Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Geometry-aware interpretable learning for predicting print time and filament consumption in fused deposition modeling

Hoa-Cuc. Nguyen
Thu Dau Mot University
Viet-Hung. Pham
Thu Dau Mot University
Bich-Ngoc. Mach
Thu Dau Mot University
Ngoc-Duong. Nguyen
Thu Dau Mot University
Tran-Thi Hong
NTT Hi-Tech Institute, Nguyen Tat Thanh University
Thanh Q. Nguyen
Baria Vungtau University
Ma-Thi-Anh Thu
thumta@bvu.edu.vn

Baria Vungtau University

Research Article

Keywords: Fused deposition modeling, additive manufacturing, interpretable machine learning, multi-output regression, mesh geometry, print time prediction, filament consumption estimation, process parameter analysis

Posted Date: August 19th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10452213/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Geometry-aware interpretable learning for predicting print time and filament consumption in fused deposition modeling

Abstract

Efficient fused deposition modeling requires more than heuristic tuning because printing outcomes are governed by a nonlinear interaction between mesh geometry and process settings. This study developed a geometry-aware interpretable learning framework to predict two operationally critical responses—print time and filament consumption—and to clarify which variables truly govern pre-print decision making in fused deposition modeling. The framework was built from a dataset of 1,500 simulated printing instances generated from 150 meshes, with 17 input variables comprising 7 geometry-derived descriptors and 10 printer-setting parameters. After categorical encoding, normalization, and outlier filtering, Random Forest, Extreme Gradient Boosting, and a stacking-based multi-output regression strategy were trained and comparatively evaluated. The results showed that the proposed framework achieved strong predictive fidelity for both targets. The stacking configuration yielded an overall mean squared error of 0.0001929, a mean absolute error of 0.0063829, a root mean squared error of 0.01361, and an R^2 of 0.9619. At the target level, the R^2 values reached 0.9645 for print time and 0.9449 for filament length, substantially outperforming the previously reported convolutional neural network–multilayer perceptron baseline. Interpretability analysis further revealed that mesh geometry dominated both outputs: volume and area were the strongest drivers of filament consumption, whereas area and volume were the principal determinants of print time. Among controllable settings, speed was most influential for filament demand, while skirt outlines, layer height, and extrusion width exerted the strongest effect on printing time. Practical validation further indicated prediction deviations within the sub-10% range for representative print cases. These findings demonstrate that geometry-aware interpretable learning can provide accurate operational prediction while supporting reduced-complexity printer configuration in fused deposition modeling.

Keywords: *Fused deposition modeling; additive manufacturing; interpretable machine learning; multi-output regression; mesh geometry; print time prediction; filament consumption estimation; process parameter analysis*

1. Introduction

Fused deposition modeling has become one of the most accessible and widely adopted additive manufacturing technologies because it combines relatively low equipment cost, broad material availability, and the capacity to fabricate geometrically diverse parts for engineering, educational, medical, and consumer-oriented applications [1]. Its practical appeal lies not only in rapid prototyping, but also in its growing role in customized and small-batch production, where flexible fabrication and short design-to-product cycles are increasingly valuable [2]. Yet, despite this broad adoption, efficient fused deposition modeling remains difficult to achieve in routine practice. The operational performance of a print is still highly dependent on operator experience, and the final configuration of a print job is often determined through repeated trial-and-error rather than through a transparent and reproducible decision process [3,4]. In practical settings, users must simultaneously balance fabrication time, material usage, and acceptable manufacturing outcomes, but these decisions are rarely supported by a framework that can systematically connect the geometry of the part to the controllable settings of the printer.

This difficulty arises because fused deposition modeling is governed by a nonlinear interaction between two sources of variation that are often treated separately [5,6]. The first is the geometry of the input mesh, which determines the structural burden imposed on the printer through scale, spatial extent, and geometric complexity [7]. The second is the configuration of the printing process, including layer-related parameters, extrusion conditions, speed, infill, and internal pattern decisions that control how the part is physically fabricated. In practice, these two components are tightly coupled [8]. A setting combination that performs acceptably for one mesh may become inefficient or suboptimal for another because the underlying deposition path, layer count, and material demand change with geometry [2,7]. This means that printer settings cannot be interpreted as universally optimal in isolation; their effect is conditioned by the geometry of the part being printed. The challenge is therefore not simply to identify good settings in general, but to determine how mesh-dependent information should guide printer configuration in a way that is both accurate and operationally useful [8].

