Key Engineering Materials
ISSN: 1662-9795, Vol. 1047, pp 73-82

Submitted: 2025-12-18
Revised: 2026-02-16
Accepted: 2026-02-25
Online: 2026-04-13

A Physics Informed Machine Learning Framework towards Density Prediction of Additively Manufactured Components

Abhilash Puthanveettil Madathil[1,a*], Olga Bylya[1,b] and Andrew Sherlock[1,c]

[1]Applied Industrial Intelligence (AI[2]), National Manufacturing Institute Scotland, University of Strathclyde, 3 Netherton Sq, Paisley, Renfrew, Glasgow, UK

[a*]abhilash.p-m@strath.ac.uk, [b]olga.bylya@strath.ac.uk, [c]a.sherlock@strath.ac.uk

*corresponding author

Keywords: laser powder bed fusion, additive manufacturing, physics-informed ML, density modelling.

Abstract. Laser powder bed fusion (LPBF) has become a key manufacturing route for high-value components, yet accurate prediction of part quality, particularly relative density, remains challenging due to complex interactions among process parameters, material properties, melt-pool physics and shielding environment. Traditional physics-based simulations offer mechanistic insight but suffer from model-form uncertainty and computational cost while purely data-driven machine learning (ML) models often lack interpretability, physical consistency and transferability across build conditions. To address these gaps, we propose a physics-informed machine learning (PIML) framework that integrates structured domain knowledge with symbolic regression to derive compact, interpretable analytical expressions for LPBF density. The framework constructs a knowledge base (KB) comprising dimensionless and normalized physics-informed process descriptors (PIPDs) that encode energy input, thermal diffusion and melt-track geometry; these descriptors form a physically consistent feature space for learning. The framework also serves as a foundation for the proposed future community-driven, knowledge-graph-based modelling of LPBF processes. The capability of the framework is demonstrated by modelling the relative density of additively manufactured maraging steel and evaluating cross-atmosphere transferability from Argon-shielded to Nitrogen-shielded builds. The resulting symbolic models provide a transparent, extensible and physically meaningful alternative to black-box ML, achieving high predictive performance under the training regime ($R = 0.964$) and strong generalization across printing atmospheres ($R = 0.896$).

1. Introduction

Additive manufacturing (AM), and laser powder bed fusion (LPBF) in particular, is now used for safety-critical applications including aerospace, energy and biomedical sectors, where properties such as relative density, porosity and defect morphology directly govern fatigue life and reliability. Yet, despite extensive process development, accurate and efficient prediction of the 3D printed part properties, especially the density of laser powder-bed fusion (LPBF) built components remains challenging because outcomes depend on complex interactions among process parameters, machine settings, and material properties.

A wide range of approaches have been explored for modelling LPBF. Physics-based models including finite-element heat transfer, thermo-mechanical solvers and multi-physics computational fluid dynamics (CFD) provide valuable mechanistic insight and can resolve melt-pool dimensions, temperature fields and stress evolution (DebRoy et al., 2018; Sames et al., 2016). However, these models remain incomplete and often inaccurate for practical process–property prediction, as the underlying thermo-fluid–powder–metallurgical interactions are not fully understood or computationally tractable at industrial scales (Tang et al., 2018). Their throughput is limited because thousands of fully transient 3D simulations would be required to span realistic process windows (Bruna-Rosso et al., 2018). Model-form uncertainty including simplified boundary conditions, homogenized powder beds and uncertain absorptivity, introduces further deviation from experimental

behavior (Zakirov et al., 2020). Moreover, these simulations provide only indirect links to porosity formation and still rely on empirical rules to detect lack-of-fusion or keyhole defects.

These limitations have motivated the adoption of data-driven machine learning (ML), which has shown promise in modelling relationships between process parameters and properties such as density, porosity and microstructure (Chen et al., 2024). Supervised regression models, including gradient boosting, neural networks, Gaussian processes and random forests, have successfully correlated scalar parameters with density or porosity (Tapia & Elwany, 2015; Tapia et al., 2016; Garcia-Moreno et al., 2020). Recently, Khalad et al. (2024) proposed a generalised ML framework with gradient boosting with particle swarm optimisation (GB-PSO) for prediction, and Shapley additive explanations (SHAP) analysis to rank the importance of volumetric energy density and spot size. Similarly, Chen et al. (2025) developed a classification model that identifies high-density regimes during metal additive manufacturing.

While these methods are more accurate and computationally efficient, they tend to learn correlations rather than physics, with no built-in guarantee of physical plausibility or reliable extrapolation beyond the training regime (Maitra et al., 2024). Even with SHAP or local interpretable model-agnostic explanations (LIME) tools applied for post-hoc explanations (Khalad et al., 2024), their performance often degrades under domain shift, for example, when transitioning across alloys, machines or atmospheres (Barrionuevo et al., 2025). Furthermore, feature sets are often ad-hoc features and lack physical constraints, leading to unrealistic predictions or trends that contradict well-established LPBF behavior (Staszewska et al., 2024; Atwya & Panoutsos, 2024).

