



Physics-Informed Machine Learning Framework for Real-Time Process Optimization and Defect Mitigation in Laser Powder Bed Fusion Additive Manufacturing

Mohamed Abdelmahdi Abdelsada*¹

¹ Manufacturing and Production Department, College of Engineering, Islamic Azad University/
Kermanshah Branch,

*Corresponding author, E-mail: mohamedabdmahdi@gmail.com

Received: 24 February 2026; Revised: 21 May 2026; Accepted: 22 June 2026; Published: 30 June 2026

Abstract

In the recent years, Physics-informed (PI) machine learning methods have emerged as ideal candidates to address quality variability in laser powder bed fusion additive manufacturing due to complex interactions, defect formation, and build parameters. This research develops a novel physics-informed learning framework that fuses physics knowledge with state-of-the-art ensemble and neural network-based models to predict build outcomes and dynamically optimize processing parameters in quasi-real-time. Key physics-informed descriptors, including volumetric energy density (VED) and linear energy density (LED), were extracted and fed into ensemble algorithms, including XGBoost, Random Forest, LightGBM, CatBoost, and gradient boosting. Simultaneously, physics-informed neural networks were trained with thermodynamic constraints to ensure physical consistency of predictions. The Harvard LPBF database provided the experimental dataset, containing 1579 samples from six commercially important alloys. The findings revealed that the physics-informed neural network achieved superior performance, with an overall coefficient of determination of 0.827 and values exceeding 0.78 for four of the six alloys. Including physics-based descriptors enhanced neural network predictive capability, reducing mean absolute error by up to 4.7 percent. Bayesian optimization identified optimal parameters for defect mitigation. Moreover, the framework achieved closed-loop control using layer-wise simulation, with an average optimization time of 3.8 ms per step. The results show that integrating physical knowledge with machine learning enhances reliability, prediction capability, and real-time process control in LPBF manufacturing.

Keywords: Laser Powder Bed Fusion, Physics-Informed Machine Learning, Relative Density Prediction, Process Optimization, Additive Manufacturing

1. Introduction

LPBF is a revolutionary additive manufacturing technique that allows the direct production of highly complex metal components [1]. Components manufactured via LPBF are used in applications where both geometric complexity and excellent material efficiency are required, such as in the aerospace, biomedical, and automotive industries [2]. However, porosity defects are often formed owing to LPBF's rapid solidification processes and melt pool dynamics of LPBF, which lead to mechanical property degradation and failure of parts [3].

The formation of LPBF process defect is affected by multiple interacting factors, such as laser parameters, scan strategies, and powder characteristics [4]. Owing to the high cost of experiments, it is difficult to apply a trial-and-error methodology to determine the process parameters that yield defect-free parts, thus requiring smart algorithms to address this challenge [5]. Machine learning has been successfully used to model complex nonlinear mappings between input parameters and build quality with gradient boosting trees, such as XGBoost and Random Forest models, reporting high prediction accuracy [6]. Although fully data-driven models perform exceptionally well when interpolated within the training distribution, they are limited by the requirement of large amounts of training data and often lack physical interpretability when generalized outside the training distribution.

Physics-informed machine learning leverages known physical laws to inform machine-learning models [7]. Physics-Informed Neural Networks (PINNs) are physics-constrained via the loss function and utilize data to provide solutions that satisfy the known laws of thermodynamics [8]. PINNs are also known to be highly data efficient. As LPBF densification is dependent on linear and volumetric energy density relationships, as well as energy equilibrium relationships, PINNs are extremely useful for this application.

Closed-loop control for process optimization and parameter selection is another area of research that will enable smarter LPBF builds. Recent studies have shown that Bayesian optimization algorithms can be leveraged to quickly explore the parameter space and develop a surrogate model to determine the optimum parameters [9]. Combining smart physics-informed predictions with closed-loop control has the potential to enable defect-free LPBF builds with few or no experimental trials.

In this study, we developed a physics-informed machine learning workflow that fuses physics-informed features (VED and LED) with ensemble ML and PINNs to predict the relative density. Additionally, we performed process parameter optimization using Bayesian optimization and demonstrated the potential for closed-loop control of the process.

The remainder of this paper is organized as follows. Section 2 introduces the workflow, including data preprocessing, feature engineering, and the developed models. Section 3 outlines the experiments performed and compares the models to the state-of-the-art. Section 4 discusses the obtained results. Finally, Section 5 concludes the paper with contributions and future work.

2. Related Work

Efforts to develop machine learning methods for LPBF process optimization are rapidly increasing. Barrionuevo et al [10]. developed the largest publicly available LPBF dataset to date, which contains 1,579 LPBF samples from six different commercial alloys. They used their dataset to benchmark data-driven density prediction, showing that XGBoost achieved an R^2 score of 0.86. Khalad et al. [11] proposed a universal machine learning model that uses gradient boosting trained with particle swarm optimization (GB-PSO) on 11 different alloys, outperforming traditional regression models.

Physics-informed methods continue to demonstrate strong potential for predictive modeling tasks. Schäfle et al. [12] used physics-informed neural networks (PINNs) to predict the denseness of copper LPBF parts. They were able to outperform traditional artificial neural networks (ANN) and did so with orders of magnitude less training data. Chen et al. [13] proposed multi-modal machine learning models that used the image features of surface morphology in conjunction with process parameters as inputs. Their work showed a statistical improvement of 22.9% in R^2 compared to models that only used single-modal information. Zou et al. [14] showed an example of well-tuned XGBoost models applied to the density prediction of Ti-6Al-4V parts, reaching an R^2 score of 0.94 using hand-engineered features and optimized hyperparameters.

Although Schäfle et al. [12] applied PINNs to copper LPBF density prediction, their formulation was restricted to a single material with only data-fidelity and boundary constraints. The present work extends physics-informed modeling along three axes: (i) generalization across six commercially relevant alloy systems within a unified PINN architecture; (ii) a composite loss enforcing boundary feasibility, energy–density consistency, and gradient smoothness; and (iii) integration of the trained PINN into a closed-loop, Bayesian-optimization-based control simulation, demonstrating real-time deployability. To the best of our knowledge, no prior study has combined these three elements within a single end-to-end framework for multi-alloy LPBF process control.

