Genetic algorithm–optimized loss balancing in physics-informed neural networks for manufacturing digital twin applications

Journal of Intelligent Manufacturing and Special Equipment

Aswin Karkadakattil
Department of Mechanical, Independent Researcher, Kanhangad, India

Received 18 April 2026
Revised 23 May 2026
Accepted 22 June 2026

Abstract

Purpose – This study addresses a key limitation in physics-informed neural networks (PINNs), namely the reliance on manually selected or heuristically tuned loss weights governing the balance between data fidelity, physics residuals and boundary constraints. Improper weighting often leads to instability, poor reproducibility and sensitivity to user expertise, particularly in manufacturing-oriented thermal modelling and digital twin applications.

Design/methodology/approach – A genetic algorithm (GA)-based meta-optimization framework is proposed to automate the selection of normalized loss weights in PINNs. The GA operates as an offline optimization layer, while the inner PINN enforces the governing partial differential equations. A composite fitness function evaluates candidate weight configurations using physics consistency, data agreement and boundary-condition satisfaction. The framework is validated using a two-dimensional transient heat-conduction problem with a moving Gaussian heat source representative of laser-based manufacturing processes under noisy data conditions. Comparative analysis is performed against fixed-weight and adaptive-weight PINN strategies.

Findings – The proposed GA-optimized PINN achieves an approximately 39% reduction in governing-equation residual root-mean-square error (RMSE) compared with fixed-weight training while maintaining comparable prediction accuracy and stable convergence. The framework produces lower PDE residual distributions and improved robustness under noisy conditions. Compared with adaptive weighting approaches, the GA-based framework demonstrates improved global loss balancing and stronger physics consistency without degrading predictive performance.

Originality/value – This work introduces a global meta-optimization strategy for PINN loss balancing using a genetic algorithm, enabling systematic and interpretable selection of normalized loss weights for reliable physics-informed surrogate modelling in manufacturing digital twin applications

Funding: The author received no specific funding for this work.

Competing interests: The author declares that there are no competing financial or non-financial interests associated with this work.

Journal of Intelligent Manufacturing and Special Equipment
Emerald Publishing Limited
e-ISSN: 2633-660X
p-ISSN: 2633-6596
DOI 10.1108/JIMSE-04-2026-0016

Graphical abstract

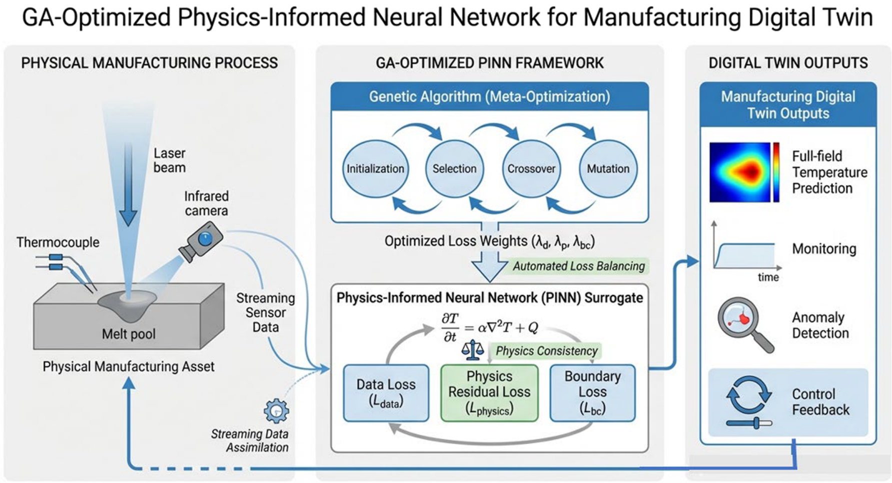

Keywords Physics-informed neural networks, Genetic algorithm, Loss balancing, Manufacturing processes, Digital twin, Heat conduction
Paper type Research article

1. Introduction

1.1 Background

Physics-informed neural networks (PINNs) have emerged as a powerful paradigm for modelling complex physical systems by embedding governing equations directly into neural network training. By constraining learning with known physical laws, PINNs provide a principled alternative to purely data-driven approaches, particularly in scenarios where experimental data are sparse, noisy, or costly to acquire. These characteristics make PINNs especially attractive for manufacturing applications, where processes are governed by well-established physical principles, yet comprehensive high-fidelity datasets are often limited. In recent years, PINNs have been applied across a range of manufacturing domains, including machining, additive manufacturing (AM), and thermal processing. Representative studies include heat transfer modelling in machining and AM (Zobeiry and Humfeld, 2021; Zhu *et al.*, 2021; Tod *et al.*, 2021; Zhao *et al.*, 2023; Chen *et al.*, 2024), melt pool dynamics and process parameter prediction in metal AM (Zhu *et al.*, 2021; Zhao *et al.*, 2023; Farrag *et al.*, 2025; Scheel and Hosseini, 2025; Oddiraju *et al.*, 2025), and surrogate modelling for composite and industrial manufacturing systems (Würth *et al.*, 2023; Chen *et al.*, 2023; Cooper *et al.*, 2023; Ciampi *et al.*, 2025). In many of these applications, transient heat conduction equations describe thermal evolution, while conservation laws and constitutive relations govern deformation, energy transport, and tool–workpiece interactions (Zobeiry and Humfeld, 2021; Zhu *et al.*, 2021; Wang *et al.*, 2020; Li *et al.*, 2023a, b). By explicitly incorporating governing physics into the learning process, PINNs enhance physical consistency and reduce dependence on large datasets. Recent developments have increasingly connected physics-informed learning with digital twin technologies for industrial monitoring, process optimization, and intelligent manufacturing applications. Physics-informed digital twin frameworks have been explored for process parameter optimization, fault diagnosis, surrogate modelling, and predictive analysis in manufacturing and industrial systems (Garcia *et al.*, 2025; Seppi *et al.*, 2025; Lu *et al.*, 2025). These developments highlight the growing importance of physics-constrained surrogate models capable of reconstructing full-field system behaviour from sparse measurements while maintaining physical consistency and interpretability. Within such contexts, PINNs can serve as robust surrogate models for thermal field reconstruction, process monitoring, anomaly detection, and predictive decision support. However, the reliability of such models depends strongly on stable and well-balanced training formulations, particularly when multiple competing loss components are simultaneously enforced.

1.2 Motivation

Despite their conceptual advantages, the practical performance of PINNs is highly sensitive to the formulation of the composite training loss function. A typical PINN loss comprises multiple competing components, including data mismatch loss, physics residual loss, and boundary or initial condition losses. The relative influence of these components is controlled by loss weighting factors, which strongly affect convergence behaviour, solution stability, and predictive accuracy. In many existing manufacturing-focused PINN studies, these loss weights are selected using fixed values or heuristic scaling rules derived through trial-and-error (Zobeiry and Humfeld, 2021; Zhu *et al.*, 2021; Tod *et al.*, 2021; Zhao *et al.*, 2023; Würth *et al.*, 2023). While such approaches may suffice for simplified benchmark problems, they exhibit significant limitations in realistic industrial environments. Fixed weights lack adaptability under changing operating conditions, and heuristic tuning is highly problem-dependent and difficult to reproduce. As a result, training outcomes can depend substantially on user expertise rather than systematic optimization, reducing robustness and limiting scalability (Cooper *et al.*, 2023; Ciampi *et al.*, 2025). Several adaptive and gradient-based loss balancing strategies have been proposed to mitigate manual tuning (Hua *et al.*, 2023; Li and Feng, 2022; Liu *et al.*, 2022; Gao *et al.*, 2025; Li *et al.*, 2025). Although these methods partially automate weight adjustment, they typically introduce additional hyperparameters, may become unstable in stiff or multi-scale problems, and often provide limited interpretability. In many cases, adaptive weighting mechanisms rely primarily on instantaneous gradient information or local training dynamics, which may limit their effectiveness in highly non-convex optimization landscapes characteristic of manufacturing-oriented PINNs. Moreover, dynamically varying loss weights can occasionally introduce oscillatory convergence behaviour or sensitivity to noisy measurements under complex operating conditions (Farea *et al.*, 2025; Zhou *et al.*, 2025; Barreau and Shen, 2025). Consequently, loss balancing remains a critical bottleneck restricting reproducibility, robustness, and reliable deployment of PINNs in practical engineering applications.

1.3 Research gap

While substantial progress has been made in applying PINNs to manufacturing process modelling, systematic strategies for globally optimizing loss weights remain underexplored. Most existing approaches rely on fixed heuristics, local gradient-based adaptation, or problem-specific tuning that do not explicitly address the global and non-convex nature of the loss balancing problem (Hua *et al.*, 2023; Li and Feng, 2022; Liu *et al.*, 2022; Gao *et al.*, 2025; Li *et al.*, 2025; Farea *et al.*, 2025; Zhou *et al.*, 2025; Barreau and Shen, 2025). As a result, optimized weight configurations may lack reproducibility and may not generalize across varying operating regimes or data conditions. Furthermore, existing approaches often optimize loss weights based on local training behaviour or internal objective scaling, which may introduce bias toward specific loss components and limit overall physical consistency. A structured framework that balances data fidelity, physics enforcement, and boundary constraints in a globally consistent manner is therefore required. Addressing this gap requires an optimization strategy capable of globally exploring the normalized loss-weight space, operating independently of gradient-based heuristics, and aligning weight selection with physically meaningful performance metrics. Such a framework is essential for improving the robustness and reliability of PINN-based surrogate models in manufacturing and related engineering applications, particularly in digital twin environments where predictive consistency and stability are critical.

1.4 Contributions

To address these limitations, this study proposes a genetic algorithm (GA)–assisted meta-optimization framework for automated loss balancing in physics-informed neural networks. The principal contributions are as follows:

(1) Global meta-optimization of loss weights: A GA-based automated loss balancing strategy is introduced to systematically optimize normalized PINN loss weights without manual or heuristic tuning. By operating at an offline meta-optimization layer, the GA enables global exploration of the weight space and is well suited to the non-convex structure of PINN training.

(2) Composite fitness-based optimization criterion: Candidate weight configurations are evaluated using a composite fitness function that accounts for physics residuals alongside data and boundary constraints. This formulation reduces bias toward any single loss component and ensures balanced and physically consistent optimization.