These limitations have motivated the development of hybrid approaches in which domain knowledge is systematically embedded into the learning process. Early demonstrations include physics-informed and hybrid ML for fused filament fabrication (Kapusuzoglu & Mahadevan, 2020) and the porosity classification framework of Liu et al. (2021), which maps machine settings to physically meaningful variables. While promising, such efforts remain fragmented, scenario-specific and lack generalization.

A key barrier to broader transferability is the absence of a unified, structured repository of LPBF process knowledge. Existing data, mechanistic insights and physical descriptors are distributed across publications, isolated datasets and model implementations, with no standardized mechanism for encoding relationships between process parameters, materials, machines and operating regimes. As a result, most ML models begin and end within narrow laboratory contexts, limiting their applicability across conditions.

These gaps motivate the physics-informed machine learning framework proposed in this work. The framework (i) constructs a structured and expandable knowledge base of LPBF process physics, (ii) defines a dimensionless and mechanistically meaningful feature space of physics-informed process descriptors (PIPDs), and (iii) learns compact symbolic expressions that capture the relationship between PIPDs and relative density. By combining the interpretability and physical grounding of physics-based modelling with the flexibility and efficiency of ML, this approach aims to deliver accurate, explainable and transferable models for LPBF density prediction. This work demonstrates that low-complexity, physics-informed symbolic models can accurately describe LPBF density while offering transparency and transferability.

2. Methodology

2.1. Physics-informed ML framework

Most existing machine-learning models for LPBF operate as black boxes: although they often achieve high numerical accuracy, they offer limited physical insight into the mechanisms governing melt-pool behavior, defect formation and density evolution. As a result, they provide little support for scientific understanding, process tuning or the development of generalisable process–property–performance relationships. A second limitation relates to deployment and reliability. Deep learning models can be difficult to calibrate, debug and validate in industrial environments, and their opaque

internal structure often undermines trust, an important concern for high-value manufacturing applications.

This study proposes an explainable and generalisable physics-informed machine-learning (PIML) framework that produces compact, physically grounded symbolic models for LPBF density and evaluates their transferability across variations in the built environment. The framework is composed of four layers that together embed domain knowledge directly into the learning process. The proposed framework consists of several layers.

First, a knowledge base (KB) acts as an extensible store of LPBF domain knowledge, capturing process variables, material properties, and physically meaningful (often dimensionless) process descriptors that form the building blocks of the model. Second, a data layer hosts the raw experimental conditions and measured performance indices (surface finish, porosity, relative densities), applies physically motivated pre-processing, and computes physics-informed process descriptors (PIPDs) from the raw inputs. Third, a symbolic regression engine searches over low-complexity analytical expressions that relate the PIPDs to relative density, while enforcing physical and dimensional constraints derived from the KB. Finally, a model interpretation and transferability assessment layer analyse and visualises the resulting expressions and evaluates their generalisation when applied to samples with previously unseen combinations of geometry and chamber atmosphere within same or different materials. The overall architecture of the framework is given in Fig. 1.

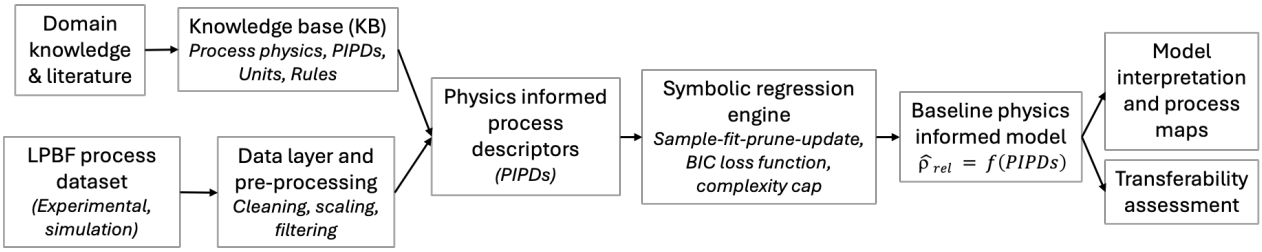

Fig. 1. Schematic of the proposed PIML framework.

2.2 Knowledge base and physics-informed process descriptors

The KB is designed as a generic, updateable repository of prior information on LPBF. It can be populated from published literature, online resources, expert experience and in-house experiments, and is not tied to any specific alloy or machine. The KB acts as the "physical backbone" of the PIML framework: it defines what constitutes a valid descriptor, how descriptors may interact, and which mechanisms are believed to dominate in typical LPBF regimes.