A substantial body of prior work has attempted to improve fused deposition modeling through parameter optimization, material selection, process refinement, and data-driven modeling [9,10]. These studies have contributed important insight into the influence of printing speed, layer structure, extrusion behavior, and material-related decisions on final manufacturing performance [11,12]. However, the literature remains fragmented in a way that limits its practical deployment for pre-print decision making. One group of studies focuses primarily on optimizing process parameters for specific materials or printing conditions, but often underrepresents the role of mesh geometry [12,13]. Another group adopts machine learning to predict isolated responses, yet these predictions are frequently treated as stand-alone outputs rather than as the basis for actionable configuration guidance [14,15]. Even where machine learning has been introduced to simplify or accelerate optimization, the broader connection between geometric burden, printer settings, and operational outputs often remains weakly formalized [16,17]. This gap is important because real printer configuration is not a single-variable tuning task; it is a coupled decision problem in which geometric information and process variables must be interpreted together.

The practical importance of this gap becomes especially clear when attention is focused on two quantities that dominate pre-print planning [18]: print time and filament consumption. These two responses directly determine scheduling feasibility, material allocation, cost estimation, and throughput planning [3,19]. For individual users and service-oriented printing environments alike, knowing how long a print will take and how much filament it will consume is often more immediately valuable than a generic claim about printing performance [20,21]. At the same time, both outputs are strongly affected by the interaction between part geometry and machine settings. A model that ignores geometry cannot fully capture the burden associated with the printed object, while a model that ignores controllable settings cannot support meaningful configuration decisions [22,23]. Moreover, high predictive accuracy alone is insufficient for deployment in engineering settings. Users also need to know which variables truly matter, which can be simplified, and which may be safely left at default values without materially affecting operational efficiency. For this reason, interpretability is not an optional add-on in the present context; it is central to the transformation of predictive modeling into deployable printer-setting intelligence.

The present study addresses this need by formulating fused deposition modeling as a geometry-aware and interpretable multi-output learning problem that jointly predicts print time and filament consumption from mesh descriptors and printer-setting variables. The proposed framework was constructed from a dataset of 1,500 simulated printing instances derived from 150

meshes, where each mesh was evaluated under multiple feasible configuration settings. The input space comprised 17 variables, including 7 geometry-derived descriptors—Area, Volume, Vertices, Faces, Height, Depth, and Width—and 10 printer-setting parameters selected from operationally meaningful control variables. Rather than treating these as disconnected inputs, the framework integrates them within a unified learning architecture so that the model can capture how the physical characteristics of the printed part and the execution logic of the printer jointly shape the two operational targets [14,20]. This formulation is more faithful to the actual structure of fused deposition modeling than parameter-only or single-output approaches, because it recognizes that time and material demand arise from the same coupled geometry–process state.

To model this relationship, the study adopts ensemble-based machine learning through Random Forest, Extreme Gradient Boosting, and a stacking-based fusion strategy [24,25]. Yet the scientific contribution of the work does not lie merely in comparing algorithms. More importantly, the framework uses interpretable learning to identify the dominant drivers of operational behavior and to distinguish between structural variables that define the printing burden and controllable variables that provide operational leverage [26]. The importance rankings reported later in the study show that mesh-related descriptors, especially area and volume, dominate both targets, while only a small subset of printer settings—such as speed, skirt outlines, layer height, and extrusion width—emerges as consistently influential. This is a consequential finding because it suggests that efficient fused deposition modeling does not require exposing users to an excessively large parameter space. Instead, it supports a reduced-complexity view of configuration in which geometry-aware prediction and variable prioritization work together to simplify decision making for non-expert users.