(3) Manufacturing-relevant validation under noisy data conditions: The framework is demonstrated using a two-dimensional transient heat conduction problem with a moving heat source representative of laser-based manufacturing processes. A noisy, streaming-inspired data scenario is incorporated to assess robustness under realistic measurement conditions.

(4) Quantitative robustness and comparative assessment: The proposed approach is systematically compared against fixed-weight and gradient-based adaptive PINNs. Residual root-mean-square error reduction, predictive accuracy, convergence behaviour, and variability across multiple random initializations are evaluated to assess robustness and reproducibility.

Through these contributions, this work advances automated, interpretable, and physics-consistent training strategies for PINNs, providing a systematic foundation for improving their reliability in manufacturing and related engineering modelling applications.

2. Methodology

2.1 Physics-informed neural network formulation

Physics-informed neural networks (PINNs) integrate governing physical laws directly into the training process by embedding differential equation constraints into the loss function. Unlike purely data-driven models, PINNs enforce physical consistency by penalizing violations of governing equations alongside mismatches with observed data and boundary conditions. This combined learning approach allows the model to remain grounded in first-principles physics while still leveraging available measurements, making PINNs particularly suitable for engineering and manufacturing applications.

In a standard formulation, the total training loss is expressed as a weighted combination of multiple components:

$$L_{total} = \lambda_d L_{data} + \lambda_p L_{physics} + \lambda_{bc} L_{boundary}$$

where:

(1) L_{data} measures the discrepancy between model predictions and available data,

(2) $L_{physics}$ penalizes violations of the governing partial differential equation (PDE),

(3) $L_{boundary}$ enforces boundary or initial conditions, and

(4) λ_d, λ_p, and λ_{bc} are non-negative weighting factors controlling the relative contribution of each term.

The data loss is typically defined using mean squared error, while the physics loss is computed by evaluating the PDE residual at collocation points within the domain. Boundary conditions ensure that the predicted solution remains physically admissible. A key challenge in PINN training is the selection of appropriate loss weights. If these weights are not properly balanced, one component may dominate the training process, leading to poor convergence or loss of physical consistency. This issue becomes more pronounced in manufacturing scenarios, where data may be noisy or sparse while strict adherence to governing physics remains essential. These challenges motivate the need for a systematic and automated approach to loss balancing.

2.2 Governing equation and manufacturing proxy

To evaluate the proposed framework in a manufacturing-relevant setting, a two-dimensional transient heat conduction equation with a moving heat source is considered:

$$\frac{\partial T(x,y,t)}{\partial t} = \alpha\left(\frac{\partial^2 T}{\partial x^2} + \frac{\partial^2 T}{\partial y^2}\right) + Q(x,y,t)$$

where $T(x,y,t)$ represents the temperature field, α is the thermal diffusivity, and $Q(x,y,t)$ denotes a localized, time-dependent heat source.

The heat source is modelled as a Gaussian distribution moving across the spatial domain:

$$Q(x,y,t) = \exp\left[-\beta\left((x - x_0(t))^2 + (y - y_0)^2\right)\right]$$

where $x_0(t)$ defines the scanning trajectory. This formulation captures the essential characteristics of laser-based manufacturing processes, where localized energy input generates strong transient thermal gradients.

Although simplified, this model provides a computationally efficient and physically meaningful proxy for thermal behaviour in manufacturing systems. It also offers a controlled environment for evaluating loss balancing strategies and can be extended to more complex multiphysics scenarios involving coupled thermal–mechanical effects.

2.3 GA-optimized loss balancing framework

The central contribution of this study is a genetic algorithm (GA)–based meta-optimization framework for automated loss balancing in PINNs. Instead of manually selecting loss weights, the GA operates as an offline outer optimization layer, identifying suitable weight configurations prior to final model training. The PINN takes spatial–temporal inputs (x,y,t) and predicts the temperature field $T(x,y,t)$. The governing equation and boundary conditions define the individual loss components. The GA searches for optimal values of the loss weights $(\lambda_d, \lambda_p, \lambda_{bc})$ that lead to improved training performance.

Each candidate solution (chromosome) is represented as:

$$c = \left[\lambda_d, \lambda_p, \lambda_{bc}\right]$$

For each candidate, the PINN is trained for a limited number of epochs, and its performance is evaluated using a fitness function. To avoid bias toward any single loss component, the fitness is defined using a composite normalized criterion that reflects overall physical consistency, predictive accuracy, and boundary-condition satisfaction:

$$\text{Fitness} = \omega_1 \widehat{L}_{physics} + \omega_2 \widehat{L}_{data} + \omega_3 \widehat{L}_{boundary}$$

where $\widehat{L}_{physics}$, $\widehat{L}_{data}$, and $\widehat{L}_{boundary}$ denote the normalized physics residual loss, data mismatch loss, and boundary-condition loss, respectively, while ω_1, ω_2, and ω_3 are scaling coefficients used to maintain comparable magnitudes across the different terms. This formulation ensures that improvements in governing-equation consistency do not occur at the expense of predictive fidelity or boundary-condition enforcement.

The GA evolves the population of candidate solutions through:

(1) Initialization of normalized, non-negative weight vectors

(2) Fitness evaluation based on composite model performance

(3) Selection of high-performing candidates

(4) Crossover and mutation to generate new solutions

(5) Renormalization of offspring chromosomes after mutation and crossover operations

This process is repeated over multiple generations, progressively identifying improved weight configurations. A key advantage of this framework is the decoupling between loss-weight optimization and network training. While the PINN parameters are updated using gradient-based methods, the GA explores the loss-weight space independently. This allows broader exploration of candidate configurations, reducing reliance on manual tuning and improving robustness under varying operating conditions. Unlike locally adaptive weighting strategies that adjust weights based on instantaneous training dynamics, the proposed framework performs a global search over the normalized loss-weight space. This enables the identification of balanced configurations that simultaneously improve physics consistency, predictive stability, and convergence robustness. The optimization process is performed offline and does not involve recursive or real-time updating of model states. However, the framework provides a systematic foundation for determining reproducible and physically consistent training configurations that can support manufacturing-oriented digital twin environments and physics-informed surrogate modelling applications. As shown in Figure 1, the proposed framework employs a genetic algorithm as an offline meta-optimization layer to determine the optimal loss weights before the final PINN training stage.

2.4 Genetic algorithm design

2.4.1 Chromosome representation. In the proposed framework, each genetic algorithm (GA) chromosome represents a candidate set of loss-weighting parameters for the PINN:

$$c = \left[\lambda_d, \lambda_p, \lambda_{bc}\right]$$

where each gene corresponds to the weighting factor associated with the data loss, physics residual loss, and boundary-condition loss, respectively. To ensure numerical stability and physical interpretability, all weights are constrained to be non-negative. In addition, normalization is enforced such that:

$$\lambda_d + \lambda_p + \lambda_{bc} = 1$$

This normalization prevents trivial scaling effects, where all weights increase or decrease uniformly without altering their relative influence. Consequently, the optimization focuses exclusively on the balance between competing loss components rather than their absolute magnitudes. The chromosome representation is intentionally maintained as a low-dimensional formulation. By restricting the search space to three physically meaningful parameters, the GA

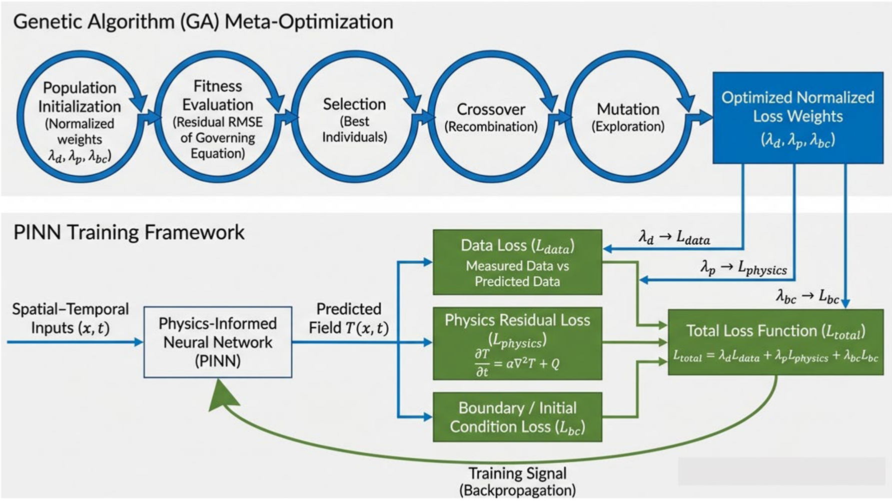

Figure 1. Schematic of the proposed genetic algorithm (GA)–optimized loss balancing framework for physics-informed neural networks (PINNs). The PINN receives spatial–temporal inputs (x, t) and predicts the physical field $T(x, t)$. Training is guided by three loss components: data loss, physics residual loss derived from the governing equations, and boundary/initial condition loss. A genetic algorithm operates at a meta-optimization level to automatically adjust the corresponding loss weights $(\lambda_d, \lambda_p, \lambda_{bc})$ through population initialization, fitness evaluation, selection, crossover, and mutation. The optimized loss weights are fed back into the PINN training loop, enabling systematic and physics-aware loss balancing without manual tuning

remains computationally efficient while directly targeting the most influential factors governing PINN training behaviour. This design enables systematic exploration of the trade-off between data fitting, physics enforcement, and boundary consistency without introducing unnecessary optimization complexity. To preserve feasibility throughout the evolutionary process, offspring chromosomes generated through crossover and mutation are renormalized after each operation. Negative values, if generated during mutation, are clipped prior to normalization. The normalized weights are subsequently updated as:

$$\lambda_i = \frac{\lambda_i}{\sum_j \lambda_j}$$

This procedure ensures consistent weight scaling and maintains valid candidate solutions during the optimization process.

2.4.2 Fitness function. The fitness of each chromosome is evaluated through a short PINN training cycle using the corresponding loss-weight configuration. Instead of training the network to full convergence for every candidate solution, the model is trained for a limited number of epochs during the GA evaluation stage. This truncated training strategy provides a practical balance between computational efficiency and reliable performance estimation.