One of the key elements of KB is the normalised and dimensionless descriptors called physics-informed process descriptors (PIPDs). They are constructed from the standard LPBF inputs such as laser power, scan speed, hatch spacing, layer thickness and beam diameter, together with relevant material and thermal properties where needed. Rather than working directly with the process parameters, these descriptors aggregate them into quantities that reflect the underlying physics of heat input, melt-pool stability and track overlap. Typical examples include normalised linear energy density (NLED), normalised areal energy density (NAED) and normalised volumetric energy density (NVED); ratios that compare hatch spacing to layer thickness or beam size; and dimensionless groups for effective beam power and scan velocity based on thermal conductivity and diffusivity. In the context of the framework, the precise choice and formulation of PIPDs can be adapted to the process and material under study, but the common objective is to span a feature space that is dimensionally consistent, mechanistically meaningful, and suitable for guiding the symbolic regression towards interpretable, physics-consistent expressions (Thomas et al., 2016).

Together, these descriptors encode energy input, spatial overlap of melt tracks, and characteristic thermal diffusion scales. They span both traditional energy-density metrics and dimensionless numbers emerging from first principles arguments and form the candidate input space for symbolic regression. Beyond numerical features, the KB can in principle store higher-level rules (e.g. "increasing NLED at fixed hatch and layer thickness should not reduce density within the conduction regime") that could later be used as additional constraints or regularisers. In this study, we focus on

descriptor-level and dimensional constraints and keep such behavioural rules as soft guidelines for post-hoc interpretation.

2.3 Concept of a hybrid KB

In the present work, the knowledge base is implemented in a lightweight manner using curated tables and Python routines that compute PIPDs and enforce basic physical constraints. This serves as a proof of concept for integrating explicit process physics into the learning pipeline. In the longer term, we envisage re-implementing the KB as a structured, graph-based database that can grow with the LPBF community. In such a graph-oriented KB, process parameters, materials, machines, build environments, PIPDs, physical mechanisms, and constraints are represented as nodes, while their dependencies, valid regimes, and provenance (e.g. specific experiments or publications) are encoded as typed edges. This design supports incremental extension: new descriptors, materials, alloys, process windows or mechanisms can be added without disturbing the existing structure, and researchers can issue rich queries such as "all dimensionless PIPDs involving laser power and thermal diffusivity for Ti-6Al-4V in conduction-mode regimes".

A complementary vector-based layer would index unstructured knowledge sources (journal articles, data descriptors, laboratory notes) using text embeddings. This layer provides a semantic retrieval interface that can surface relevant passages (for example, new correlations or revised parameter ranges) and propose them as candidate nodes and edges in the knowledge graph for expert curation. A collaborative web graphical user interface (GUI) could then act as the front end through which LPBF researchers, engineers and practitioners contribute new findings (e.g. additional PIPDs, updated mechanism maps, machine-specific corrections) directly into the KB, with versioning and attribution.

When building a physics-informed model for a particular use case, the PIML framework would query this evolving KB for the subset of descriptors, correlations and rules relevant to the chosen alloy, machine, build atmosphere and operating range. In this way, the same global, community-maintained knowledge graph can support many local modelling tasks, while each new study feeds back structured information, gradually strengthening the shared, physics-based understanding of LPBF. The overall conceptual framework is given in Fig. 2.