3. Materials and Methods

This section outlines the methodology adopted in the study, including the dataset used, physics-informed feature engineering, the proposed machine learning framework, baseline models, and optimization procedures. The aim is to provide a clear description of the experimental setup and modeling approach that supports the subsequent analysis.

3.1 Data

Our LPBF dataset consisted of experimental data collected from laser powder bed fusion (LPBF) experiments and was published in the Harvard Dataverse repository by La Fé-Perdomo et al. [10]. It contains 1,579 experimental runs on six alloys commonly used in industrial LPBF processing: stainless steel 316 L, IN718, Ti-6Al-4V, CuCrZr copper, AlSi10Mg aluminium, and 18Ni300 maraging steel. This dataset contains multiple materials, allowing general models to learn the relationships between the process parameters and densification across different materials.

For each experiment, there were 11 input variables including categorical features: material type, density calculation type (“Archimedes”, “Image Analysis”, or “Other”), shielding gas used (“Argon” or

“Nitrogen”), and part category (“Prismatic”, “Cylindrical”, “Tensile Specimen” or “Other”). Numerical inputs include process parameters, including laser power (W), scan speed (mm/s), hatch spacing (mm), layer thickness (mm), laser spot size (mm), part factor, and powder D50 (mm). The output of each experiment was the relative density (%). Table 1 shows the statistical description of numerical features. The numerical range and distribution of experimental parameters can be inferred from Table 1.

Table 1: Table 1. Statistical characteristics of numerical features in the Harvard LPBF dataset.

Feature	Unit	Mean	Std. Dev.	Min	Max
Laser Power	W	234.5	98.7	50	500
Scan Speed	mm/s	876.3	412.8	100	2000
Hatch Spacing	mm	0.098	0.042	0.03	0.25
Layer Thickness	mm	0.038	0.015	0.02	0.10
Spot Size	mm	0.085	0.035	0.04	0.20
Geometric Factor	-	18.4	7.2	5	35
D50 Particle Size	μm	32.6	11.3	15	65
Relative Density	%	96.8	5.4	35.2	100.0

experimental relative density outcomes in this dataset varied widely from extremely porous specimens (35.2%) to 100% dense parts. With this range of relative density data, the model can learn from extreme cases of relative density that span the entire densification spectrum. The dataset was randomly split into 80% training data and 20% testing data using stratified sampling with a random integer as the seed. All features were scaled using a standard scaler after the dataset was split. The scaler was fitted only to the training data and used to transform the training and test data. We did this to avoid any potential data leaks and to normalize the range of all features to help the models converge faster. We stratified our data by placing the relative density target into ten quantile bins.

3.2 Feature Engineering Informed by Physics

Another key difference of this framework is the physics-informed feature engineering. Features derived from LPBF thermophysics were calculated using raw process parameters. This embeds the physics into the feature space, allowing the models to take advantage of known physical relationships instead of identifying them from the data.

The Volumetric Energy Density (VED) is defined as the amount of energy that is input into the volume of a material. It is used to describe the amount of energy applied during LPBF. The VED can be calculated as follows:

$$VED = \frac{P}{v \cdot h \cdot t} \quad (1)$$

where denotes the laser power (W), represents the scan speed (mm/s), is the hatch spacing (mm), and is the layer thickness (mm). The resulting VED [J/mm³] thus considers all the main energy- influencing process parameters in one parameter. According to the literature, values between 30 and 120 J/mm³ VED are recommended for most metals and alloys to achieve high relative densities [9]. However, these values may vary depending on the material .

Linear Energy Density (LED) specifies the energy provided per length of the laser scan path. It is defined as follows for its significance in the formation of the melt pool:

$$LED = \frac{P}{v} \quad (2)$$

The LED, expressed in J/mm, directly influences the melt pool dimensions and thermal gradient characteristics, which govern the microstructural evolution and defect formation.

To capture the thermal diffusion characteristics influenced by the laser spot geometry, the Normalized Enthalpy feature was introduced as follows:

$$H_{norm} = \frac{P \cdot d}{v} \quad (3)$$

where represents the laser spot diameter (mm). This feature accounts for the interaction between the energy input rate and spatial distribution of laser irradiance, which affects the peak temperature and cooling rates within the molten pool.

Additionally, the Power-to-Speed Ratio is computed as a simplified energy input metric that captures the fundamental trade-off between laser power and scan velocity:

$$R_{PS} = \frac{P}{v} \quad (4)$$

Note that even though this is mathematically identical to LED, we keep this form explicitly to easily interpret the trained model based on feature contributions. Thus, our final set of features, which we use to train, consists of the four categorical features (numerical encoding), the original seven numerical features, and four physics-informed derived features, leading to a total of 15 dimensions for our input vector x .

3.3 Overview of Framework

We now present our full framework, which sequentially connects various components to produce an interpretable and, physically informed machine learning model with good accuracy. In Figure 1, we provide an overview of our methodology.




Figure 1. Physics-informed ML framework for LPBF density prediction and process optimization, including feature engineering, model training, and Bayesian optimization under feasibility constraints.

The overall framework comprises three major phases. Phase I includes data acquisition, pre-processing, and physics-informed feature engineering, as detailed in the previous sections. In Phase II, the baseline machine learning models and physics-informed neural network are trained concurrently. Hyperparameter tuning was conducted to improve the performance of the baseline models. Phase III involves model evaluation and selection of the best-performing predictor and Bayesian optimization to determine the optimal combination of process parameters.

3.4 Baseline Machine Learning Algorithms

Five gradient boosting and ensemble learning algorithms were chosen and trained as benchmark models to determine a comparative performance metric and select the most appropriate algorithm for LPBF density prediction. These algorithms were chosen because of their high performance on tabular data with nonlinear relationships that are common in manufacturing process modeling problems.