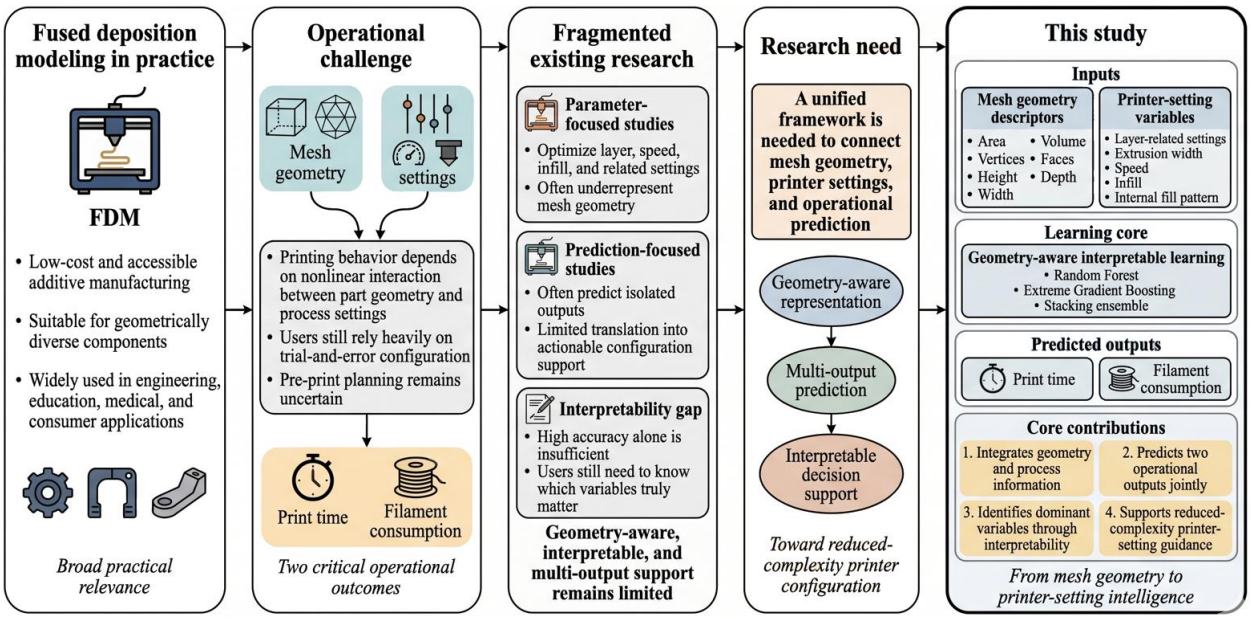


Figure 1. *Conceptual overview of the research motivation, knowledge gap, and core contributions of the proposed geometry-aware interpretable learning framework for predicting print time and filament consumption in fused deposition modeling*

Figure 1 synthesizes this logic by linking the practical relevance of fused deposition modeling with the unresolved difficulty of geometry-dependent printer configuration and the central contribution of the present study. As the figure illustrates, the problem begins with the broad utility

of FDM in real manufacturing contexts, but quickly narrows to a core operational challenge [9,11]: print behavior depends on the interaction between mesh geometry and process settings, while existing research has often treated parameter optimization, prediction, and interpretability as separate directions. The framework proposed here was designed precisely to bridge that divide. By integrating geometry-aware representation, multi-output prediction, and interpretable variable ranking in a single workflow, the study moves beyond traditional process tuning and establishes a more deployable pathway from mesh geometry to printer-setting intelligence [12,27]. In this sense, the novelty of the work lies not only in predictive performance, but in the way prediction is converted into operationally meaningful guidance for fused deposition modeling.

Accordingly, the contribution of this paper is threefold. First, it establishes a geometry-aware learning framework that jointly models mesh descriptors and printer-setting variables for the prediction of print time and filament consumption. Second, it demonstrates that interpretability can reveal a reduced set of materially important configuration variables, thereby clarifying the true operational control surface of the process. Third, it provides the foundation for simplified and more accessible printer configuration support by showing how geometry-aware multi-output learning can transform raw printing data into actionable pre-print intelligence. Through this perspective, the study positions efficient fused deposition modeling not merely as a parameter optimization task, but as an interpretable decision problem that connects part geometry, process execution, and operational planning within a unified additive manufacturing framework.