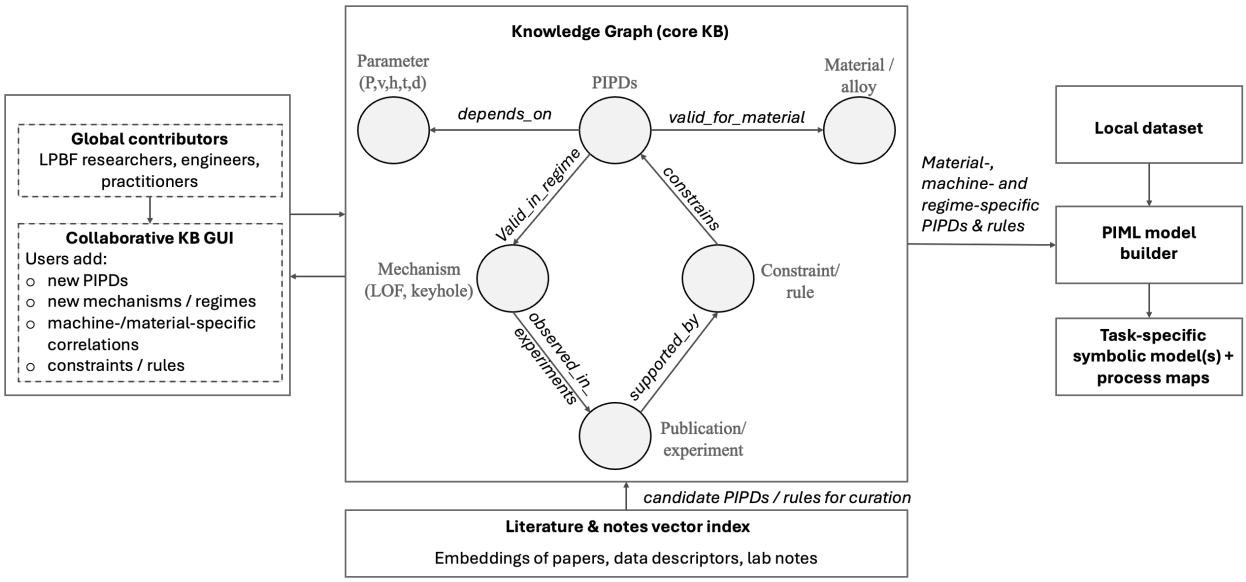

Fig. 2. Conceptual framework of the community-driven knowledge base.

2.4 Symbolic regressor

Symbolic regression is used to learn a closed-form mapping $\hat{\rho}_{rel} = f(x)$ from the vector of PIPDs 'x' to the relative density. We employ the Qlattice algorithm, which treats candidate expressions as computation graphs composed of simple operations (addition, subtraction, multiplication, division, exponential, logarithm, hyperbolic tangent, etc.) and searches over this space in a probabilistic way.

2.4.1 Search procedure and objective

At a high level, the Qlattice-based symbolic regression proceeds through an iterative sample–fit–prune–update cycle. In each iteration, the algorithm first samples a population of candidate model structures, represented as computation graphs, from its current probability distribution over topologies. This sampling is biased by the knowledge base, which favours PIPDs and operator combinations with physical justification and excludes dimensionally inconsistent constructions. For every sampled structure, the numerical parameters (such as coefficients and exponents) are then optimised on the training data using gradient-based backpropagation, with the mean squared error between predicted and measured relative density as the loss. Once fitted, the candidate models are ranked using the Bayesian Information Criterion (BIC)

$$\mathrm{BIC} = k\ln(n) - 2\ln(L)$$

where k is the number of free parameters, n is the number of training samples and L is the maximised likelihood; lower BIC values correspond to a better balance between accuracy and parsimony. Models that perform poorly or become unnecessarily complex are pruned away, while the best-scoring candidate is always retained. Finally, the Qlattice updates its internal sampling distribution to increase the probability of re-using structural motifs (specific PIPDs or operator patterns) found in high-scoring models, while preserving some exploration to discover new structures. This process is repeated until a convergence criterion is met or a maximum number of iterations is reached, yielding a set of compact, physics-consistent symbolic models.

2.4.2 Complexity control and embedded feature selection

To keep the resulting analytical expressions interpretable, we cap the maximum graph complexity (measured as the number of edges or operations) at 10, which implicitly limits how many PIPDs and interactions can appear in any single model. Because each candidate structure only uses a subset of the available descriptors and the BIC penalises additional parameters, the algorithm naturally performs embedded feature selection: descriptors that do not meaningfully improve the fit are discarded, and highly correlated PIPDs tend to compete such that only one representative typically remains in the final expression. To avoid converging onto a single, near-duplicate model, we retain a small ensemble of structurally distinct but high-performing candidates; comparing these allows us both to check the robustness of the discovered trends (for example, a monotonic dependence on an energy-density-type descriptor) and to recognise alternative but physically equivalent formulations (such as trading NVED for a combination of NLED and layer thickness).

After selecting a baseline symbolic model on the reference train/test split, we then examine its intra-material transferability by applying it, without any re-training, to samples with different printed geometries (e.g. prismatic versus cylindrical specimens) and to samples produced under different shielding atmospheres (e.g. nitrogen versus argon), where such combinations exist in the data. Performance on these held-out subsets is compared against the in-distribution test performance, and any degradation in accuracy is interpreted in light of the underlying descriptors - for instance, whether additional geometry- or atmosphere-sensitive PIPDs are required, or whether the existing descriptors already encode the key changes in melt-pool behaviour. In this way, the methodology delivers not only a predictive model for relative density, but also an explicit, physics-based analytical expression that can act as a reusable baseline for future studies involving other clusters, materials, or the incorporation of in-process sensing data.

3. Case Study

3.1 Source dataset

We use the open LPBF relative-density dataset compiled by Barrionuevo et al. (2025) which aggregates 1,579 experiments from 85 publications on commercial metallic alloys processed by LPBF. The dataset includes process parameters (laser power, scan speed, hatch spacing, layer thickness, spot size), powder characteristics (D50, median particle diameter)), a geometric factor, material type, shielding gas, printed geometry and the measured relative density.