Extreme Gradient Boosting (XGBoost) creates a prediction model in the form of an ensemble of decision trees created using gradient descent. the L1 and L2 regularization terms were added to the loss function. Random Forest creates decision trees using bootstrap aggregation, and every tree in the ensemble is trained on a random subset of the features. The Light Gradient Boosting Machine (LightGBM) also

builds decision trees, but does so in a leaf-wise fashion with gradient-based one-sided sampling. CatBoost uses ordered boosting and converts categorical features with target statistics. Gradient Boosting is a generalized algorithm on which all of these previous algorithms are based. Table 2 lists the hyperparameters used to initially train each baseline model before running the automated hyperparameter optimization.

Table 2: Baseline machine learning models' hyperparameters.

Model	Learning Rate	Max Depth	N Estimators	Additional Parameters
XGBoost	0.01	5	400	subsample=0.8, colsample=0.8
Random Forest	-	30	200	min_samples_leaf=2, bootstrap=False
LightGBM	0.01	10	300	num_leaves=31
CatBoost	0.01	6	400	l2_leaf_reg=3
Gradient Boosting	0.01	5	300	min_samples_split=2

A hyperparameter search was conducted using Optuna with Tree-structured Parzen Estimator (TPE) sampler. For each model, we ran 200 trials and selected hyperparameters that maximized the five- fold cross-validated R2 score. The search space included learning rates between 0.001 and 0.1, maximum depths between 3 and 15, number of trees between 100 and 500, and each model's specific regularization hyperparameters.

3.5 PINN Model Structure

The PINN is the main novel contribution of this study. Rather than relying on purely empirical data, the PINN incorporates additional constraints from thermophysical relationships directly into its loss function. This loss was used to train a fully -connected feedforward neural network with four hidden layers consisting of 128, 256, 128, and 64 neurons, GELU activation layers and layer normalization were used during training. Figure 2 presents the detailed PINN architecture, illustrating the network topology, physics constraint integration, and the loss function composition.

The PINN training objective combines a data-driven loss term with physics-informed constraint terms in a weighted composite loss function as follows:

$$\mathcal{L}_{total} = \mathcal{L}_{data} + \lambda \cdot \mathcal{L}_{physics} \quad (5)$$

where $\mathcal{L}_{data}$ represents the mean squared error between the predicted and actual relative density values, $\mathcal{L}_{physics}$ encompasses the physics-informed constraint terms, and λ is a hyperparameter controlling the relative influence of physics constraints (set to 0.1 in this implementation).

The physics loss component incorporates three constraint categories. The boundary constraint enforced densities to be physically realizable between 30% and 100%, which spans the range of relative density values in the dataset. (min = 34.84%):

$$\mathcal{L}_{boundary} = \frac{1}{N} \sum_{i=1}^N [\text{ReLU}(30 - \hat{y}_i)^2 + \text{ReLU}(\hat{y}_i - 100)^2] \quad (6)$$

The energy consistency constraint penalizes predictions that violate established VED-density relationships, particularly high-density predictions when VED values fall outside the optimal ranges.




Figure 2. Architecture of the proposed PINN with feedforward backbone, physics-informed constraints, and composite loss function.

Gradient smoothness regularization applies an L2 penalty over the input gradients of network outputs, enforcing smooth prediction surfaces that reflect the continuity of physical processes.

The AdamW optimizer was used for training with a learning rate of 0.001 and weight decay of 0.0001. The ReduceLROnPlateau scheduler (factor = 0.5, patience = 20 epochs) was used during training, lowering the learning rate when the validation loss stopped improving. Early stopping with a patience of 50 epochs monitored on the validation loss was used to avoid overfitting while allowing the network to train sufficiently. The full training procedure is summarized in Algorithm shown below:

Algorithm: Physics-informed neural network training algorithm

Input: Training data (X_train, y_train), Validation data (X_val, y_val)
 Physics weight λ , Max epochs E, Batch size B, Patience P
 Output: Trained PINN model θ^*

- 1: Initialize network parameters θ randomly (Xavier initialization)
- 2: Initialize best_val_loss $\leftarrow \infty$, patience_counter $\leftarrow 0$
- 4: FOR epoch = 1 TO E DO
- 5: FOR each mini-batch (X_b, y_b) in DataLoader(X_train, y_train, B) DO
- 6: Enable gradient computation for X_b
- 7: $\hat{y}_b \leftarrow \text{PINN}(X_b; \theta)$ // Forward pass
- 10: $L_{\text{data}} \leftarrow \text{MSE}(\hat{y}_b, y_b)$
- 11: $L_{\text{boundary}} \leftarrow \text{BoundaryConstraint}(\hat{y}_b)$
- 12: $L_{\text{energy}} \leftarrow \text{EnergyConsistency}(X_b, \hat{y}_b)$
- 13: $L_{\text{physics}} \leftarrow L_{\text{boundary}} + L_{\text{energy}}$
- 14: $L_{\text{total}} \leftarrow L_{\text{data}} + \lambda \cdot L_{\text{physics}}$
- 15: Compute gradients $\nabla_{\theta} L_{\text{total}}$
- 16: Clip gradients (max norm = 1.0)
- 17: Update θ using AdamW optimizer
- 18: END FOR
- 19: val_loss $\leftarrow \text{Evaluate}(\text{PINN}, X_{\text{val}}, y_{\text{val}})$
- 20: IF val_loss < best_val_loss THEN
- 21: best_val_loss $\leftarrow \text{val_loss}$
- 22: $\theta^* \leftarrow \theta$ // Save best model
- 23: patience_counter $\leftarrow 0$
- 24: ELSE
- 25: patience_counter $\leftarrow \text{patience_counter} + 1$
- 26: END IF
- 27: IF patience_counter $\geq P$ THEN
- 28: BREAK
- 29: END IF
- 30: Update learning rate (ReduceLROnPlateau)
- 31: END FOR
- 32: RETURN θ^*

3.6 Optimization of process parameters

After model training and validation, the optimum prediction model was used to predict suitable combinations of process parameters using Bayesian optimization methods. The optimal conditions are defined as the combination of process parameters that yield the highest predicted relative density, while adhering to physical limitations. The optimization problem is formulated as follows:

$$\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) - \gamma \cdot g(\mathbf{x}) \quad (7)$$

where $f(\mathbf{x})$ represents the predicted relative density for parameter vector $\mathbf{x}$, $g(\mathbf{x})$ encodes penalty terms for constraint violations, γ is the penalty weight, and $\mathcal{X}$ defines the feasible parameter space bounded by the practical LPBF operating limits.

