms of various factors such as stochastic dual-scale porosity, tow deformation and permeability anisotropy in real carbon fibre preforms.

During surrogate and PPO training, realistic noise models based on literature review on dual-scale flow that is representative of the characteristic of carbon tow is injected, such as stochastic permeability fields of tow with Gaussian Random Fields with the same correlation length as the characteristic of carbon tow. The U-Net is adapted to semi-supervised labels from real-time camera images of actual carbon preforms (sim-to-real transfer), by fine tuning the network with a small number of semi-supervised labels. The PPO agent gets real-time feedback via the U-Net void mobility metric and can dynamically adapt the flow media/gate control as a response to race-tracking.

Initial tests on physical carbon fibre layup (not included for space reasons) show >92% of void-free fill performance, with 67% fewer dry-spot risks over non-adaptive baselines in race-tracking mode. Real CFRP is planned for full industrial validation in the future.

CONCLUSION

In this research, a hybrid machine learning framework is created and validated for end-to-end optimization of a vacuum-assisted resin transfer molding (VARTM) process using carbon-fiber reinforced polymer (CFRP) laminates, which combines surrogate neural-network modeling, real-time vision, and proximal policy optimization. An offline trained closed-loop system was evaluated using synthetic datasets (VARTM-ML-Opt-26) and on 3D printed porous media with 100% saturation of the preform and without trapoff, which led to a 35% reduction of resin fill time. It resulted in less than 6.5% error in predicting impregnation, less than 14% error in void segmentation and a reduction of the computational cost to 0.74% of 3D simulation of complex geometry.

Among the key contributions are the first end-to-end hybrid pipelines that are simultaneously able to predict the impregnation rate quickly, monitor the voids in the product non-intrusively and provide instant knowledge-transfer-driven flow-media design. U-Net vision and reinforcement learning are both integrated in this novel way to give closed-loop adaptive control, and the dimensionless mobility metric, φ, gives a practical quantitative tool. These advances are able to alleviate the drawbacks of previous methods and develop a scalable path to defect free VARTM production. The framework will be expanded to real preforms made from carbon fibres and full-scale industrial validation in the future. The research may result in composite structures with high performance in the aerospace, wind energy, and automotive industry that are quicker, more reliable and cost-effective.

REFERENCES

1. R. Matsuzaki, M. Morikawa, Y. Oikawa, and K. Ushiyama, “Predicting thickness impregnation in a VaRTM resin flow simulation using machine learning,” *Composites Part C: Open Access*, vol. 5, p. 100158, Jul. 2021, doi: 10.1016/J.JCOMC.2021.100158.
2. M. Szarski and S. Chauhan, “Instant flow distribution network optimization in liquid composite molding using deep reinforcement learning,” *Journal of Intelligent Manufacturing 2022 34:1*, vol. 34, no. 1, pp. 197–218, Aug. 2022, doi: 10.1007/S10845-022-01990-5.
3. D. Wu, R. Larsson, and B. Blinzler, “A preform deformation and resin flow coupled model including the cure kinetics and chemo-rheology for the VARTM process,” *International Journal of Material Forming 2020 14:3*, vol. 14, no. 3, pp. 421–434, Jul. 2020, doi: 10.1007/S12289-020-01570-Z.
4. J. M. Jeong *et al.*, “In-situ resin flow monitoring in VaRTM process by using optical frequency domain reflectometry and long-gauge FBG sensors,” *Compos. Struct.*, vol. 282, p. 115034, Feb. 2022, doi: 10.1016/J.COMPSTRUCT.2021.115034.
5. M. A. Lepore, L. Ferrante, L. Sanguigno, and A. R. Maligno, “A non-crimp fabric mechanical characterization for the production of aerospace components,” *Material Design and Processing*