In the original study, the authors performed extensive exploratory data analysis, calculated mutual information between each feature and relative density, and proposed a K-means clustering into four clusters to group experiments with similar process conditions and material characteristics. They further demonstrated that the dataset is suitable for training machine-learning regressors such as XGBoost and Random Forest. These models are however black box in nature and don't reveal the underlying physical mechanisms which influence the density and porosity of additively manufactured parts. For this study, we worked with a subset of the data (1 cluster, material - 18Ni300) in order to demonstrate (a) obtaining a clean baseline physics-informed symbolic model, and (b) studying model transfer across atmosphere (atm 1- Argon, atm 2- Nitrogen) within the same material.

3.2 Preprocessing

The pre-processing pipeline begins with filtering and basic cleaning. Only entries with complete information for the required process parameters, material label, geometry and shielding atmosphere are retained. Any records with inconsistent combinations of process parameters are removed, guided by the parameter ranges reported in the original data descriptor.

From the cleaned subset, we construct a compact set of physics-informed process descriptors (PIPDs) that encode normalized energy input and geometric overlap between melt tracks. As a first step, the core LPBF inputs: laser power P, scan speed S, hatch spacing h, layer thickness t and spot size d, are individually normalized, yielding the dimensionless variables P_n, S_n, h_n, $t_n$and d_n. On this basis, we define three energy-density-type descriptors: a linear energy density P/S, an areal energy density $P/(S\,h)$, and a volumetric energy density $P/(S\,h\,t)$. Each of these is again normalized to produce the PIPDs NLED, NAED and NVED. In parallel, we compute a set of geometric and overlap descriptors: the hatch-to-layer ratio $HT = h/t$; the dimensionless hatch $NH = h/d$; the track overlap fraction, $\mathrm{o} = 1 - h/d$; and a dimensionless layer thickness $\tau_t = t/d$. Together, these twelve PIPDs form the primary feature space passed to the symbolic regression engine. The resulting modelling data frame therefore contains the PIPDs, the target relative density, and auxiliary columns for shielding atmosphere (Atm) and original row index, the latter being used only for stratification and transferability studies rather than as predictors in the baseline symbolic model.

4. Results and Discussion

The proposed physics-informed symbolic modelling framework was first trained on the Argon-shielded subset of the 18Ni300 LPBF dataset. According to the proposed PIML framework, input process parameter space is transformed to a PIPD space of 12 dimensionless physical process descriptors as described in section 2.2 and section 3.2.

Using a maximum model complexity of six, which restricts the expression to at most three physics-informed process descriptors (PIPDs), the symbolic regression produced a compact analytical model that captured the dominant structure–property trends. The resulting expression

$$\hat{\rho}_{\mathrm{rel}} = -17.6\, e^{-2(0.27-11.2S_n)^2} - 2\, e^{-2(1.34-2\,\tau_t)^2} + 17.6\, \tanh(13.3\, NLED - 0.6) + 81.4 \quad (1)$$

showed excellent agreement with the experimental relative density under Argon conditions (R^2= 0.9288; R = 0.964). When this same model was transferred, without retraining, to the Nitrogen-shielded dataset, it retained good predictive power (R^2 = 0.8015; R = 0.896), with only moderate broadening of residuals (Fig. 3). This behaviour suggests that the model encodes atmosphere-independent physical relationships, although secondary effects related to shielding gas do introduce some performance degradation.