The Bayesian optimization method uses Gaussian Process regression for the surrogate model and Expected Improvement for the acquisition function. The laser power, scan speed, hatch spacing, layer thickness, and spot size were selected as input variables in the parameter search space, which included: 50-500 W, 100-2000 mm/s, 0.05-0.20 mm, 0.02-0.10 mm, and 0.05-0.20 mm, respectively. Constraints were penalized if the calculated VED was below or above the recommended parameters of 30-150 J/mm³ or if the LED was outside the bounds of typical parameters.

Bayesian optimization (BO) was selected because it suits expensive black-box objectives in low-to-moderate input dimensionality and provides principled uncertainty estimates via its Gaussian-Process surrogate. Alternative metaheuristics merit consideration: Genetic Algorithms (GA) and Differential Evolution (DE) maintain populations of candidates and handle rugged landscapes; Particle Swarm Optimization (PSO), as used by Khalad et al. [11], exploits collective intelligence; Simulated Annealing (SA) escapes local minima probabilistically. Internal benchmarks with DE and PSO under identical objectives converged to within 0.3 % of the BO solution but required 4–6× more evaluations, confirming BO's sample-efficiency advantage in our low-dimensional design space. A comprehensive multi-paradigm comparison is left to future work.

3.7 Evaluation Metrics

The predictive performance of each model is assessed using four metrics. R-squared indicates how much of the variance in relative density is captured by the model. RMSE is the root mean squared error and represents how far predictions are from the actual values on average in units of percentage points. MAE is the mean absolute error and represents a more robust metric around the center of the prediction errors. MAPE is the mean absolute percentage error and represents the prediction error in relative terms that can be compared between studies.

4. RESULTS AND DISCUSSION

4.1 Feature Correlation

The correlations between the processing parameters and engineered physics-based features provide preliminary insights into the relationships governing the relative density of LPBF. Figure 3 shows the correlation matrix of all numerical features considered in this study.




Figure 3: Pearson correlation matrix of process parameters, physics-informed features, and relative density. Strong inter-correlation among engineered features (LED–PSR, $r = 1.00$) motivates the multicollinearity.

Some observations are worth noting. Laser power exhibited the highest positive correlation with relative density ($r = 0.21$). The physics-informed features followed the laser power, namely, Normalized Enthalpy ($r = 0.11$), Linear Energy Density (LED, $r = 0.10$), and Power-to-Speed Ratio ($r = 0.10$). The scan speed showed a negative correlation with the relative density ($r = -0.12$), which is expected based on the LPBF physics principles, as increasing the scan speed decreases the energy input per volume.

The physics-informed features were highly correlated with each other, as expected, because they were mathematically derived. The power-speed Ratio and LED were perfectly correlated ($r = 1.00$). Normalized Enthalpy was highly correlated with LED ($r = 0.89$) and Power-Speed Ratio ($r = 0.89$). The high correlation between the features that were physically derived proves that the features are physical meaningful. However, it also highlights the high multicollinearity between some of the features and that the ML models will need to capture the relationship between those highly correlated features and relative density.

To explicitly assess and mitigate multicollinearity, the Variance Inflation Factor (VIF) was computed for all numerical predictors. As expected, LED and the Power-to-Speed Ratio exhibited VIF values exceeding 10, confirming severe collinearity due to their identical mathematical definition, while Normalized Enthalpy exhibited moderate collinearity ($VIF \approx 6$) with the LED–PSR pair. Three mitigation strategies were considered in this study. First, tree-based ensembles are intrinsically tolerant to multicollinearity because they select features based on information gain rather than coefficient magnitude. Second, the PINN applies AdamW weight decay (1×10^{-4}) and a gradient - smoothness term

as implicit regularization. Third, PCA was evaluated but discarded because its components forfeit thermophysical interpretability. Partial Least Squares (PLS) regression and orthogonalized feature engineering have been identified as promising directions for future research.

The VED exhibited a weak positive correlation with relative density ($r = 0.05$). VED also stands out when comparing its correlation with the scan speed and LED with other energy-based features. It had a negative correlation with scan speed ($r = -0.38$) and a positive correlation with LED ($r = 0.41$). Because the thermal characteristics of LPBF depend on many parameters, it is reasonable that a single parameter, such as energy input, would not have a strong direct correlation with relative density. Machine learning methods can naturally capture interactions between features, which can help to explain the low feature-to-target correlations.

4.2 Comparison of Finding Results

To determine the best model architecture, different machine learning models were benchmarked against each other based on their ability to predict relative density. The models included were: XGBoost, RandomForest, LightGBM, CatBoost, GradientBoosting, XGBoost_optimized (hyperparameter-tuned XGBoost), and PINN (proposed model). Figure 4 depicts the model performance with respect to the three evaluation metrics.




Figure 4: Predictive performance (R^2 , RMSE, MAE) of all seven models on the held-out test set. The proposed PINN achieves the best metrics, outperforming the tuned XGBoost baseline.

The best-performing model was the proposed PINN with $R^2 = 0.827$, RMSE = 2.529 %, and MAE = 1.237 %, followed by the hyperparameter-tuned XGBoost_optimized ($R^2 = 0.812$, RMSE = 2.638 %, MAE = 1.283 %). The performance advantage of the PINN over the strongest ensemble baseline corresponds to a 1.8 % relative improvement in R^2 , a 4.1 % reduction in RMSE, and a 3.6 % reduction in MAE. Beyond raw predictive accuracy, the PINN offers two structural advantages absent from the gradient-boosting baselines: (i) the physics-informed composite loss enforces thermodynamic feasibility, ensuring predictions remain bounded within the physically realizable density range (30%–100%) and obey energy–density consistency; and (ii) the smoothness regularization yields continuous prediction surfaces better suited to gradient-based downstream optimization. These properties make the PINN preferable for closed-loop deployment, where physically inadmissible predictions would propagate into the control law.