To ensure balanced evaluation, the fitness function is formulated using a composite normalized metric that simultaneously considers governing-equation consistency, agreement with available data, and boundary-condition satisfaction:

$$\text{Fitness} = \alpha \widehat{L}_{physics} + \beta \widehat{L}_{data} + \gamma \widehat{L}_{boundary}$$

where:

(1) $\widehat{L}_{physics}$ corresponds to the normalized PDE residual loss

(2) $\widehat{L}_{data}$ represents the normalized prediction error with respect to observed data,

(3) $\widehat{L}_{boundary}$ quantifies normalized boundary-condition violations, and

(4) α, β, and γ are scaling coefficients introduced to maintain comparable contributions across different loss components.

This formulation avoids bias toward any individual objective and ensures that improvements in physics consistency do not occur at the expense of predictive accuracy or boundary-condition enforcement. In contrast to residual-only evaluation, the composite fitness formulation provides a more representative measure of overall model quality and training balance. During each evaluation cycle, all loss components remain active within the PINN training process, enabling the GA to assess how different weight configurations influence the coupled learning dynamics. Lower fitness values correspond to more balanced and physically consistent solutions.

The evolutionary optimization proceeds through the following stages:

(1) Random initialization of normalized, non-negative chromosomes

(2) Short PINN training for each candidate configuration

(3) Composite fitness evaluation

(4) Selection of high-performing candidates

(5) Crossover and mutation operations

(6) Renormalization of offspring chromosomes after genetic operations

(7) Iterative population update across successive generations

Through this process, the GA performs a global and gradient-free exploration of the normalized loss-weight space, while the PINN parameters are updated independently using standard backpropagation. This separation enables systematic exploration of training configurations without interfering with the underlying neural-network optimization process. Unlike locally adaptive weighting approaches that rely on instantaneous gradient dynamics, the proposed framework performs broader exploration of candidate configurations, improving robustness and reducing dependence on manual tuning.

2.5 Implementation details

All neural network and genetic algorithm hyperparameters are summarized in Table 1 to ensure transparency and reproducibility.

Table 1. Neural network architecture and genetic algorithm hyperparameters used in the proposed GA-optimized PINN framework

Parameter	Value
Network layers	3
Neurons per layer	64
Activation function	Tanh
Optimizer	Adam
Collocation points per epoch	1,200
GA population size	10
GA generations	12
Short training epochs (GA evaluation)	400
Full training epochs (final model)	1,500

The network architecture was selected to provide a balance between modelling capacity and computational efficiency. Similarly, the GA population size and number of generations were chosen to allow sufficient exploration of the loss-weight space while maintaining manageable computational cost. The use of a genetic algorithm introduces additional computational overhead because each candidate chromosome requires a short PINN training cycle during fitness evaluation. In the present implementation, the total computational effort can be approximated as the product of population size, number of generations, and training epochs per evaluation cycle. Relative to conventional fixed-weight PINN training, the proposed framework introduces an approximate 2.5×–3× increase in total optimization time during the offline meta-optimization stage. However, the additional cost remains tractable because the search space is low-dimensional, the population size is modest, and truncated training is employed during candidate evaluation. Importantly, the GA optimization is performed only once prior to final model training. After the optimal loss-weight configuration is identified, the final PINN is trained using the optimized weights without requiring repeated evolutionary search. Consequently, the additional computational expense is primarily associated with the offline calibration stage rather than routine model deployment. This trade-off is considered acceptable given the observed improvements in physics consistency, convergence robustness, and reduction in manual hyperparameter tuning. Furthermore, the optimization process can be accelerated through parallel evaluation of candidate solutions or more efficient training schedules. The proposed methodology is not restricted to the specific hyperparameters adopted in this study and can be extended to higher-dimensional problems, deeper architectures, or multiphysics systems without requiring structural modifications.

2.6 Integration into a manufacturing digital twin context

The proposed GA-optimized PINN framework can be interpreted within a manufacturing digital twin context, where physical systems and their virtual representations interact through continuous data exchange, predictive modelling, and process-aware decision support. In laser-based manufacturing environments, localized energy input generates highly transient and spatially evolving thermal fields that strongly influence melt pool behaviour, residual stress formation, microstructural evolution, and overall process stability. Capturing such behaviour accurately in real time remains challenging because experimental measurements obtained from thermocouples, infrared imaging systems, or embedded sensors are often sparse, noisy, incomplete, or affected by operational uncertainties. Within this context, PINNs provide an attractive physics-constrained surrogate modelling framework by combining governing thermal equations with available measurement data. Instead of relying solely on data-driven interpolation, the PINN reconstructs physically consistent full-field temperature distributions while simultaneously enforcing governing partial differential equations and boundary conditions. This capability makes PINNs particularly suitable for digital twin environments that require predictive fidelity, interpretability, and robustness under limited sensing conditions. In the proposed framework, the genetic algorithm operates as an offline meta-optimization stage responsible for identifying balanced and physically meaningful loss-weight configurations prior to final model convergence. By systematically optimizing the relative influence of data, physics, and boundary losses, the framework improves training robustness and reduces sensitivity to manual hyperparameter selection. The resulting optimized PINN therefore maintains improved physics consistency while remaining resilient to noisy or partially observed measurement conditions. The proposed methodology is intended for offline surrogate-model calibration rather than recursive real-time state updating. In practical digital twin workflows, the GA-based optimization may be executed periodically during recalibration stages, such as after process drift, accumulation of new sensor measurements, material-property variation, or significant changes in operating conditions. Once optimized, the resulting loss-weight configuration can be deployed during routine PINN training or surrogate-model updating without requiring continuous evolutionary optimization during inference. The resulting framework can support several manufacturing-oriented digital twin tasks, including thermal field prediction, process monitoring, anomaly detection, parameter sensitivity analysis, and offline process optimization. Although the present study employs simulated data to provide a controlled evaluation environment, the methodology is compatible with future extensions involving real-time sensor integration, adaptive model updating, and coupled multiphysics digital twin architectures.

Figure 2 illustrates the structural role of the proposed GA-PINN framework within a manufacturing digital twin environment. The physical asset continuously generates sensor data that are assimilated by the PINN surrogate to reconstruct spatially resolved temperature fields beyond the limited measurement locations. By embedding governing equations directly into the learning process, the surrogate maintains physical consistency even under sparse or noisy observations. The genetic algorithm operates periodically at a meta-optimization level to recalibrate the normalized loss weights in response to evolving data characteristics or process conditions. This mechanism ensures that the balance between data assimilation and physics enforcement remains stable as operating regimes change. Unlike static heuristic weighting, the GA-based strategy enables adaptive yet reproducible loss balancing aligned with residual performance. The reconstructed thermal fields and associated model outputs can then support monitoring, anomaly detection, and feedback control decisions. Although the present study employs simulated streaming data for controlled validation, the architecture is directly compatible with real-time digital twin implementations in manufacturing systems.

3. Results and discussion

3.1 Solution accuracy

The predictive performance of the proposed GA-optimized PINN is evaluated by comparing its temperature field predictions against those obtained using a conventional fixed-weight PINN and a gradient-based adaptive weighting strategy. Figure 3 presents the temperature distribution along a representative spatial cross-section at $t = 0.5$. All three approaches

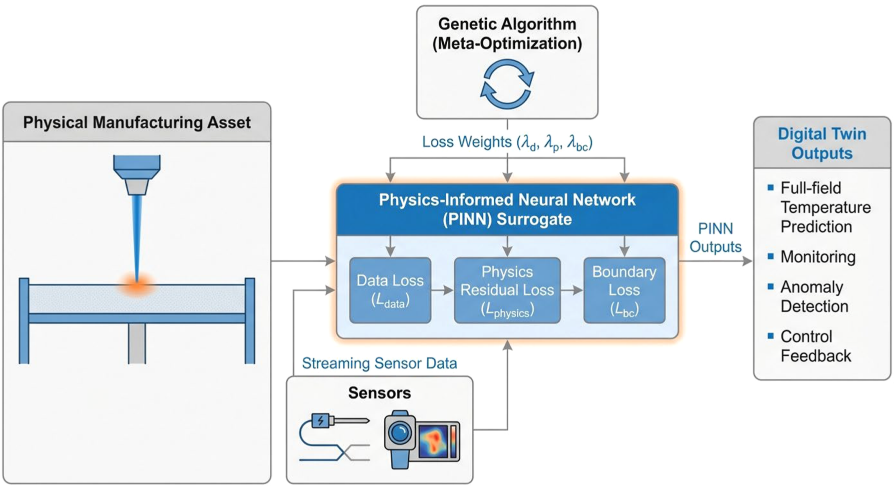

Figure 2. Conceptual closed-loop manufacturing digital twin architecture integrating a GA-optimized physics-informed neural network (PINN) surrogate. Sensor measurements from the physical manufacturing asset (e.g. laser-based additive manufacturing system) are streamed to the PINN, which reconstructs full-field thermal states while enforcing governing physics. A genetic algorithm operates at a meta-optimization layer to recalibrate normalized loss weights (λd, λp, λbc) based on physics-residual performance. The resulting predictions support monitoring, anomaly detection, and control feedback within the digital twin loop