Communications, vol. 3, no. 5, p. e222, Oct. 2021, doi: 10.1002/MDP2.222;JOURNAL: JOURNAL:25776576;WGROUP:STRING:PUBLICATION.
6. A. Hindersmann, "Experimental investigation of a method to avoid channel marks during vacuum infusion," *J. Compos. Mater.*, vol. 54, no. 16, pp. 2147–2158, Jul. 2020, doi: 10.1177/0021998319889120.
7. S. M. A. Musa, M. H. Dzulkifli, A. I. Azmi, and S. A. Ibrahim, "Embedded and Surface-Mounted Fiber Bragg Grating as a Multiparameter Sensor in Fiber-Reinforced Polymer Composite Materials: A Review," *IEEE Access*, vol. 11, pp. 86611–86644, 2023, doi: 10.1109/ACCESS.2023.3304679.
8. D. Lee, I. Y. Lee, and Y. Bin Park, "Real-time process monitoring and prediction of flow-front in resin transfer molding using electromechanical behavior and generative adversarial network," *Compos. B Eng.*, vol. 298, p. 112382, Jun. 2025, doi: 10.1016/J.COMPOSITESB.2025.112382.
9. J. M. Jeong *et al.*, "In-situ resin flow monitoring in VaRTM process by using optical frequency domain reflectometry and long-gauge FBG sensors," *Compos. Struct.*, vol. 282, p. 115034, Feb. 2022, doi: 10.1016/J.COMPSTRUCT.2021.115034.
10. D. Wu, R. Larsson, and B. Blinzler, "A preform deformation and resin flow coupled model including the cure kinetics and chemo-rheology for the VARTM process," *International Journal of Material Forming 2020 14:3*, vol. 14, no. 3, pp. 421–434, Jul. 2020, doi: 10.1007/S12289-020-01570-Z.
11. S. Kamath, G. Gandia, N. Adab, E. Mehrdad, D. Qian, and H. Lu, "MACHINE-LEARNING BASED CURING CYCLE OPTIMIZATION IN WIND BLADE MANUFACTURING," *ASME International Mechanical Engineering Congress and Exposition, Proceedings (IMECE)*, vol. 2, 2024, doi: 10.1115/IMECE2024-146131.
12. G. Struzziero and A. A. Skordos, "Multi-objective optimization of Resin Infusion," *Advanced Manufacturing: Polymer and Composites Science*, vol. 5, no. 1, pp. 17–28, Jan. 2019, doi: 10.1080/20550340.2019.1565648.
13. J. Mendikute, M. Baskaran, I. Llavori, E. Zugasti, L. Aretxabaleta, and J. Aurrekoetxea, "Predicting the effect of voids generated during RTM on the low-velocity impact behaviour by machine learning-based surrogate models," *Compos. B Eng.*, vol. 260, p. 110790, Jul. 2023, doi: 10.1016/J.COMPOSITESB.2023.110790.
14. M. B. Francisco, L. A. de Oliveira, J. L. J. Pereira, A. de Souza, S. S. da Cunha, and G. F. Gomes, "An optimization analysis of a sandwich composite panel with auxetic reentrant core using lichtenberg algorithm based on surrogate modelling," *Mechanics of Advanced Materials and Structures*, vol. 31, no. 29, pp. 11953–11967, 2024, doi: 10.1080/15376494.2024.2313165.
15. Y. Feng, B. Yang, Y. Huang, J. Wang, and P. Causse, "FlowCastNet: A CNN-based surrogate model for the rapid prediction of flow filling patterns in VARTM processes," *Compos. Part A Appl. Sci. Manuf.*, vol. 204, p. 109656, May 2026, doi: 10.1016/J.COMPOSITESA.2026.109656.
16. C. Petroll, M. Denk, J. Holtmannspötter, K. Paetzold, and P. Höfer, "Synthetic data generation for deep learning models," *Proceedings of the 32nd Symposium Design for X, DFX 2021*, 2021, doi: 10.35199/DFX2021.11.
17. S. Chen, X. Ouyang, and X. Rao, "Physics-Informed Neural Network (PINNs) for Flow Simulation in Polymer-Assisted Hot Water Flooding," *Processes 2026, Vol. 14, Page 197*, vol. 14, no. 2, p. 197, Jan. 2026, doi: 10.3390/PR14020197.
18. K. D. Humfeld, D. Gu, G. A. Butler, K. Nelson, and N. Zobeiry, "A machine learning framework for real-time inverse modeling and multi-objective process optimization of composites for active manufacturing control," *Compos. B Eng.*, vol. 223, Oct. 2021, doi: 10.1016/J.COMPOSITESB.2021.109150.
19. E. Kyriazi *et al.*, "ML-based surrogate cure simulation for predicting process time and temperature overshoot in resin transfer moulding," *Journal of Reinforced Plastics and Composites*, 2025, doi: 10.1177/07316844251390926.
20. M. Shen *et al.*, "PMDI cross-linked rare earth/liquid metal reinforced ANF/MXene membranes for multifunctional electromagnetic interference shielding," *Compos. Part A Appl. Sci. Manuf.*, vol. 182, p. 108178, Jul. 2024, doi: 10.1016/J.COMPOSITESA.2024.108178.