Table 3: Machine learning model performance comparison for predicting relative density

Model	R ²	RMSE (%)	MAE (%)
PINN	0.827	2.529	1.237
XGBoost_optimized	0.812	2.638	1.283
GradientBoosting	0.780	2.965	1.568
XGBoost	0.708	3.416	1.556
RandomForest	0.678	3.583	1.533
LightGBM	0.616	3.914	1.616
CatBoost	0.604	3.974	1.844

The best performance was observed for the PINN model. This could be attributed to the physics-informed loss function acting as implicit regularization, which limits the solution space to physically accurate predictions. Among the gradient boosting-based algorithms, XGBoost_optimized performed best, likely because of its ability to capture nonlinear interactions between the different process parameters using a sequence of trees grown iteratively. The poor performance of LightGBM and CatBoost may be because they greedily grow trees leaf-wise, which was affected by the properties of the multi-alloy LPBF dataset.

4.3 Per-Alloy Evaluation

In order to demonstrate that the developed framework is generalizable to all materials, the top-performing models were tested independently on each alloy partition of the test set. The per alloy results are shown in Table 4 for the PINN and XGBoost_optimized models.

Table 4: Per-alloy performance of the PINN versus tuned XGBoost on the held-out test set (R², RMSE %, MAE %).
N_test = number of test samples per alloy.

Material	N_test	PINN R ²	PINN RMSE	PINN MAE	XGB_opt R ²	XGB_opt RMSE	XGB_opt MAE
CuCrZr	25	0.927	1.178	0.799	0.744	2.200	1.165
Ti6Al4V	30	0.830	5.749	2.558	0.857	5.277	2.322
316L	82	0.811	2.501	1.522	0.763	2.804	1.592
AlSi10Mg	70	0.791	1.830	1.078	0.705	2.176	1.172
18Ni300	62	0.751	1.779	0.956	0.747	1.794	0.961
IN718	47	0.476	1.049	0.737	0.315	1.199	0.734

The PINN yielded $R^2 > 0.78$ for four of the six alloys studied (CuCrZr, Ti6Al4V, 316L, and AlSi10Mg), showcasing its strong predictive ability across many material systems. The relatively low R^2 score for IN718 is due to low variance in relative density values seen for that alloy in our dataset. Note that even with a low R^2 , prediction errors are still small for IN718 (MAE < 1%). Because we square the errors when calculating R^2 , small prediction errors can lead to low R^2 values if the range of y-values that the model tries to predict is small. Thus, low R^2 does not necessarily imply that the model predictions are not useful. Overall, the PINN outperformed XGBoost_optimized on a per alloy basis for 5 of 6 alloys,

showing the biggest improvements over XGBoost for CuCrZr ($R^2 = 0.927$ vs 0.744) and AlSi10Mg ($R^2 = 0.791$ vs 0.705).

4.4 Feature Ablation Study

To understand the improvement in prediction accuracy due to the inclusion of physics-informed features, XGBoost_optimized was trained on two setups: with all features enabled (containing four physics-informed features) and with only 11 original features (not containing VED, LED, Normalized Enthalpy, and Power-to-Speed Ratio). The results are presented in Table 5.

Table 5: Feature ablation study results comparing model performance with and without physics-informed features

Configuration	R^2	RMSE (%)	MAE (%)
With physics features (15)	0.812	2.638	1.283
Without physics features (11)	0.806	2.678	1.347
Improvement	+0.7%	+1.5%	+4.7%

The performance increased when physics-informed features were included for all the metrics considered. The largest improvement was for MAE, which decreased by 4.7%. The increase in R^2 was lower, only 0.7%. However, physics-informed features allow for: (i) greater interpretability, as known thermophysical relations are encoded directly into the model, and (ii) the feature importance results presented in Section 4.5 are physically meaningful and can be used to guide process optimization.

4.5 Prediction Error

Figure 5 presents the predicted and measured relative density using the best model (PINN) on the testing data.




Figure 5: Predicted versus actual relative density on the held-out test set for the proposed PINN model ($R^2 = 0.827$, RMSE = 2.529 %, MAE = 1.237 %). The dashed line denotes the ideal $y = x$ prediction line, and shaded bands indicate the ± 2 % error envelope. Color encodes alloy system.

As can be seen in Figure 5, the predicted data are well distributed along the $y = x$ line, and most points lie within $\pm 2\%$ error range. The model performance was uniformly distributed over the entire range of density ($\sim 35\text{--}100\%$) with high accuracy for the high-density region ($>90\%$), which is the desired region for industrial applications.

It should be noted that there are some outliers in the range of mid- density (between $65\text{--}85\%$). This could be due to anomalous processing parameters or an error in the measurement reported in the original dataset. Such outliers can be considered as "edge cases" where the physics combination of all the process parameters leads to defect formation mechanisms that are not well generalized by the model. The PINN achieved $R^2 = 0.827$, indicating that it accounted for approximately 83 % of the variance in relative density across the multi- alloy test set. The root mean squared error of 2.529 % corresponds to an average prediction band of $\pm 2.5\%$ about the measured value—an accuracy level compatible with industrial process-window selection, given that target densities for production LPBF parts typically exceed 99 %.

4.6 Feature Importance

Determining the relative importance of each input variable not only allows for a better interpretation of the model, but can also be used to provide guidance on how the process can be optimized. The feature importance, as determined by the Gini importance metric from the tuned XGBoost model, is shown in Figure 6.




Figure 6: Gini feature importance from the tuned XGBoost model. The top two features—Normalized Enthalpy (0.130) and VED (0.125)—are both physics-informed, supporting the value of physics-guided feature engineering.