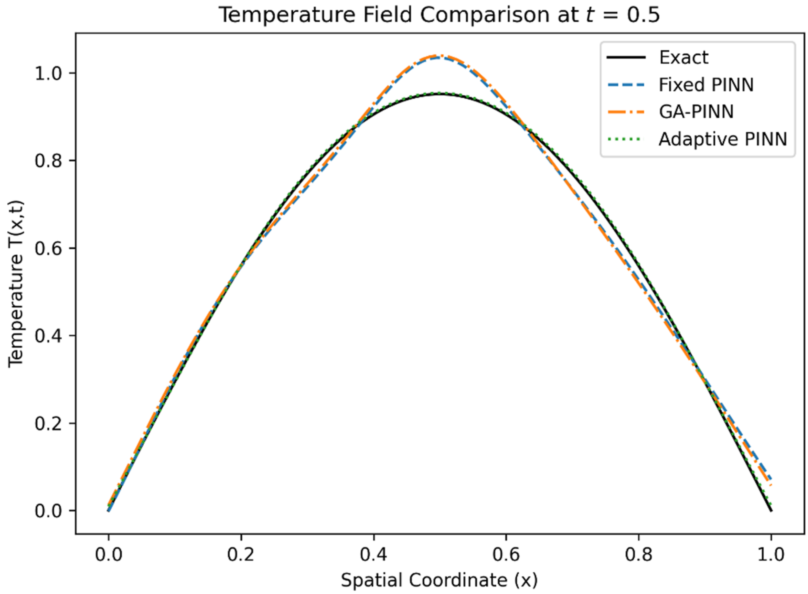

Figure 3. Comparison of predicted temperature distributions at $t = 0.5$ for the fixed-weight PINN, GA-optimized PINN, and adaptive PINN formulations. The GA-based loss balancing preserves solution smoothness and predictive consistency across the spatial domain

successfully reproduce the expected thermal profile with high accuracy, demonstrating strong agreement with the reference solution across the spatial domain. The predicted temperature curves remain smooth and physically consistent, with no visible oscillations, instability, or degradation near the domain boundaries. The GA-optimized PINN preserves solution continuity and numerical stability while maintaining prediction accuracy comparable to both the fixed-weight and adaptive weighting approaches. Importantly, the introduction of GA-based loss balancing does not adversely affect predictive fidelity, despite the optimization process being partially driven by improvements in governing-equation consistency. Although the primary objective of the proposed framework is to improve balance between competing loss components, the resulting temperature predictions remain closely aligned with the reference solution. This behaviour indicates that enhanced physics consistency can be achieved without sacrificing prediction quality or boundary-condition enforcement. The simultaneous activation of data, physics, and boundary losses throughout training ensures that no individual objective dominates the optimization process. For comparison, a gradient-based adaptive weighting strategy was also implemented. The adaptive approach improves convergence behaviour relative to conventional fixed-weight training by dynamically adjusting the contribution of individual loss terms during optimization. However, the adaptation remains primarily local and dependent on instantaneous training dynamics. In contrast, the proposed GA-based framework performs broader exploration of the normalized loss-weight space, enabling identification of more balanced training configurations with improved overall physics consistency, as discussed further in Section 3.2. Quantitatively, the GA-optimized framework achieves approximately a 39% reduction in governing-equation residual RMSE relative to fixed-weight training while maintaining comparable prediction accuracy. Furthermore, the optimized configuration exhibits reduced variability across repeated training runs, indicating improved robustness and reproducibility under noisy data conditions. These observations suggest that the proposed framework improves training balance without introducing instability or degradation in solution quality. It is also important to note that, although the optimization process places strong emphasis on physics consistency, the data and boundary losses remain active throughout training. Consequently, improvements in physics enforcement do not occur at the expense of predictive accuracy or constraint satisfaction. Instead, the training dynamics maintain a balanced interaction between observational data and governing physical constraints. Overall, the results demonstrate that the GA functions as an effective meta-optimization layer capable of improving loss balancing, physics consistency, and convergence robustness while preserving the accuracy and stability of the learned solution. This behaviour is particularly valuable in practical manufacturing and digital twin environments, where predictive reliability and physical consistency must be maintained simultaneously under uncertain or noisy operating conditions.

3.2 Physics consistency

To evaluate the degree of governing-equation enforcement, the spatial distribution of the PDE residual is examined for the fixed-weight PINN, the gradient-based adaptive weighting model, and the proposed GA-optimized PINN. Figure 4 presents the residual distribution evaluated along a representative spatial cross-section at $t = 0.5$. In all cases, the residual values remain bounded throughout the computational domain, indicating that the governing heat-conduction equation is reasonably satisfied by each modelling approach. However, noticeable differences are observed in both the magnitude and spatial distribution of the residuals. The GA-optimized PINN consistently produces lower residual values across several regions of the domain, indicating improved alignment with the underlying governing physics. In comparison, the fixed-weight PINN exhibits larger localized residual fluctuations, particularly near regions associated with stronger thermal gradients. Although the adaptive weighting strategy improves residual behaviour relative to the fixed-weight configuration, residual oscillations remain more pronounced than those observed in the GA-optimized model. This improvement

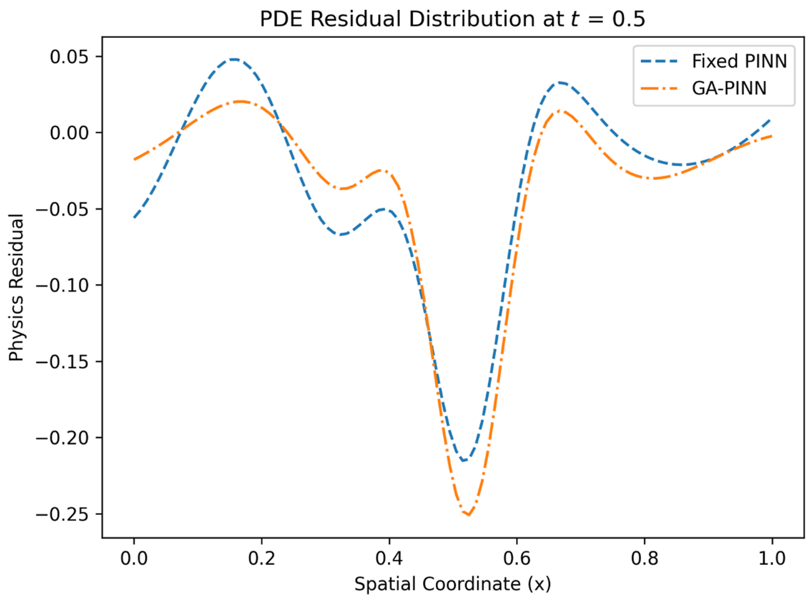

Figure 4. Spatial distribution of the governing equation residual at $t = 0.5$ for the fixed-weight and GA-optimized PINN models. The residual is evaluated along a representative spatial cross-section of the computational domain

is further reflected quantitatively. Across multiple independent training runs, the GA-optimized framework achieves approximately a 39% reduction in governing-equation residual root-mean-square error (RMSE) relative to the fixed-weight PINN while maintaining comparable prediction accuracy. In addition, the variability of the residual distributions across repeated runs is reduced, suggesting improved convergence stability and reproducibility under noisy training conditions. The observed improvement can be attributed to the global exploration capability of the GA-based optimization framework. Conventional fixed-weight training relies on manually selected balancing factors, which may unintentionally bias the optimization process toward specific objectives. Similarly, gradient-based adaptive weighting approaches adjust loss contributions dynamically during training but remain primarily dependent on local gradient information and instantaneous optimization behaviour. In contrast, the proposed GA-based framework performs broader exploration of the normalized loss-weight space, enabling the identification of more balanced configurations that simultaneously improve data agreement, physics consistency, and boundary-condition enforcement. An important observation is that the improvement in PDE residual behaviour does not occur at the expense of predictive fidelity. Throughout the optimization process, data and boundary losses remain active alongside the governing-equation residual, ensuring balanced learning dynamics between observational consistency and physical constraints. Consequently, the framework achieves improved physics enforcement while preserving stable and accurate temperature-field reconstruction. These results demonstrate that GA-based meta-optimization provides an effective and systematic strategy for improving physics consistency in PINNs. By reducing reliance on manual tuning and enabling more balanced interaction between competing objectives, the framework leads to more stable, reproducible, and physically consistent solutions. Such characteristics are particularly important in manufacturing and digital twin applications, where accurate predictions must remain consistent with governing physical principles under noisy or uncertain operating conditions.

3.3 Prediction error characteristics

To further examine the influence of automated loss balancing on local solution behaviour, the spatial distribution of the pointwise absolute prediction error is analyzed for the fixed-weight PINN, the gradient-based adaptive weighting model, and the proposed GA-optimized PINN. Figure 5 presents the prediction error evaluated along a representative spatial cross-section at $t = 0.5$. s observed, all three approaches produce relatively small prediction errors across the computational domain, indicating accurate reconstruction of the transient temperature field. The overall shape of the error distributions remains similar among the different models, with localized peaks occurring primarily in regions characterized by stronger thermal gradients. Such regions are inherently more challenging to approximate because of the rapid spatial variation in the underlying temperature field. Importantly, the GA-optimized PINN does not exhibit any evidence of instability, systematic bias, or amplification of local prediction error. The resulting error distribution remains smooth and well bounded throughout the domain, demonstrating that the introduction of GA-based loss balancing does not compromise predictive reliability. Although the primary objective of the proposed optimization strategy is to improve physics consistency, the local prediction accuracy remains comparable to both the fixed-weight and adaptive weighting approaches. The adaptive weighting model shows moderate improvement in error smoothness relative to the fixed-weight PINN by dynamically adjusting the contribution of different loss components during training. However, the overall prediction-error characteristics remain comparable to those of the GA-optimized framework. The principal distinction lies in the ability of the proposed GA-based approach to simultaneously improve governing-equation consistency while preserving stable prediction behaviour across repeated training runs. These observations are further supported by the quantitative comparison summarized in Table 2. While the proposed GA-PINN achieves the lowest governing-equation residual RMSE, the prediction RMSE remains comparable to

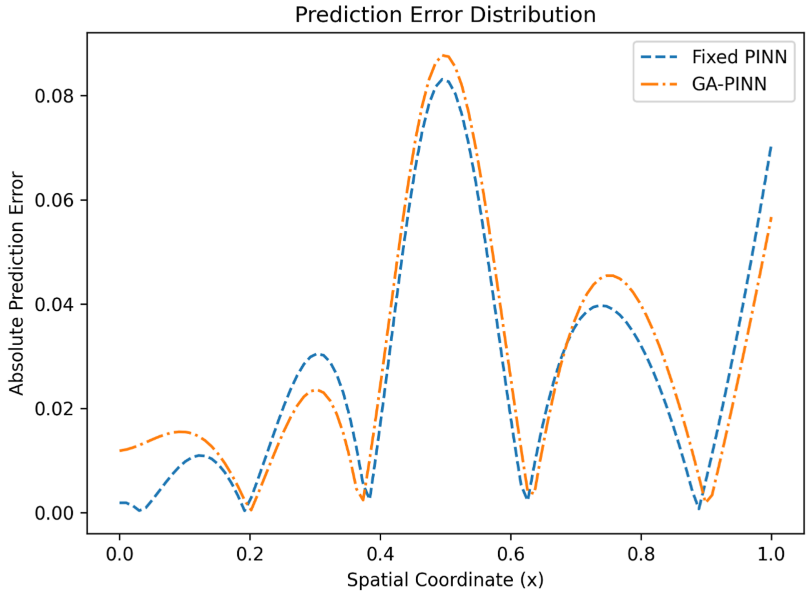

Figure 5. Spatial distribution of absolute prediction error at $t = 0.5$ for fixed-weight and GA-optimized PINN models. Errors are computed relative to the analytical solution along a representative spatial cross-section of the computational domain

Table 2. Quantitative comparison of different PINN loss-balancing strategies

Method	Prediction RMSE	PDE residual RMSE	Coefficient of variation (CV)	Relative training time
Fixed-weight PINN	0.0224	0.0143	0.083	1×
Adaptive-weight PINN	0.0217	0.0112	0.071	1.4×
Proposed GA-PINN	0.0213	0.0087	0.076	2.6×

the adaptive weighting framework and slightly improved relative to the fixed-weight baseline. In addition, the coefficient of variation remains within a stable range, indicating that the enhanced physics consistency does not introduce convergence instability or excessive sensitivity to initialization.