2. Related Work

2.1. Process-centric optimization in fused deposition modeling

A large share of the fused deposition modeling literature has been built on the assumption that printing performance can be improved primarily through parameter optimization [23,28]. This stream has generated an extensive body of work on layer thickness, printing speed, extrusion temperature, raster strategy, infill density, and shell-related variables, typically with the aim of improving dimensional accuracy, surface quality, or mechanical behavior [29]. Even in recent studies, the dominant logic remains process-centric [30,31]: the printing problem is framed as one of identifying better operating settings under fixed material and specimen conditions. Dev and Srivastava [27], for example, used a hybrid response surface methodology and genetic algorithm strategy to optimize flexural strength in FDM, while Nikooharf et al. [32] formulated fused filament fabrication parameter design under uncertainty by combining multiphysics modeling, uncertainty quantification, and optimization. More recently, Deka and Hall [33] proposed a predictive-modeling-plus-response-surface framework to optimize FDM parameters when experimental datasets are limited or non-systematically collected. Collectively, these studies have strengthened the methodological foundation of process optimization in extrusion-based additive manufacturing by showing that parameter effects are nonlinear, coupled, and often too complex for one-factor-at-a-time analysis [13].

However, the process-centric tradition also exhibits a structural limitation that becomes more evident when the research objective shifts from post-print property improvement to pre-print operational prediction [22]. Most parameter-optimization studies are conducted on a restricted family of geometries, standard tensile or flexural specimens, or fixed benchmark parts. Under those conditions, the apparent influence of process settings may be amplified because the geometric burden of the part is effectively controlled or treated as background information [34]. This makes such studies highly valuable for local process understanding, but less effective for generalized

printer-setting support across diverse meshes [12,35]. As recent reviews have emphasized, additive manufacturing performance emerges from the interaction of material, geometry, process physics, and execution logic rather than from settings alone. Nikooharf et al. [32] show that machine learning in polymer additive manufacturing increasingly spans property prediction, process planning, monitoring, and optimization, precisely because no single parameter family can fully characterize the behavior of the process. Zhang et al. [23] similarly review machine learning-driven 3D printing as an ecosystem involving process optimization, monitoring, and motion planning, underscoring that process parameters are only one part of a larger decision structure. Within that broader perspective, process-centric optimization remains essential, but it is no longer sufficient if the goal is to predict time and material demand before fabrication begins for geometrically diverse parts.

This observation is especially relevant to the present study because the operational outputs of interest—print time and filament consumption—are not only controlled by how the printer is configured, but also by what the printer is asked to fabricate [11]. A part with larger area, greater volume, or more complex geometric structure imposes a fundamentally different deposition burden from a compact or simpler mesh, even under the same nominal setting combination. As a result, a purely parameter-focused literature provides only a partial answer to the problem addressed here [13]. It clarifies the importance of process tuning, but it does not adequately explain how those settings should be interpreted once part geometry is allowed to vary substantially across print jobs [10]. The present work is therefore positioned not against this literature, but beyond it: it accepts the importance of parameter optimization, while arguing that the operationally meaningful unit of analysis in FDM is the geometry–process interaction rather than the process parameter set alone.

2.2. Data-driven prediction and the shift from tuning to inference

A second, more recent stream of work has moved from direct optimization toward data-driven prediction [25]. In this literature, the central question is no longer only which parameter combination is optimal, but whether machine learning can infer process behavior, part properties, or manufacturing outcomes from observed data [36]. This shift is important because it reflects a broader recognition that additive manufacturing systems are too nonlinear, too multivariate, and too context-dependent to be captured reliably by low-order empirical rules alone [9]. Deka and Hall [33] developed predictive models of mechanical properties for FDM parts using data from multiple sources and explicitly introduced environmental information to improve robustness. Nguyen et al. [18] proposed a data-driven machine learning approach to 3D printing process optimization that attempted to connect model design and process settings more directly than traditional trial-and-error workflows. Wang et al. [28] advanced this direction further by using machine learning to infer 3D printing parameter settings for desired mechanical properties, effectively moving from forward prediction toward an inverse-design logic. These studies collectively demonstrate that machine learning can do more than support process tuning; it can reshape the way the configuration problem itself is posed.