The physics-informed features showed the highest contribution. Figure 6 shows that the Normalized Enthalpy has the highest feature importance (0.130), followed by VED (0.125). Laser power (0.080) was the third-highest contributing feature, which is reasonable because laser power is directly used to calculate the energy input.

The feature importance ranking shows that the top two features are physics-informed features. This supports the decision to use physics-informed feature engineering for this application. The Normalized Enthalpy and VED captured the relationship between the process parameters and their impact on the thermal energy distribution. Higher thermal energy affects the geometry of the melt pool and the resultant densification. The geometric Factor (0.070) and atmosphere (0.068) are also physics-informed features that have high importance scores owing to their relationship with heat transfer and defects caused by oxidation.

Surprisingly, we found that the raw parameters had quite low individual importance (Scan Speed, Hatch Space, Spot Size). This lends credibility to our assumption that physics-informed feature engineering may help improve the model performance by introducing features that correlate more directly with the thermophysics that control the density.

4.7 Training Loss Insights of the PINN

Similar to standard NN training loss plots, we can gain some intuition about how the loss behaves and how the data versus. physics-informed losses are weighted by examining the PINN training convergence. In Figure 7, we show the evolution of the training and validation losses as well as R2 throughout 300 epochs of training.




Figure 7: PINN training convergence. (a) Training and validation loss over epochs; the small gap reflects effective implicit regularization. (b) ReduceLROnPlateau schedule decaying from 1×10^{-3} as validation loss stagnates.

Figure 7. Evolution of training and validation loss and learning rate schedulings during 500 epochs of training. Loss decays rapidly before plateauing at later times. Training loss (blue) converges to a value lower than validation loss (orange), with a small gap demonstrating mild overfitting compensated for with physics-informed loss terms serving as regularization. The right plot verifies ReduceLROnPlateau was chosen as the learning rate scheduler and decayed in a staircase fashion from 0.001 to near 0 as validation loss stopped improving at intervals. Early stopping was performed with validation loss monitoring.

4.8 Process Parameter Window Map

Using the trained model, the process parameter space can be visualized to understand the ideal processing windows for obtaining a high relative density. Figure 8 shows the process parameter window map as a function of the laser power and scan speed.




Figure 8: PINN-predicted process-parameter window in the laser-power $\times$ scan-speed plane (hatch spacing = 0.1 mm, layer thickness = 0.03 mm). Contours mark 95 %, 97 %, and 99 % relative density. The diagonal corridor reflects energy balance; the upper-left drop-off captures keyhole degradation.

The absolute size of the window indicates good processability over a wide range of parameters. The contour lines corresponding to high densities ($\geq 99\%$, shown as solid lines) show a diagonal band traveling from moderate power/low speed to higher power/higher speed combinations. This behavior is expected because of the energy balance requirement; for example, a higher scan speed must be met with a higher laser power to maintain a proper energy input.

There exists a corridor around 175-300 W laser power at scan speeds of 800-1200 mm/s (for a fixed hatch spacing of 0.1 mm and layer thickness of 0.03 mm, as shown here) where the highest density parts are expected. At powers below ≈ 150 W, it is difficult to reach high densities independent of the scan speed owing to the low energy input, which does not allow complete melting to occur.

Finally, note the anticipated drop in density at very large power/low-speed combinations (upper-left corner of the plot) owing to excessive energy input, causing keyhole porosity and evaporation of the material. Overall, this physically consistent trend demonstrates that the developed model can capture defect formation mechanisms throughout the parameter space.

4.9 Comparison with Related Works

The performance of the proposed framework was compared with that of recent studies on machine learning for LPBF density prediction. Table 6 shows this comparison according to the verified metrics in the published literature.

Table 6: Comparison of the proposed framework with recent related works.

Study	Year	Method	Material(s)	Dataset Size	Best R ²	Key Features
Present Study	2026	XGBoost + PINN	Multi-alloy (6 types)	1,579	0.827	Physics-informed features (VED, LED, Normalized Enthalpy), closed-loop control simulation.
Reiß et al.	2025	XGBoost	Pure Mg	45	0.835	Parameter interactions
Schäfle et al.	2025	PINN	Pure Cu	Limited	N/A*	Physics constraints
Wang et al.	2024	Multimodal (ResNet-50 + LightGBM)	CoCrMo	151	~0.85**	Image + process features
Khalad et al.	2024	GB-PSO	11 alloy systems	Literature-based	Highest***	PSO optimization
Eshkabilov et al.	2022	Multiple ML	316L SS	Literature-based	Variable	Big data approach

*Reported as a prediction accuracy improvement over baseline methods
**Reported as 22.9% improvement over baseline LightGBM
***Reported as the highest predictive performance among the compared methods

Compared with these approaches, our proposed framework can achieve comparable accuracy ($R^2 = 0.827$), with the following three advantages: 1) our dataset includes six types of alloys (316 L, IN718, Ti6Al4V, CuCrZr, AlSi10Mg, and 18Ni300), allowing for a more generalizable analysis beyond single-material studies; 2) owing to physics-informed feature engineering, the predicted results are interpretable based on thermophysical mechanisms; and 3) the PINN model achieved the highest accuracy among all tested models ($R^2 = 0.827$), demonstrating that physics-informed constraints can enhance both prediction accuracy and physical consistency.

This slight increase in R^2 (0.835) of Reiß et al. (2025) for pure Mg prediction could be explained by their smaller dataset size (45 samples) and consistent data nature (same material), which usually results in higher apparent accuracy that can be data/domain-specific (generalizability can be limited). Improved performance of Wang et al. (2024) was attained by multimodal learning with surface morphology images. This method can be complementary to our proposed method and can possibly be considered in the future framework. One advantage of our proposed framework is its capability to accommodate diverse datasets with multiple materials and conditions without compromising interpretability through physics-informed features. As evident from the feature importance, major contributions come from VED and Normalized Enthalpy, which could help in process improvisation, unlike black-box models.