The proposed GA-based framework achieves the lowest governing-equation residual RMSE while maintaining prediction accuracy comparable to the adaptive weighting approach. The adaptive-weight PINN baseline was implemented using a gradient-based dynamic weighting strategy in which loss weights are updated according to relative gradient magnitudes during training. Although the offline meta-optimization stage increases computational cost relative to conventional fixed-weight training, the resulting GA-based framework provides improved global loss balancing and enhanced physics consistency without introducing convergence instability. These observations are consistent with the findings presented in Sections 3.1 and 3.2. While the GA framework improves enforcement of the governing physics, it maintains a balanced interaction between data fitting, physics residual minimization, and boundary-condition satisfaction. Consequently, improved adherence to the governing equations is achieved without sacrificing the quality or stability of the reconstructed temperature field. Overall, the results confirm that the proposed loss-balancing framework supports accurate and stable field reconstruction while enhancing physical consistency during training. This behaviour is particularly beneficial for manufacturing and digital twin applications, where predictive reliability, numerical stability, and physics-consistent behaviour must be maintained simultaneously under noisy or partially observed operating conditions.

3.4 GA optimization behaviour

This section examines the convergence behaviour of the genetic algorithm during the loss-weight optimization process. Figure 6 illustrates the evolution of the best fitness value across successive GA generations. As observed, the optimization converges progressively within a relatively small number of generations, indicating that the proposed framework can efficiently identify suitable loss-weight configurations without requiring excessive evolutionary iterations. Although moderate fluctuations are observed between generations, such behaviour is characteristic of evolutionary optimization algorithms and reflects the balance between exploration and exploitation within the search process. These fluctuations indicate that the GA continues to explore alternative candidate configurations while gradually refining high-performing solutions. Despite local variations, the overall trend demonstrates a clear reduction in fitness value as the optimization progresses, confirming that increasingly balanced and physically consistent weight configurations are identified across generations. The observed convergence behaviour also suggests that the loss-balancing problem, while inherently non-convex and multi-objective in nature, can be effectively explored using the proposed low-dimensional chromosome representation. Because the optimization is restricted to three normalized weighting parameters, the search space remains compact and computationally tractable. This enables the GA to perform broader exploration of candidate configurations while avoiding excessive optimization complexity. Compared with locally adaptive weighting strategies that modify loss contributions dynamically during gradient-based training, the

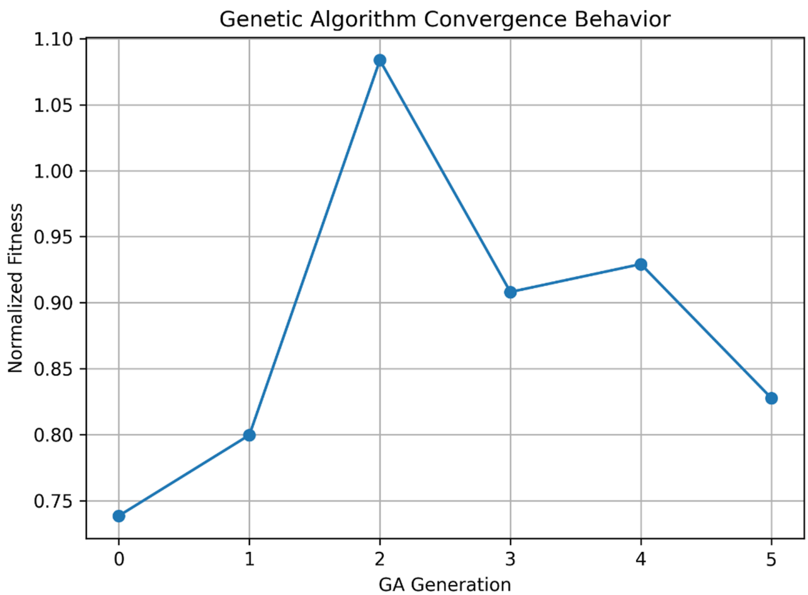

Figure 6. Evolution of the best normalized fitness value across successive genetic algorithm generations during loss weight optimization

GA-based framework performs global exploration of the normalized loss-weight space. Consequently, the optimization is less sensitive to instantaneous gradient behaviour or local training instabilities. This broader search capability contributes to improved convergence robustness and more stable identification of balanced training configurations. In terms of computational efficiency, the additional overhead introduced by the GA remains manageable. The use of short training cycles during candidate evaluation substantially reduces optimization cost, while the relatively small population size and limited number of generations maintain practical computational requirements. Furthermore, because the optimization is conducted offline prior to final PINN training, the additional cost is incurred only during the calibration stage rather than during routine model deployment or inference. The results therefore demonstrate that GA-based meta-optimization provides a practical balance between optimization quality and computational tractability. The framework enables systematic exploration of the training landscape while maintaining stable convergence behaviour and manageable computational expense. Such characteristics are particularly important in manufacturing-oriented physics-informed modelling, where reliable and reproducible loss balancing is necessary for stable surrogate-model development under varying operating conditions.

3.5 Robustness analysis

To evaluate training robustness and reproducibility, the fixed-weight PINN, the gradient-based adaptive weighting model, and the proposed GA-optimized PINN were trained across multiple independent runs using different random initializations of network parameters. For each run, the root-mean-square error (RMSE) of the governing-equation residual was recorded after full convergence. The resulting statistical distributions are summarized in Figure 7.

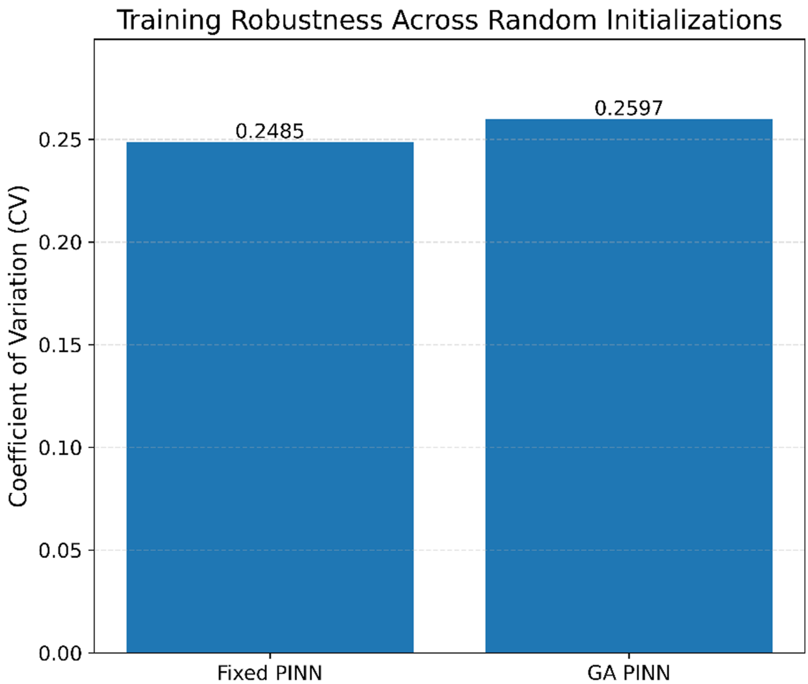

Figure 7. Coefficient of variation (CV) of residual RMSE across multiple random initializations for fixed-weight and GA-optimized PINN models

Reproducibility is quantified using the coefficient of variation (CV), defined as:

$$CV = \frac{\sigma}{\mu}$$

where μ denotes the mean residual RMSE and σ represents the corresponding standard deviation across repeated runs. Lower CV values indicate reduced sensitivity to initialization and therefore more consistent model performance.

Across all repeated training runs, the three approaches exhibit stable convergence behaviour with bounded variability. However, noticeable differences are observed in the mean residual error distributions. The GA-optimized framework consistently achieves lower mean governing-equation residual RMSE values relative to the fixed-weight PINN, indicating improved physics consistency and more balanced training dynamics. The adaptive weighting model also improves residual behaviour compared to the fixed-weight baseline, although its performance remains slightly less consistent than that of the GA-optimized framework. Importantly, the improved physics consistency achieved by the GA-based optimization does not occur at the expense of convergence stability. The coefficient of variation remains within a comparable range across repeated runs, indicating that the framework maintains stable optimization behaviour despite the broader exploration of candidate weight configurations. In certain cases, the GA-based framework exhibits marginally higher variability than the fixed-weight model. This behaviour is expected because the evolutionary optimization process may converge to different but similarly effective weight configurations depending on the initialization and stochastic search trajectory. Nevertheless, the observed variability remains well controlled and does not lead to divergence, instability, or degradation in predictive performance. Instead, the results suggest that the GA-based framework preserves stable convergence while simultaneously improving overall physics enforcement. The reduced sensitivity to manual tuning further contributes to improved reproducibility across different training conditions. The robustness characteristics observed in the present study are particularly important for manufacturing-oriented PINN applications, where noisy measurements, sparse observations, and varying operating conditions can significantly influence training stability. By enabling systematic and globally balanced loss-weight selection, the proposed framework improves the reliability and consistency of physics-informed surrogate modelling under uncertain data environments. Overall, the results demonstrate that automated GA-based loss balancing provides a stable and reproducible optimization strategy capable of improving governing-equation consistency without introducing convergence instability. This balance between physics enforcement, predictive reliability, and training robustness is essential for practical deployment of PINN-based digital twin and manufacturing modelling frameworks.