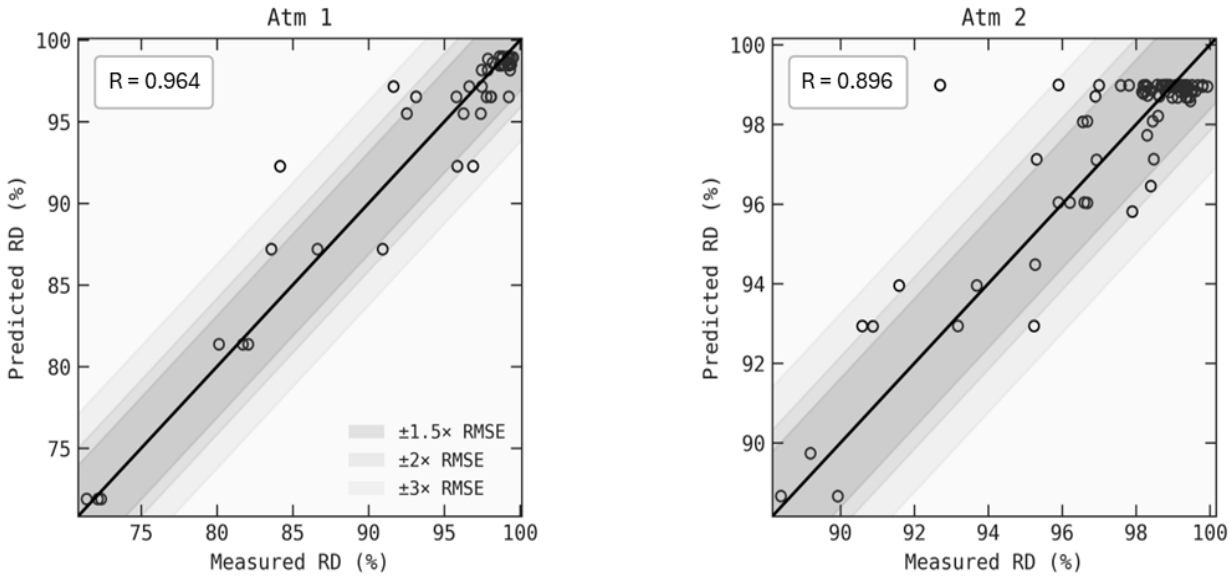

Fig. 3. (a) Training performance of model 1 (with 3 PIPDs) in Argon print atmosphere (b) model transferability to Nitrogen atmosphere.

To further investigate model compactness and interpretability, the symbolic regression was repeated with a stricter complexity limit of five, effectively restricting the solution to two PIPDs. Under this constraint, the symbolic regression produced a remarkably concise expression involving only the normalised linear energy density (NLED) and the normalised scan speed (S_n):

$$\hat{\rho}_{\mathrm{rel}} = 18.6 \tanh(12.5\, NLED - 0.5) + 4.63 \tanh(15.9\, S_n - 1.1) + 76 \quad (2)$$

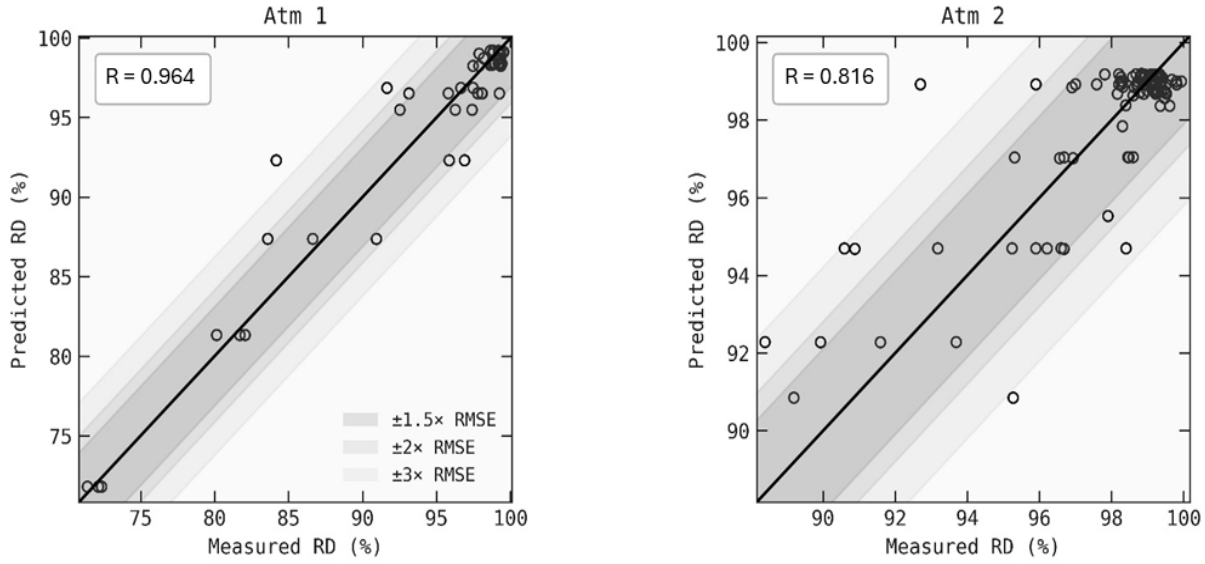

Fig. 4. (a) Training performance of model 2 (with just 2 PIPDs) in Argon print atmosphere (b) model transferability to Nitrogen atmosphere.

Despite its simplicity, this two-term model preserved nearly the same predictive accuracy on the Argon dataset (R^2 = 0.9293; R = 0.964), confirming that most of the variance in melt-pool quality and final density for 18Ni300 can be captured through two physically motivated descriptors. When evaluated on Nitrogen builds, the performance decreased moderately (R^2= 0.6656; R = 0.816), indicating that the omitted descriptor in the higher-complexity model likely captured subtle atmosphere-dependent behaviour such as variations in absorptivity or plume-driven instability (Fig. 4). Nonetheless, the two-descriptor formulation remains predictive and offers a highly interpretable baseline model.