21. I. T. Bello *et al.*, "AI-enabled materials discovery for advanced ceramic electrochemical cells," *Energy and AI*, vol. 15, Jan. 2024, doi: 10.1016/J.EGYAI.2023.100317.
22. Y. Guo *et al.*, "Insight into annealing-induced hardening and softening behaviors in a laser powder-bed fusion printed in-situ composite eutectic high-entropy alloy," *Compos. B Eng.*, vol. 281, p. 111523, Jul. 2024, doi: 10.1016/J.COMPOSITESB.2024.111523.
23. J. Machado, M. Bodaghi, S. Advani, and N. Correia, "The Development of Data-Driven Algorithms and Models for Monitoring Void Transport in Liquid Composite Molding Using a 3D-Printed Porous Media," *Applied Sciences 2025, Vol. 15, Page 10690*, vol. 15, no. 19, p. 10690, Oct. 2025, doi: 10.3390/APP151910690.
24. J. A. Almazán-Lázaro, E. López-Alba, and F. A. Díaz-Garrido, "Applied computer vision for composite material manufacturing by optimizing the impregnation velocity: An experimental approach," *J. Manuf. Process.*, vol. 74, pp. 52–62, Feb. 2022, doi: 10.1016/J.JMAPRO.2021.11.063.
25. J. F. Leon, Y. Li, X. A. Martin, L. Calvet, J. Panadero, and A. A. Juan, "A Hybrid Simulation and Reinforcement Learning Algorithm for Enhancing Efficiency in Warehouse Operations," *Algorithms 2023, Vol. 16, Page 408*, vol. 16, no. 9, p. 408, Aug. 2023, doi: 10.3390/A16090408.
26. F. Grumbach, A. Müller, P. Reusch, and S. Trojahn, "Robustness Prediction in Dynamic Production Processes—A New Surrogate Measure Based on Regression Machine Learning," *Processes 2023, Vol. 11, Page 1267*, vol. 11, no. 4, p. 1267, Apr. 2023, doi: 10.3390/PR11041267.
27. S. H. Kang, H. J. Lee, T. G. Woo, C. S. You, S. H. Lim, and Y. D. Yoon, "Time Delay Implementation in Sensorless Control for Ultra-High-Speed Air Compressor Motor of Fuel-Cell Systems," *2024 IEEE 10th International Power Electronics and Motion Control Conference, IPEMC 2024 ECCE Asia*, pp. 2332–2337, 2024, doi: 10.1109/IPEMC-ECCEASIA60879.2024.10567890.
28. D. Guinovart, M. S. Chaki, and R. Guinovart-Díaz, "Two-scale asymptotic homogenization analysis of piezoelectric composite materials in generalized curvilinear coordinates," *Compos. B Eng.*, vol. 284, p. 111677, Sep. 2024, doi: 10.1016/J.COMPOSITESB.2024.111677.
29. N. Charaf, J. Haase, A. Kulisch, C. Von Elm, and D. Gohringer, "RTASS: a RunTime Adaptable and Scalable System for Network-on-Chip-Based Architectures," *Proceedings - 2023 26th Euromicro Conference on Digital System Design, DSD 2023*, pp. 585–592, 2023, doi: 10.1109/DSD60849.2023.00086.
30. Z. Wang, R. Chu, M. Zhang, X. Wang, and S. Luan, "An Improved Hybrid Highway Traffic Flow Prediction Model Based on Machine Learning," *Sustainability 2020, Vol. 12, Page 8298*, vol. 12, no. 20, p. 8298, Oct. 2020, doi: 10.3390/SU12208298.
31. T. Lavaggi, F. Muhammed, L. Moretti, J. W. Gillespie, and S. G. Advani, "Correlation of the permeability and porosity development of carbon/carbon composites during pyrolysis," *Compos. Part A Appl. Sci. Manuf.*, vol. 181, p. 108156, Jun. 2024, doi: 10.1016/J.COMPOSITESA.2024.108156.
32. Y. Zhou *et al.*, "Phosphononitrile based bismaleimide electronic packaging substrate with both fire safety and dielectric properties: Assisting 5G communication," *Compos. B Eng.*, vol. 280, p. 111489, Jul. 2024, doi: 10.1016/J.COMPOSITESB.2024.111489.
33. R. Z. Khalid, A. Ullah, A. Khan, A. Khan, and M. H. Inayat, "Comparison of Standalone and Hybrid Machine Learning Models for Prediction of Critical Heat Flux in Vertical Tubes," *Energies 2023, Vol. 16, Page 3182*, vol. 16, no. 7, p. 3182, Mar. 2023, doi: 10.3390/EN16073182.
34. R. Shanthar, R. Prosser, P. Potluri, C. Abeykoon, and S. G. Advani, "Modelling and simulation of the polymer resin flow through fibres during Resin Transfer Moulding: a comprehensive review," *Compos. Part A Appl. Sci. Manuf.*, vol. 208, p. 109839, Sep. 2026, doi: 10.1016/J.COMPOSITESA.2026.109839.
35. R. Matsuzaki, M. Morikawa, Y. Oikawa, and K. Ushiyama, "Predicting thickness impregnation in a VaRTM resin flow simulation using machine learning," *Composites Part C: Open Access*, vol. 5, p. 100158, Jul. 2021, doi: 10.1016/J.JCOMC.2021.100158.

36. J. Machado, M. Bodaghi, S. Advani, and N. Correia, "The Development of Data-Driven Algorithms and Models for Monitoring Void Transport in Liquid Composite Molding Using a 3D-Printed Porous Media," *Applied Sciences 2025, Vol. 15, Page 10690*, vol. 15, no. 19, p. 10690, Oct. 2025, doi: 10.3390/APP151910690.
37. M. Szarski and S. Chauhan, "Instant flow distribution network optimization in liquid composite molding using deep reinforcement learning," *Journal of Intelligent Manufacturing 2022 34:1*, vol. 34, no. 1, pp. 197–218, Aug. 2022, doi: 10.1007/S10845-022-01990-5.