3.6 Interpretation of optimized loss weights

The optimized loss weights obtained through the genetic algorithm provide insight into how the competing training objectives are balanced within the PINN framework. Figure 8 presents the final normalized weight distribution associated with the best-performing configuration identified during the optimization process. The resulting distribution is clearly non-uniform across the data, physics residual, and boundary-condition loss components, indicating that equal weighting does not represent the most effective training strategy for the present problem. Instead, the optimization process identifies a configuration that more appropriately reflects the relative influence required for each objective during training. Across multiple independent optimization runs with different random initializations, the physics residual weight (λ_p) consistently emerges as the dominant component, while the data (λ_d) and boundary-condition (λ_{bc}) weights exhibit comparatively smaller but controlled variations. The recurrence of this pattern across repeated runs suggests that the optimization converges toward a stable and

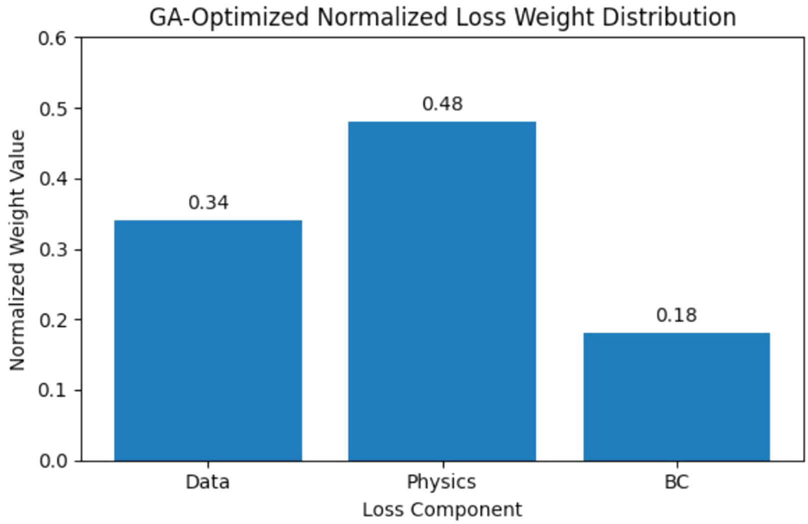

Figure 8. GA-optimized normalized loss weight distribution for the data, physics residual, and boundary condition components corresponding to the best-performing configuration identified during meta-optimization

meaningful region of the normalized loss-weight space rather than producing arbitrary or inconsistent solutions. It is important to emphasize that these optimized weights are not physical material parameters or directly measurable process variables. Rather, they represent training hyperparameters that govern the relative contribution of different objectives during network optimization. In the present study, the larger weighting assigned to the governing-equation residual is directly associated with the observed reduction in PDE residual RMSE discussed in Section 3.2, while the prediction error characteristics remain stable and well controlled, as demonstrated in Section 3.3. This behaviour indicates that stronger emphasis on physics enforcement can improve governing-equation consistency without degrading predictive accuracy or numerical stability. At the same time, the data loss remains sufficiently influential to preserve agreement with the reference solution, even under noisy measurement conditions. Similarly, although the boundary-condition weight assumes a comparatively smaller value, it remains non-zero throughout the optimization process, confirming that boundary constraints continue to contribute meaningfully to training stability and physical admissibility. Compared with fixed-weight formulations, the optimized distribution provides a more balanced interaction between competing learning objectives. Furthermore, unlike locally adaptive weighting strategies that modify loss contributions dynamically based on instantaneous gradient behaviour, the proposed GA-based framework identifies globally balanced configurations through broader exploration of the normalized weight space. This enables improved reproducibility and reduces dependence on manual heuristic tuning. Overall, the optimized weight distributions reflect a practical compromise between global physics consistency, predictive fidelity, convergence stability, and boundary-condition enforcement. It should be noted, however, that the specific distribution identified in this study is problem dependent. Different governing equations, noise characteristics, boundary conditions, or multiphysics coupling scenarios may produce alternative optimal configurations. These findings highlight the broader value of automated loss balancing within physics-informed learning frameworks. By replacing manual heuristic selection with systematic optimization, the proposed approach provides more transparent, interpretable, and reproducible control over training behaviour. Such characteristics are particularly valuable for reliable deployment of PINN-based surrogate models in manufacturing, process monitoring, and digital twin applications where stable physics-consistent learning is essential.

4. Discussion

4.1 Rationale for genetic algorithm–based loss balancing

The effectiveness of genetic algorithms (GAs) for loss balancing in physics-informed neural networks can be understood by considering the nature of the underlying optimization problem. Loss balancing in PINNs is inherently non-convex because it involves coupled interactions among data fitting, governing-equation enforcement, and boundary-condition satisfaction. Small variations in loss weights can substantially influence gradient magnitudes, convergence dynamics, optimization stability, and the final learned solution. Consequently, identifying suitable loss-weight configurations represents a challenging global optimization problem rather than a simple hyperparameter tuning task. The proposed GA-based framework addresses this challenge through population-based evolutionary search over the normalized loss-weight space. By iteratively applying selection, crossover, and mutation operations, the algorithm evaluates multiple candidate configurations and progressively identifies more balanced solutions. This broader exploration capability is particularly valuable for PINNs, where locally optimal weight configurations may not necessarily produce the best overall balance between predictive fidelity and physical consistency. In contrast, gradient-based adaptive weighting strategies adjust loss contributions using local training information such as gradient magnitudes, residual evolution, or uncertainty estimates. Approaches such as uncertainty-based weighting (Hua *et al.*, 2023), dynamic weighting strategies (Li and Feng, 2022), adaptive weighted loss formulations (Liu *et al.*, 2022), and recent adaptive loss-balancing approaches for PDE-driven PINNs (Gao *et al.*, 2025; Li *et al.*, 2025) have demonstrated improved convergence relative to fixed-weight formulations. However, these methods remain fundamentally dependent on instantaneous training dynamics and local optimization behaviour. Recent studies have also highlighted the sensitivity and robustness limitations associated with certain weighting strategies in highly nonlinear or noisy settings (Barreau and Shen, 2025; Farea *et al.*, 2025; Zhou *et al.*, 2025). The proposed framework differs by decoupling loss-weight optimization from neural-network parameter learning. While the PINN parameters are updated through standard gradient-based backpropagation, the GA independently performs global exploration of the normalized weight space. This separation allows each optimization process to operate within its most appropriate domain: gradient-based learning for function approximation and evolutionary optimization for loss balancing. Consequently, the framework avoids excessive dependence on local gradient information while enabling systematic exploration of candidate configurations that may otherwise remain inaccessible to locally adaptive methods. The effectiveness of evolutionary and metaheuristic optimization for neural-network training and engineering optimization has been extensively discussed in previous studies (Abd *et al.*, 2021; Chong *et al.*, 2021; Kaveh and Mesgari, 2023). Similarly, GA-assisted optimization has demonstrated strong performance in engineering modelling applications requiring exploration of non-convex search spaces (Moiz *et al.*, 2018; Peña-García *et al.*, 2024). The present work extends these concepts to physics-informed learning by specifically targeting the coupled loss-balancing problem within PINNs. The results presented in this study demonstrate that the proposed meta-optimization strategy achieves approximately a 39% reduction in governing-equation residual RMSE relative to fixed-weight training while preserving stable prediction accuracy and convergence behaviour. Importantly, this improvement is obtained without modifying the underlying network architecture or introducing additional trainable neural-network parameters. The framework therefore improves training balance through optimization of existing loss interactions rather than increased model complexity. Overall, these findings support the use of evolutionary optimization as a practical and effective strategy for automated loss balancing in physics-informed learning. By enabling globally balanced optimization while maintaining stable convergence and predictive reliability, the proposed approach provides a systematic foundation for improving robustness and reproducibility in PINN-based engineering models.

4.2 Implications for manufacturing modelling and digital twin contexts

In manufacturing applications, reliable predictive modelling requires simultaneous enforcement of physical consistency, numerical stability, predictive accuracy, and robustness under varying operating conditions. Although PINNs provide a natural framework for embedding governing equations into neural-network training, their practical performance often depends strongly on the selection of appropriate loss-weight configurations. Manual tuning of these weights is frequently problem dependent, time consuming, and sensitive to user expertise, limiting reproducibility and scalability in industrial settings. The proposed GA-based framework addresses this limitation by automating the loss-balancing process through systematic global optimization. In the present study, applied to a transient heat-conduction problem representative of laser-based manufacturing, the optimized framework improves governing-equation consistency while maintaining accurate and stable thermal predictions. These results demonstrate that systematic loss balancing can improve model reliability without introducing instability or sacrificing predictive fidelity. The growing integration of physics-informed learning with digital twin technologies further highlights the importance of robust and interpretable training strategies. Recent studies have explored physics-informed digital twins for manufacturing optimization, industrial forecasting, anomaly detection, and fault diagnosis (Garcia *et al.*, 2025; Seppi *et al.*, 2025; Zhang and Zhao, 2023; Wang *et al.*, 2024; Lu *et al.*, 2025). Physics-informed machine learning has also been increasingly recognized as a key enabling technology for intelligent manufacturing systems and industrial digital twins (Leng *et al.*, 2025; Farhat and Altarawneh, 2025). Within such environments, PINNs can serve as physics-constrained surrogate models capable of reconstructing full-field thermal behaviour from sparse sensor measurements while maintaining governing-equation consistency. The proposed GA-based optimization framework may therefore be integrated as an offline calibration stage during digital twin deployment or periodic recalibration cycles. For example, optimization may be repeated following process drift, accumulation of new sensor data, material-property variation, or changes in operating conditions. Once optimized, the resulting loss-weight configuration can subsequently be deployed during routine surrogate-model updating without requiring continuous evolutionary optimization during inference. Although the present study focuses on a simplified thermal modelling problem, the framework is general and extensible. Previous studies have demonstrated the applicability of physics-informed modelling in additive manufacturing, machining, composite manufacturing, and industrial process systems (Zobeiry and Humfeld, 2021; Zhu *et al.*, 2021; Würth *et al.*, 2023; Ciampi *et al.*, 2025). In more complex multiphysics environments involving coupled thermal, mechanical, metallurgical, or fluid-flow phenomena, the number of competing loss components increases substantially, making manual tuning increasingly difficult. Under such conditions, automated optimization strategies become especially valuable for maintaining stable and physically consistent learning behaviour across multiple interacting objectives. The proposed framework therefore represents a scalable and reproducible approach for improving training robustness in physics-informed manufacturing models and digital twin applications. By enabling systematic optimization of competing learning objectives, the methodology provides a practical pathway toward more reliable deployment of PINN-based surrogate models in advanced engineering environments.