4.1 Physical Insights

The symbolic model distilled by the PIML framework highlights two descriptors (NLED and Sn) as the dominant physical drivers governing densification in 18Ni300. Such compact relations enable rapid process-window estimation, reduce reliance on computational simulations, and support interpretable control strategies in production environments. The outcome is consistent with the fundamental LPBF physics linking melt-pool stability, thermal diffusion and defect formation. The plot (Fig. 5) depicts equation 2 (3d surface) embedded with experimental datapoints. Normalized linear energy density combines laser power and scan speed into a single physically interpretable measure of energy delivered per unit track length. Its influence in the model reflects the classical requirement for sufficient melt depth and track overlap to avoid lack-of-fusion. The saturating form of the NLED term corresponds to a well-known physical phenomenon: once the melt pool becomes large enough to reflow and remelt adjacent tracks, further increases in energy do not substantially improve densification due to thermal self-regulation and increased convective loss.

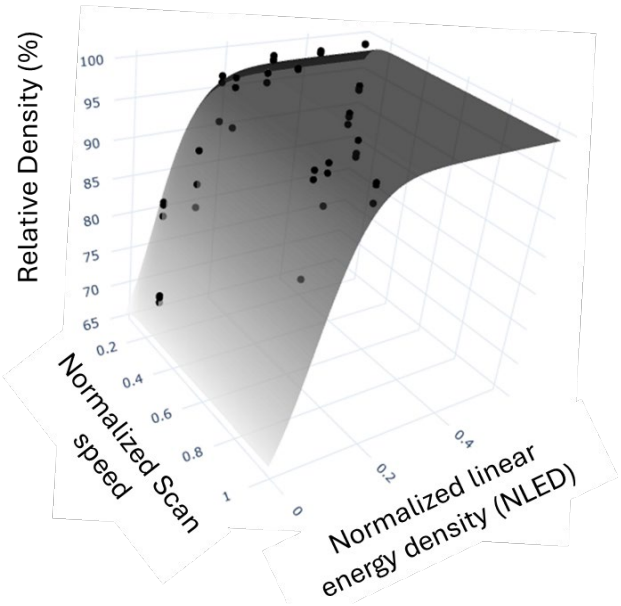

Fig. 5. Trained symbolic model plotted as 3D surface with experimental datapoints embedded.

While effective within the explored parameter space, the current descriptor set does not explicitly encode plume dynamics or absorptivity variation, which likely contributed to reduced performance under Nitrogen. Future work will incorporate these mechanisms into an expanded knowledge base and explore integration with in-process sensing modalities.

5. Conclusion

This work introduced a novel physics-informed machine learning (PIML) framework that unifies explicit LPBF process physics with symbolic regression to produce compact, interpretable and physically consistent models for relative density. By embedding domain knowledge into a structured space of dimensionless and normalised physics-informed process descriptors (PIPDs), the framework constrains learning to mechanistically meaningful relationships rather than opaque correlations. This design allows symbolic regression to operate within a physically coherent hypothesis space, yielding analytical expressions that retain the transparency of traditional process maps while achieving predictive performance comparable to state-of-the-art ML.

Applied to LPBF of 18Ni300, the framework generated concise symbolic models that captured the underlying density–process trends under Argon while demonstrating strong transferability to Nitrogen. The ability to recover accurate expressions using as few as two PIPDs highlights the framework's capacity to reveal minimal yet physically grounded forms that govern melt-pool behaviour. A key contribution of the framework is its extensibility: both the PIPD set, and the embedded physical rules can expand systematically as new scientific insights emerge. This opens the

path toward a community-driven, graph-based knowledge base, where researchers can contribute new descriptors, constraints, mechanistic hypotheses and experimental evidence. Such an evolving KB would allow future symbolic models to be trained on a richer, curated representation of LPBF physics, further improving interpretability, robustness and cross-condition generalization.

List of Acronyms

Acronym	Full form
AM	Additive manufacturing
CFD	Computational fluid dynamics
D50	Median powder particle diameter
GUI	Graphical user interface
KB	Knowledge base
LPBF	Laser powder bed fusion
LIME	Local Interpretable Model-agnostic Explanations
ML	Machine learning
NAED	Normalised areal energy density
NLED	Normalised linear energy density
NVED	Normalised volumetric energy density
PIML	Physics-informed machine learning
PIPD(s)	Physics-informed process descriptor(s)
SHAP	Shapley Additive exPlanations

Data Availability

The study was carried out using an open-source dataset given in in the Harvard Dataverse repository https://doi.org/10.7910/DVN/VPBQK8.