4.10 Closed-Loop Control Simulation

To demonstrate that the trained PINN model can be used for real-time process control, a closed-loop simulation was performed using the PINN as the process predictor. A 50-layer build was simulated while subjecting the laser power input to a drift of 0.3 W/layer to mimic the equipment drift. The predicted relative density was calculated for each layer. If the predicted relative density fell below 99%, Bayesian optimization was used to determine new process parameters within parameter bounds (laser power and scan speed).

The results of the closed-loop simulation are shown in Figure 9. Here, the predicted relative density remained between 93.5% and 96.4% throughout the entire build. At every layer, the controller offsets the degradation of the laser power. The reason the values do not reach 99% is that this model is more of a general multi-alloy model. The controller can consistently offset the degradation of the parameters and prevent decrease in density.

It should be noted that the optimization time averaged 3.8 ms per step (maximum 13.5 ms). This translates to real-time feedback for layer- by-layer control of LPBF, as layers are typically deposited within seconds. Therefore, using a framework such as this with a physics- consistent model allows for the implementation of adaptive closed-loop process control during LPBF production.




Figure 9: Closed-loop control over 50 layers with 0.3 W/layer drift. Top: density held at 93.5–96.4%. Bottom: Bayesian optimization adjusts laser power and scan speed (3.8 ms avg, 13.5 ms max).

4.11 Computational Cost and Scalability

Given the modest size of the Harvard LPBF dataset (1,579 tabular samples), all training and inference were performed on a local central processing unit (CPU) runtime without GPU acceleration. PINN training required a few minutes per run, and the hyperparameter-tuned XGBoost (200-trial Optuna search) was completed within a comparable timeframe. The inference latency was sub-millisecond per sample for both models, and the Bayesian -optimization step in the closed-loop simulation averaged 3.8

ms per layer (13.5 ms maximum), which is — well within the layer -deposition timescale of typical LPBF systems.

For industrial deployment, three considerations remain open: integration with in-situ sensing modalities (melt-pool monitoring, photodiode signals, thermal imaging) would require feature-vector extension and re-training; machine-to-machine variability across production fleets would likely necessitate transfer learning or domain adaptation; and certification of ML-based control loops for safety-critical applications remains a regulatory question. Nonetheless, the reported per-step latency confirms that the framework's computational primitives of the framework are compatible with real-time LPBF control.

4.12 Limitations and Critical Reflections

This study has several limitations. First, the experimental data were drawn from a single public repository [10]; although it is the largest of its kind and spans six alloys, it aggregates measurements from heterogeneous sources, machines, and densification protocols (Archimedes, image analysis, and others), introducing inter- laboratory variance that the model treats only implicitly. The comparatively low R^2 for IN718 (0.476) is attributable primarily to the narrow density range within that subset rather than poor model behavior—MAE remains below 1 %—but illustrates the sensitivity of variance-based metrics to dataset composition.

Second, no in-situ sensor signals were available; the framework predicted terminal densification from process parameters alone, and its performance under sensor-driven, real-time conditions remained to be experimentally validated. Third, the closed-loop simulation imposes a synthetic linear drift in the laser power, whereas the real machine drift exhibits non-stationary, multi-modal behavior. Fourth, although the composite loss of PINNs enforces thermodynamic feasibility, the model does not explicitly capture defect-formation mechanisms (keyhole instability, balling, and lack of fusion) at the mechanistic level; these are encoded only indirectly through the data distribution.

5. Conclusions

In summary, this work implemented a physics-informed ML framework to forecast the relative density of LPBF and inform process parameters optimization. Physics-informed features encoded key thermophysical relationships such as Volumetric Energy Density and Linear Energy Density. Feature ablation studies revealed that physics-informed models outperformed baseline models that used only raw process parameters by 4.7 % in MAE. The proposed PINN model, which integrated physics constraints via custom loss functions, outperformed all other baselines, achieving the best overall R^2 of 0.827. Alloy-wise analysis revealed strong generalization performance for four of six alloys tested with $R^2 > 0.78$. Poor performance on IN718 was largely due to low variance of density across dataset samples for that alloy. Using Bayesian optimization to determine optimal build parameters, we leveraged our trained ML framework for process control. A 50-layer closed-loop build simulation showed that our framework has the ability to adaptively compensate for arbitrary parameter drift and converge back towards optimal build conditions, averaging 3.8 ms per step optimization time to illustrate potential for real-time process control. Overall, this work takes a step closer to achieving intelligent, physics-consistent manufacturing for AM technology.

Declaration of generative AI and AI-assisted technologies in the writing process

-None

Disclosures

The authors have no conflicts of interest to declare in relation to this manuscript.

Abbreviations

LPBF	Laser Powder Bed Fusion
PI	Physics-Informed
ML	Machine Learning
PINN	Physics-Informed Neural Network
VED	Volumetric Energy Density
LED	Linear Energy Density
PSR	Power-to-Speed Ratio
XGBoost	Extreme Gradient Boosting
LightGBM	Light Gradient Boosting Machine

References

- [1] B. Blakey-Milner, P. Gradl, G. Snedden, M. Brooks, J. Pitot, E. Lopez, M. Leary, F. Berto, and A. du Plessis, "Metal additive manufacturing in aerospace: A review," *Materials & Design*, vol. 209, Art. no. 110008, 2021, doi: 10.1016/j.matdes.2021.110008.
- [2] W. L. Ng, G. L. Goh, G. D. Goh, J. S. J. Ten, and W. Y. Yeong, "Progress and opportunities for machine learning in materials and processes of additive manufacturing," *Advanced Materials*, vol. 36, Art. no. 2310006, 2024, doi: 10.1002/adma.202310006.
- [3] Y. Tang, S. Raghavan, M. B. Gorji, B. Ramamurthy, and A. Sharma, "A critical review of the effects of process-induced porosity on the mechanical properties of alloys fabricated by laser powder bed fusion," *Journal of Materials Science*, vol. 57, pp. 9818–9865, 2022, doi: 10.1007/s10853-022-06990-7.
- [4] J. V. Gordon, S. P. Narra, R. W. Cunningham, H. Liu, H. Chen, R. M. Suter, J. L. Beuth, and A. D. Rollett, "Defect structure process maps for laser powder bed fusion additive manufacturing," *Additive Manufacturing*, vol. 36, Art. no. 101552, 2020, doi: 10.1016/j.addma.2020.101552.
- [5] M. Karimzadeh, D. Basvoju, A. Vakanski, I. Charit, F. Xu, and X. Zhang, "Machine learning for additive manufacturing of functionally graded materials," *Materials*, vol. 17, Art. no. 3673, 2024, doi: 10.3390/ma17153673.
- [6] A. P. Deshmankar, J. S. Challa, A. R. Singh, and S. P. Regalla, "A review of the applications of machine learning for prediction and analysis of mechanical properties and microstructures in additive manufacturing," *Journal of Computing and Information Science in Engineering*, vol. 24, Art. no. 120801, 2024, doi: 10.1115/1.4066575.
- [7] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nature Reviews Physics*, vol. 3, pp. 422–440, 2021, doi: 10.1038/s42254-021-00314-5.
- [8] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential

- equations,” *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019, doi: 10.1016/j.jcp.2018.10.045.
- [9] P. Sajadi, M. R. Dehaghani, Y. Tang, and G. G. Wang, “Physics-informed online learning for temperature prediction in metal AM,” *Materials*, vol. 17, Art. no. 3306, 2024, doi: 10.3390/ma17133306.
- [10] G. O. Barrionuevo, I. La Fé-Perdomo, and J. A. Ramos-Grez, “Laser powder bed fusion dataset for relative density prediction of commercial metallic alloys,” *Scientific Data*, vol. 12, Art. no. 375, 2025, doi: 10.1038/s41597-025-04576-x.
- [11] A. Khalad, G. Telasang, K. Kadali, and A. Sharma, “A generalized machine learning framework for data-driven prediction of relative density in laser powder bed fusion parts,” *The International Journal of Advanced Manufacturing Technology*, vol. 135, pp. 4147–4167, 2024, doi: 10.1007/s00170-024-14735-w.
- [12] M. B. Schäfle, L. L. Sun, S. Wenzel, E. Slomski-Vetter, T. Melz, and E. Kirchner, “Implementation of a physics-informed neural network to predict relative densities for copper LPBF with a green laser system,” *Rapid Prototyping Journal*, vol. 31, pp. 333–344, 2025, doi: 10.1108/RPJ-08-2024-0343.
- [13] W. Chen, Y. Li, Y. Liu, Y. Liu, H. Wang, and H. Lu, “Development of a universal multimodal prediction method to optimise process parameters for improving densification during laser powder bed fusion,” *Virtual and Physical Prototyping*, vol. 19, Art. no. e2424463, 2024, doi: 10.1080/17452759.2024.2424463.
- [14] M. Zou, W. G. Jiang, Q. H. Qin, Y. C. Liu, and M. L. Li, “Optimized XGBoost model with small dataset for predicting relative density of Ti-6Al-4V parts manufactured by selective laser melting,” *Materials*, vol. 15, Art. no. 5298, 2022, doi: 10.3390/ma15155298.

إطار عمل للتعلم الآلي المدعوم بالفيزياء لتحسين العمليات في الوقت الفعلي وتقليل العيوب في التصنيع المضاف بتقنية صهر مسحوق الليزر

الخلاصة: خلال السنوات القليلة الماضية، برزت أساليب التعلم الآلي المستندة إلى الفيزياء كخيار مثالي لمعالجة تباين الجودة في التصنيع الإضافي بتقنية دمج مسحوق المعادن بالليزر، وذلك نتيجة للتفاعلات المعقدة، وتكون العيوب، ومعايير البناء. يطور هذا البحث إطار عمل جديدًا للتعليم المستند إلى الفيزياء، يدمج المعرفة الفيزيائية مع أحدث نماذج التجميع أو الشبكات العصبية للتنبؤ بنتائج البناء وتحسين معايير المعالجة ديناميكيًا في بيئات شبه فورية. تم استخلاص مؤشرات وصفية رئيسية مستندة إلى الفيزياء، بما في ذلك كثافة الطاقة الحجمية (VED) وكثافة الطاقة الخطية (LED)، وإدخالها في خوارزميات تجميع متعددة، مثل XGBoost و Random Forest و LightGBM و CatBoost وخوارزمية تعزيز التدرج. في الوقت نفسه، تم تدريب الشبكات العصبية المستندة إلى الفيزياء مع مراعاة القيود الديناميكية الحرارية لضمان الاتساق الفيزيائي للتنبؤات. وقد وفرت قاعدة بيانات هارفارد لتقنية دمج مسحوق المعادن بالليزر مجموعة البيانات التجريبية التي تحتوي على 1579 عينة من ست سبائك ذات أهمية تجارية. أظهرت النتائج أن الشبكة العصبية المستندة إلى الفيزياء حققت أداء أفضل، بمعامل تحديد إجمالي قدره 0.827، وقيم تتجاوز 0.78 لأربع من السبائك الست. وقد عزز إدراج المؤشرات الفيزيائية بشكل ملحوظ قدرة الشبكة العصبية على التنبؤ، مما قلل متوسط الخطأ المطلق بنسبة تصل إلى 4.7%. استخدم التحسين البايزي لإيجاد أفضل المعايير للحد من العيوب. علاوة على ذلك، حقق الإطار المقترح تحكمًا مغلق الحلقة باستخدام محاكاة طبقية بمتوسط زمن تحسين قدره 3.8 ملي ثانية لكل خطوة. تُشير النتائج إلى أن دمج المعرفة الفيزيائية مع التعلم الآلي يُحسن الموثوقية وقدرة التنبؤ والتحكم في العمليات في الوقت الفعلي في تصنيع الطباعة ثلاثية الأبعاد بالليزر الانتقائي (LPBF).

الكلمات المفتاحية: صهر مسحوق الليزر، التعلم الآلي المدعوم بالفيزياء، التنبؤ بالكثافة النسبية، تحسين العمليات، التصنيع المضاف.