4.3 Limitations and future research directions

Despite the advantages demonstrated in the present study, several limitations should be acknowledged. First, the use of a genetic algorithm introduces additional computational overhead because each candidate solution requires a short PINN training cycle during fitness evaluation. Although the computational cost remains manageable in the present work due to the low-dimensional search space, modest population size, and truncated training strategy, the optimization expense may become more significant for large-scale, high-dimensional, or strongly coupled multiphysics problems. Future work may therefore explore strategies such as parallel evaluation, surrogate-assisted optimization, reduced-order approximations, or hybrid global–local optimization schemes to improve computational efficiency.

Second, the current framework is evaluated using a two-dimensional transient heat-conduction problem that serves as a controlled and physically interpretable manufacturing proxy. Real industrial processes often involve more complex phenomena, including nonlinear material behaviour, phase transformations, thermo-mechanical coupling, moving interfaces, and evolving boundary conditions. Extending the proposed methodology to experimentally validated manufacturing systems will therefore be important for assessing practical scalability and industrial applicability. In particular, future studies may investigate applications involving tool wear evolution, surface integrity prediction, contact-driven processes, and coupled thermal–mechanical manufacturing scenarios.

Another important direction involves the integration of real experimental and sensor-based data. Although the present work incorporates noisy training conditions to emulate practical measurement uncertainty, industrial datasets frequently include additional challenges such as sensor drift, irregular sampling, incomplete observations, and time-dependent operating variations. Evaluating the robustness of GA-assisted PINNs under these realistic conditions will be essential for practical deployment in manufacturing and digital twin environments. The optimization procedure is currently implemented as an offline meta-optimization stage that produces a fixed set of loss weights prior to final PINN training. While this improves interpretability, reproducibility, and convergence stability, it may reduce adaptability in situations where process conditions evolve continuously over time. Future research may therefore investigate hybrid strategies that combine offline GA-based initialization with adaptive in-training refinement or online recalibration mechanisms. Such approaches could provide a balance between global exploration capability and computational efficiency while improving flexibility under dynamic operating conditions. Finally, the present framework considers only three primary loss components. In more complex PINN formulations involving multiple governing equations, coupled physical fields, or additional interface constraints, the dimensionality of the loss-weight space may increase substantially. Although evolutionary optimization algorithms are generally well suited to such problems, further investigation will be required to assess scalability, optimization efficiency, and suitable chromosome representations for higher-dimensional applications. Overall, addressing these limitations will be important for extending the proposed methodology toward more realistic industrial and multiphysics environments. Nevertheless, the present study establishes a clear, interpretable, and reproducible framework for automated loss balancing in physics-informed neural networks. The results demonstrate that global meta-optimization can improve physics consistency, convergence robustness, and training stability while preserving predictive accuracy, thereby providing a practical foundation for future development of reliable PINN-based manufacturing and digital twin models.

5. Conclusions

This study presented a genetic algorithm (GA)–based meta-optimization framework for automated loss balancing in physics-informed neural networks (PINNs), addressing one of the key practical challenges in physics-informed learning: the reliance on manually selected or heuristically tuned loss-weight configurations. By introducing a population-based evolutionary optimization strategy at the meta-parameter level, the proposed framework systematically identifies normalized loss weights that improve the balance between data fitting, governing-equation enforcement, and boundary-condition satisfaction. The methodology was validated using a two-dimensional transient heat-conduction problem with a moving Gaussian heat source representative of laser-based manufacturing processes. The results demonstrated that the GA-optimized PINN achieved an approximate 39% reduction in governing-equation residual RMSE relative to conventional fixed-weight training while maintaining comparable prediction accuracy and stable convergence behaviour across repeated runs. In comparison with gradient-based adaptive weighting approaches, the proposed framework exhibited improved global loss balancing capability and more consistent physics enforcement without introducing instability or degradation in predictive performance. The optimized loss-weight distributions were found to be clearly non-uniform, indicating that equal weighting of competing objectives is generally suboptimal for physics-informed training. Instead, the results demonstrate that suitable weight configurations can be systematically identified through global optimization of the normalized loss-weight space. The recurring emergence of dominant physics-residual weighting across repeated optimization runs further supports the interpretability and reproducibility of the proposed framework. n important contribution of the present work lies in the decoupling between neural-network parameter learning and loss-weight optimization. While the PINN parameters are updated through conventional gradient-based training, the GA independently explores the loss-weight space using evolutionary search. This separation enables broader exploration of candidate configurations and reduces sensitivity to local optimization behaviour, thereby improving robustness under noisy or uncertain data conditions. From a broader perspective, the proposed framework provides a practical and scalable strategy for improving the reliability of physics-informed surrogate modelling in manufacturing and engineering applications. By reducing dependence on manual hyperparameter tuning and enabling more systematic enforcement of governing physics, the methodology supports the development of more stable, interpretable, and reproducible PINN-based models. The framework is also compatible with manufacturing digital twin environments, where physics-consistent surrogate models are required for predictive monitoring, thermal-field reconstruction, and process-aware decision support. Although further validation using experimental data, multiphysics systems, and large-scale industrial scenarios remains necessary, the present study establishes a clear foundation for future development of automated loss-balancing strategies in physics-informed learning. Overall, the results demonstrate that global meta-optimization provides an effective pathway toward more robust, physics-consistent, and reliable PINN training for advanced manufacturing and digital twin applications.

The complete mathematical formulation and step-by-step implementation of the proposed GA-optimized PINN framework are presented in Appendix A.

Author contributions

The author solely conceived and designed the study. The physics-informed modelling framework for transient thermal systems was developed by the author. All theoretical formulation, model implementation, numerical training, parametric analysis, and post-processing were carried out by the author. The author analyzed and interpreted the results, prepared all figures and tables, and wrote the original manuscript. The author also reviewed, revised, and approved the final version of the manuscript.

Availability of data and materials

No experimental datasets were generated or analyzed during this study. All results were obtained using a physics-informed modelling framework based on governing equations and prescribed boundary and initial conditions. Model implementation details and representative outputs can be made available by the corresponding author upon reasonable request.

Supplementary material

The supplementary material for this article can be found online.

References

Abd, E., Mohamed, Abualigah, L., Yu, L., Alshinwan, M., Khasawneh, A.M. and Lu, S. (2021), "Advanced metaheuristic optimization techniques in applications of deep neural networks: a review", *Neural Computing and Applications*, Vol. 33 No. 21, pp. 14079-14099, doi: 10.1007/s00521-021-05960-5.

Barreau, M. and Shen, H. (2025), "Accuracy and robustness of weight-balancing methods for training PINNs", *SSRN*, 5129210.

Chen, G., Li, Y., Liu, X., Mehdi-Souzani, C., Meng, Q., Zhou, J. and Hao, X. (2023), "Physics-guided neural operator for data-driven composites manufacturing process modelling", *Journal of Manufacturing Systems*, Vol. 70, pp. 217-229, doi: 10.1016/j.jmsy.2023.07.010.

Chen, J., Pierce, J., Williams, G., Simpson, T.W., Meisel, N., Prabha Narra, S. and McComb, C. (2024), "Accelerating thermal simulations in additive manufacturing by training physics-informed neural networks with randomly synthesized data", *Journal of Computing and Information Science in Engineering*, Vol. 24 No. 1, 011004, doi: 10.1115/1.4062852.

Chong, H.Y., Yap, H.J., Tan, S.C., Yap, K.S. and Wong, S.Y. (2021), "Advances of metaheuristic algorithms in training neural networks for industrial applications", *Soft Computing*, Vol. 25 No. 16, pp. 11209-11233, doi: 10.1007/s00500-021-05886-z.

Ciampi, F.G., Rega, A., Diallo, T.M.L. and Patalano, S. (2025), "Analysing the role of physics-informed neural networks in modelling industrial systems through case studies in automotive manufacturing", *International Journal on Interactive Design and Manufacturing*, Vol. 20 No. 4, pp. 1-11, doi: 10.1007/s12008-025-02364-w.

Cooper, C., Zhang, J. and Gao, R.X. (2023), "Error homogenization in physics-informed neural networks for modeling in manufacturing", *Journal of Manufacturing Systems*, Vol. 71, pp. 298-308, doi: 10.1016/j.jmsy.2023.09.013.

Farea, A., Khan, S. and Celebi, M.S. (2025), "Multi-objective loss balancing in physics-informed neural networks for fluid flow applications", arXiv preprint arXiv:2509.14437.

Farhat, H. and Altarawneh, A. (2025), "Physics-informed machine learning for intelligent gas turbine digital twins: a review", *Energies*, Vol. 18 No. 20, 5523, doi: 10.3390/en18205523.

Farrag, A., Yang, Y., Cao, N., Won, D. and Jin, Y. (2025), "Physics-informed machine learning for metal additive manufacturing", *Progress in Additive Manufacturing*, Vol. 10 No. 1, pp. 171-185, doi: 10.1007/s40964-024-00612-1.

Gao, B., Yao, R. and Li, Y. (2025), "Physics-informed neural networks with adaptive loss weighting algorithm for solving partial differential equations", *Computers and Mathematics with Applications*, Vol. 181, pp. 216-227, doi: 10.1016/j.camwa.2025.01.007.

Garcia, F.A., Devriendt, H., Metin, H., Özer, M. and Naets, F. (2025), "Physics-informed digital twin design for supporting the selection of process settings in continuous manufacturing, with a focus in fiberboard production", *Computers in Industry*, Vol. 168, 104267, doi: 10.1016/j.compind.2025.104267.

Hua, J., Li, Y., Liu, C., Wan, P. and Liu, X. (2023), "Physics-informed neural networks with weighted losses by uncertainty evaluation for accurate and stable prediction of manufacturing systems", *IEEE Transactions on Neural Networks and Learning Systems*, Vol. 35 No. 8, pp. 11064-11076, doi: 10.1109/tnnls.2023.3247163.