Declaration of Competing Interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this study.

Acknowledgement

The authors would like to acknowledge the funding from Innovate UK D3M_CoLAB and D3M-CoLab 2.

References

[1] DebRoy, T., Wei, H. L., Zuback, J. S., Mukherjee, T., Elmer, J. W., Milewski, J. O., … Zhang, W. (2018). Additive manufacturing of metallic components: Process, structure and properties. *Progress in Materials Science, 92*, 112–224.

[2] Kan, W. H., Chiu, L. N. S., Lim, C. V. S., … Lim, C. V. S. (2022). A critical review on the effects of process-induced porosity on the mechanical properties of alloys fabricated by laser powder bed fusion. *Journal of Materials Science, 57*, 9818–9865.

[3] Bruna-Rosso, C., Demir, A. G., & Previtali, B. (2018). Selective laser melting finite element modeling: Validation with high-speed imaging and lack of fusion defects prediction. *Materials & Design, 156*, 143–153.

[4] Tang, C., Tan, J. L., & Wong, C. H. (2018). A numerical investigation on the physical mechanisms of single-track defects in selective laser melting. *International Journal of Heat and Mass Transfer, 126*, 957–968.

[5] Zakirov, A., Belousov, S., Bogdanova, M., … Shakhov, F. (2020). Predictive modeling of laser and electron beam powder bed fusion additive manufacturing of metals at the mesoscale. *Additive Manufacturing, 35*, 101236.

[6] Chen, K., Zhang, P., Yan, H., Chen, G., Sun, T., Lu, Q., ... & Shi, H. (2024). A review of machine learning in additive manufacturing: Design and process. *The International Journal of Advanced Manufacturing Technology*, *135*(3), 1051-1087.

[7] Tapia, G., & Elwany, A. H. (2015). Prediction of porosity in slm parts using a mars statistical model and bayesian inference, 2015 *International Solid Freeform Fabrication Symposium*.

[8] Tapia, G., Elwany, A. H., & Sang, H. (2016). Prediction of porosity in metal-based additive manufacturing using spatial Gaussian process models. *Additive Manufacturing*, *12*, 282-290.

[9] García-Moreno, A. I., Alvarado-Orozco, J. M., Ibarra-Medina, J., & Martínez-Franco, E. (2020). Image-based porosity classification in Al-alloys by laser metal deposition using random forests. *The International Journal of Advanced Manufacturing Technology*, *110*(9), 2827-2845.

[10] Khalad, A., Telasang, G., Kadali, K., Zhang, P. N., Xu, W., & Chinthapenta, V. (2024). A generalized machine learning framework for data-driven prediction of relative density in laser powder bed fusion parts. *International Journal of Advanced Manufacturing Technology, 135*, 4147–4167.

[11] Chen, Y.-M., Lu, J.-L., Yu, D., … Wang, J.-C. (2025). Accurate identification of high relative density in laser-powder bed fusion across materials using a machine learning model with dimensionless parameters. *Acta Metallurgica Sinica (English Letters), 38*, 1645–1656.

[12] Maitra, V., Arrasmith, C., & Shi, J. (2024). Introducing explainable artificial intelligence to property prediction in metal additive manufacturing. *Manufacturing Letters*, *41*, 1125-1135.

[13] Kapusuzoglu, B., & Mahadevan, S. (2020). Physics-informed and hybrid machine learning in additive manufacturing: Application to fused filament fabrication. *JOM, 72*(12), 4695–4705.

[14] Liu, R., Liu, S., & Zhang, X. (2021). A physics-informed machine learning model for porosity analysis in laser powder bed fusion additive manufacturing. *International Journal of Advanced Manufacturing Technology, 113*(7), 1943–1958.

[15] Staszewska, A., Patil, D.P., Dixith, A.C. et al. A machine learning methodology for porosity classification and process map prediction in laser powder bed fusion. Prog Addit Manuf 9, 1901–1911 (2024). https://doi.org/10.1007/s40964-023-00544-2.

[16] Atwya, M., & Panoutsos, G. (2024). In-situ porosity prediction in metal powder bed fusion additive manufacturing using spectral emissions: A prior-guided machine learning approach. *Journal of Intelligent Manufacturing, 35*(6), 2719–2742.

[17] Barrionuevo, G.O., La Fé-Perdomo, I. & Ramos-Grez, J.A. Laser powder bed fusion dataset for relative density prediction of commercial metallic alloys. *Sci Data* **12**, 375 (2025). https://doi.org/10.1038/s41597-025-04576-x.

[18] Thomas, M., Baxter, G. J., & Todd, I. (2016). Normalised model-based processing diagrams for additive layer manufacture of engineering alloys. *Acta Materialia*, *108*, 26-35.