Even so, the data-driven literature remains uneven in two ways that are highly relevant for the current paper. First, many studies continue to focus on mechanical properties, geometric accuracy, or defect-related responses rather than on the operational outputs that drive pre-print planning [16,17]. This is understandable, since strength, stiffness, roughness, and dimensional fidelity are central manufacturing objectives. Yet for many real users, knowing the expected printing time and filament demand before launching a job is equally critical because these outputs govern schedule, throughput, and material logistics [14]. Second, much of the predictive literature remains

either single-output or weakly connected to decision support [15,37]. A model may achieve excellent prediction of one response, but still provide limited assistance if it does not reveal how multiple operational targets are jointly shaped by the same underlying input state [24]. The reviews by Nikooharf et al. [32] and Zhang et al. [23] both point to this broader fragmentation: machine learning applications in additive manufacturing are expanding rapidly, but the field still contains many highly specialized models built around narrow tasks rather than unified operational frameworks. Ördek et al. [12] make a similar point from a manufacturing-wide perspective, arguing that machine-learning support systems often remain overly specialized and limited in scope, which constrains their transferability and practical deployment.

This is precisely where the present study departs from much of the existing predictive literature [38]. Rather than treating process settings as the only meaningful predictors or focusing on a single post-print quality metric, the current work frames FDM as a geometry-aware multi-output problem in which print time and filament consumption are learned jointly from mesh descriptors and controllable settings [9,12]. This formulation is motivated by the observation that both outputs arise from the same geometry–process state, even though their sensitivities are not identical. In that sense, the study aligns with the broader data-driven movement in additive manufacturing, but reorients it toward operational prediction and pre-print intelligence. It asks not only whether outcomes can be predicted, but whether those predictions can be organized around the specific information a user needs before fabrication begins. That shift from generic predictive performance to operational inference is one of the defining distinctions of the proposed framework.

2.3. Interpretability, configuration support, and the unresolved gap

A third strand of recent work has emphasized that prediction alone is not enough. As machine learning has become more capable in manufacturing contexts, interpretability has emerged as a critical requirement for systems that are intended to support real decisions rather than merely produce post hoc estimates [21,39]. This is especially true in additive manufacturing, where many input variables correspond to physically meaningful settings that operators can directly adjust [40]. If a model predicts an outcome but does not clarify which variables matter most, which interactions dominate, or which settings can be simplified, then its practical usefulness remains limited. Kharate et al. [41] illustrate this clearly in their explainable-AI study of FDM-based biocomposites, where the purpose of the analysis is not only to improve prediction but also to expose the relationship between process parameters and material properties. More broadly, the recent review by Ördek et al. (2024) argues that manufacturing-oriented machine-learning systems often suffer from excessive specialization and weak integration of their inputs and outputs, reinforcing the need for models that are both informative and decision-relevant. The interpretability challenge is therefore not secondary; it is one of the main conditions for turning manufacturing analytics into deployable support tools [20].

Figure 2 synthesizes this literature landscape into three converging but still insufficiently integrated directions. The first direction, process-centric optimization, has clarified how individual FDM parameters affect manufacturing responses, but often under fixed or narrowly varying geometries [27,41]. The second direction, data-driven prediction, has shown that machine learning can capture nonlinear behavior more effectively than conventional empirical modeling, but many of these models remain focused on isolated outputs or property-specific tasks [18,32]. The third direction, interpretable manufacturing analytics, has begun to reveal how model outputs can be translated into engineering understanding, yet this strand has only recently started to influence additive manufacturing workflows in a systematic way [26,42]. As **Figure 2** makes clear, the central

gap does not lie in the absence of optimization, prediction, or explanation as separate ideas. It lies in the lack of a unified geometry-aware, multi-output, and interpretable framework that can connect mesh descriptors, printer settings, and operational pre-print outcomes within a single decision structure.