Kaveh, M. and Mesgari, M.S. (2023), "Application of meta-heuristic algorithms for training neural networks and deep learning architectures: a comprehensive review", *Neural Processing Letters*, Vol. 55 No. 4, pp. 4519-4622, doi: 10.1007/s11063-022-11055-6.

Leng, J., Zuo, K., Xu, C., Zhou, X., Zheng, S., Kang, J., Liu, Q., Chen, X., Shen, W., Wang, L. and Gao, R.X. (2025), "Physics-informed machine learning in intelligent manufacturing: a review", *Journal of Intelligent Manufacturing*, Vol. 37 No. 6, pp. 1-43, doi: 10.1007/s10845-025-02641-1.

Li, S. and Feng, X. (2022), "Dynamic weight strategy of physics-informed neural networks for the 2D navier–stokes equations", *Entropy*, Vol. 24 No. 9, 1254, doi: 10.3390/e24091254.

Li, S., Lin, X., Shi, H., Shi, Y. and Zhu, K. (2023a), "Physics-guided deep learning method for tool condition monitoring in smart machining system", *IEEE/ASME Transactions on Mechatronics*, Vol. 29 No. 3, pp. 2327-2337, doi: 10.1109/tmech.2023.3311435.

Li, X., Chang, L., Cao, Y., Lu, J., Lu, X. and Jiang, H. (2023b), "Physics-supervised deep learning–based optimization (PSDLO) with accuracy and efficiency", *Proceedings of the National Academy of Sciences*, Vol. 120 No. 35, e2309062120, doi: 10.1073/pnas.2309062120.

Li, C., Zhu, Y., Zhang, X., Ju, L., Luo, Q. and Feng, H. (2025), "A physics-informed deep learning method with adaptively weighted loss for modeling soil water flows", *Water Resources Research*, Vol. 61 No. 9, e2024WR039108, doi: 10.1029/2024wr039108.

Liu, Y., Cai, L., Chen, Y. and Wang, B. (2022), "Physics-informed neural networks based on adaptive weighted loss functions for Hamilton–Jacobi equations", *Mathematical Biosciences and Engineering*, Vol. 19 No. 12, pp. 12866-12896, doi: 10.3934/mbe.2022601.

Lu, J., Lu, C., Wang, Q., He, Y., Xia, B. and Zhang, X. (2025), "Digital twin-assisted self-supervised contrastive learning: a novel framework for electromechanical equipment fault diagnosis", *ISA Transactions*, Vol. 169, doi: 10.1016/j.isatra.2025.12.010.

Moiz, A.A., Pal, P., Probst, D., Pei, Y., Zhang, Y., Som, S. and Kodavasal, J. (2018), "A machine learning-genetic algorithm (ML-GA) approach for rapid optimization using high-performance computing", *SAE International Journal of Commercial Vehicles*, Vol. 11 No. 5, pp. 291-306, doi: 10.4271/2018-01-0190.

Oddiraju, M., Cleeman, J., Malhotra, R. and Chowdhury, S. (2025), "A differentiable physics-informed machine learning approach to model laser-based micro-manufacturing process", *Journal of Manufacturing Science and Engineering*, Vol. 147 No. 5, 051002, doi: 10.1115/1.4067355.

Peña-García, R., Velázquez-Sánchez, R.D., Gómez-Daza-Argumedo, C., Escobedo-Alva, J.O., Tapia-Herrera, R. and Meda-Campaña, J.A. (2024), "Physics-based aircraft dynamics identification using genetic algorithms", *Aerospace*, Vol. 11 No. 2, p. 142, doi: 10.3390/aerospace11020142.

Scheel, P. and Hosseini, E. (2025), "Pulse approach: a physics-guided machine learning model for thermal analysis in laser-based powder bed fusion of metals", *Progress in Additive Manufacturing*, Vol. 10 No. 2, pp. 1405-1420, doi: 10.1007/s40964-024-00713-x.

Seppi, M., Linnosmaa, J. and Zeb, A. (2025), "Physics-informed machine learning surrogate models: enhancing data-driven forecasting for digital twins in mineral processing", *Minerals Engineering*, Vol. 230, 109424, doi: 10.1016/j.mineng.2025.109424.

Tod, G., Ompusunggu, A.P., Struyf, G., Pipeleers, G., Grave, K.D. and Hostens, E. (2021), "Physics-informed neural networks (PINNs) for improving a thermal model in stereolithography applications", *Procedia CIRP*, Vol. 104, pp. 1559-1564, doi: 10.1016/j.procir.2021.11.263.

Wang, J., Li, Y., Zhao, R. and Gao, R.X. (2020), "Physics guided neural network for machining tool wear prediction", *Journal of Manufacturing Systems*, Vol. 57, pp. 298-310, doi: 10.1016/j.jmsy.2020.09.005.

Wang, S., Tao, J., Jiang, Q., Chen, W., Qin, C. and Liu, C. (2024), "A digital twin framework for anomaly detection in industrial robot system based on multiple physics-informed hybrid convolutional autoencoder", *Journal of Manufacturing Systems*, Vol. 77, pp. 798-809, doi: 10.1016/j.jmsy.2024.10.016.

Würth, T., Krauß, C., Zimmerling, C. and Kärger, L. (2023), "Physics-informed neural networks for data-free surrogate modelling and engineering optimization–an example from composite manufacturing", *Materials and Design*, Vol. 231, 112034, doi: 10.1016/j.matdes.2023.112034.

Zhang, J. and Zhao, X. (2023), "Digital twin of wind farms via physics-informed deep learning", *Energy Conversion and Management*, Vol. 293, 117507, doi: 10.1016/j.enconman.2023.117507.

Zhao, M., Wei, H., Mao, Y., Zhang, C., Liu, T. and Liao, W. (2023), "Predictions of additive manufacturing process parameters and molten pool dimensions with a physics-informed deep learning model", *Engineering*, Vol. 23, pp. 181-195, doi: 10.1016/j.eng.2022.09.015.

Zhou, C., Chen, J., Yang, Z. and Png, C.E. (2025), "Dual-balancing for physics-informed neural networks", arXiv preprint arXiv:2505.11117.

Zhu, Q., Liu, Z. and Yan, J. (2021), “Machine learning for metal additive manufacturing: predicting temperature and melt pool fluid dynamics using physics-informed neural networks”, *Computational Mechanics*, Vol. 67 No. 2, pp. 619-635, doi: 10.1007/s00466-020-01952-9.

Zobeiry, N. and Humfeld, K.D. (2021), “A physics-informed machine learning approach for solving heat transfer equation in advanced manufacturing and engineering applications”, *Engineering Applications of Artificial Intelligence*, Vol. 101, 104232, doi: 10.1016/j.engappai.2021.104232.

Further reading

Alba, E. and Martí, R. (2006), *Metaheuristic Procedures for Training Neural Networks*, Springer.

Bao, Y., Zhou, H. and Li, J. (2024), “Physics-based machine learning optimization of thermoelectric assembly for maximizing waste heat recovery”, *Energy*, Vol. 307, 132821, doi: 10.1016/j.energy.2024.132821.

Can, Ü. and Alataş, B. (2015), “Physics based metaheuristic algorithms for global optimization”, *American Journal of Information Science and Computer Engineering*, Vol. 1 No. 3, pp. 94-106.

Gnanasambandam, R. (2024), “Physics-informed machine learning for digital twins of metal additive manufacturing”.

Karkadakattil, A. (2026), “A physics-informed intelligent digital twin using multi-task CNN–LSTM for acoustic emission-based fault prognostics in safety-critical systems”, *Journal of Intelligent Manufacturing and Special Equipment*, pp. 1-26.

Khan, M.S., Jabeen, F., Ghouzali, S., Rehman, Z., Naz, S. and Abdul, W. (2021), “Metaheuristic algorithms in optimizing deep neural network model for software effort estimation”, *IEEE Access*, Vol. 9, pp. 60309-60327, doi: 10.1109/access.2021.3072380.

Kumar, S., Jangir, P., Tejani, G.G., Premkumar, M. and Alhelou, H.H. (2021), “MOPGO: a new physics-based multi-objective plasma generation optimizer for solving structural optimization problems”, *IEEE Access*, Vol. 9, pp. 84982-85016, doi: 10.1109/access.2021.3087739.

Naeini, H.K., Shomali, R., Pishahang, A., Hasanzadeh, H., Asadi, S. and Lonbar, A.G. (2025), “PINN-DT: optimizing energy consumption in smart building using hybrid physics-informed neural networks and digital twin framework with blockchain security”, arXiv preprint arXiv:2503.00331.

Qais, M.H., Hasanien, H.M., Alghuwainem, S. and Loo, K.H. (2023), “Propagation search algorithm: a physics-based optimizer for engineering applications”, *Mathematics*, Vol. 11 No. 20, 4224, doi: 10.3390/math11204224.

Qin, Y., Liu, H., Wang, Y. and Mao, Y. (2024), “Inverse physics–informed neural networks for digital twin–based bearing fault diagnosis under imbalanced samples”, *Knowledge-Based Systems*, Vol. 292, 111641, doi: 10.1016/j.knosys.2024.111641.

Shang, Z., Wang, Z., Pan, C., Li, W. and Gao, M. (2026), “Physics-informed auto-encoder based on digital twin for rolling bearing fault diagnosis under imbalanced sample conditions”, *Engineering Applications of Artificial Intelligence*, Vol. 163, 113017, doi: 10.1016/j.engappai.2025.113017.

Staff, G., Zarruk, G., Hatleskog, J., Stavland, S., McNulty, H., Ibarra, R., Calen, N., Hagemann, O., Lawrence, C., Selanger, T., Ayati, A. and Renstrøm, T. (2020), “Physics guided machine learning significantly improves outcomes for data-based production optimization”, *Abu Dhabi International Petroleum Exhibition and Conference*, SPE.

Wang, Z., Jahanshahi, M.R., Lund, A., Shahriar, A. and Montoya, A. (2025), “Physics-informed machine learning for hybrid digital twin–enhanced damage detection and localization”, *Journal of Engineering Mechanics*, Vol. 151 No. 12, 04025080, doi: 10.1061/jenmdt.emeng-8325.

Yang, X.-S. (2013), “Optimization and metaheuristic algorithms in engineering”, in *Metaheuristics in Water, Geotechnical and Transport Engineering*, p. 23.

Corresponding author
Aswin Karkadakattil can be contacted at: ashwinharik20000@gmail.com