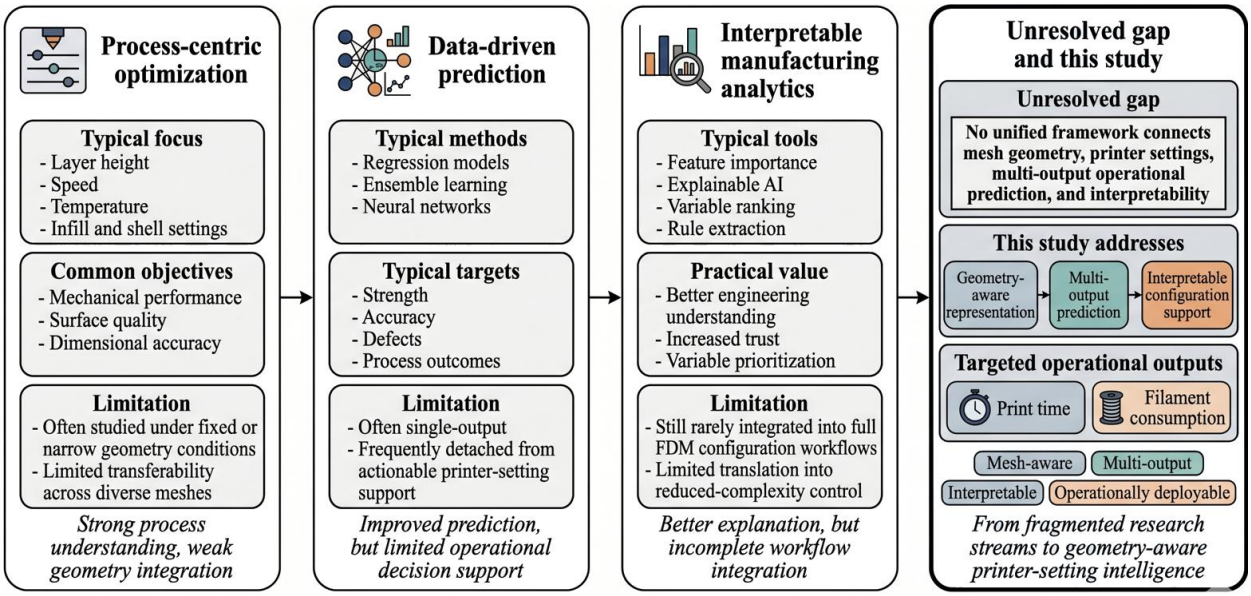


Figure 2. *Conceptual synthesis of the related-work landscape and the research gap addressed by the proposed geometry-aware interpretable multi-output learning framework*

This gap directly motivates the present study. Existing parameter-focused work does not fully address part-dependent variability. Existing predictive work does not consistently prioritize the two operational outputs that matter most before printing begins [43,44]. Existing interpretability-oriented studies have not yet established a reduced-complexity control view in which only the subset of materially influential printer settings is emphasized under geometry-aware conditions [21,45]. The contribution of the current paper is therefore not simply to add another machine-learning model to the additive manufacturing literature. Rather, it is to establish a framework in which geometry-aware representation, multi-output operational prediction, and interpretable variable ranking are designed as mutually reinforcing components. By doing so, the study shifts the discussion from "which model predicts best under a given dataset" to "how mesh geometry can be transformed into actionable printer-setting intelligence for efficient FDM planning." That is the research space left open by prior studies, and it is the space this paper is designed to occupy.

3. Materials and Data Description

3.1. Printing environment and dataset construction

The empirical foundation of this study was established within a sustained fused deposition modeling workflow that had been operated and documented since 2019. Rather than being assembled from isolated laboratory trials, the dataset emerged from a long-running printing environment in which print samples, parameter settings, material usage, execution time, and related production records were continuously accumulated through routine fabrication activities [20,46]. This context is important because it gives the dataset a practical rather than purely demonstrative character. The study did not begin with a synthetic benchmark designed only for algorithm comparison; instead, it drew on an operational record structure created to support a

broader transition from expert-dependent printer setup toward a more automated and reproducible configuration process. In the original study context, the printing workflow combined in-house machine deployment, model preparation, and expert-informed process control, thereby producing a data resource that was both technically grounded and operationally meaningful.

The selected manufacturing context was fused deposition modeling with polylactic acid as the primary material system [2,8]. This choice was not incidental. Fused deposition modeling remains the most practically accessible additive manufacturing route for desktop, laboratory, and small-scale production environments, while PLA provides stable printability, wide availability, and low operational complexity relative to many other thermoplastic systems [16,17]. These characteristics make the data especially suitable for a study focused on operational prediction rather than on highly specialized material behavior. The aim was not to characterize one unusual process condition in depth, but to construct a generalizable representation of how ordinary FDM print jobs behave when geometry and process settings vary together [9,31]. In this sense, the dataset was designed to reflect the real decision space faced by users who must configure a printer before fabrication begins and who require reliable estimates of time and filament demand. As illustrated conceptually in **Figure 3**, the data foundation of the study can therefore be understood as a three-part system: an operational FDM environment, a controlled process-parameter space, and a geometry-rich set of printable meshes that together produce a learning-ready record of coupled printing behavior [27,41].

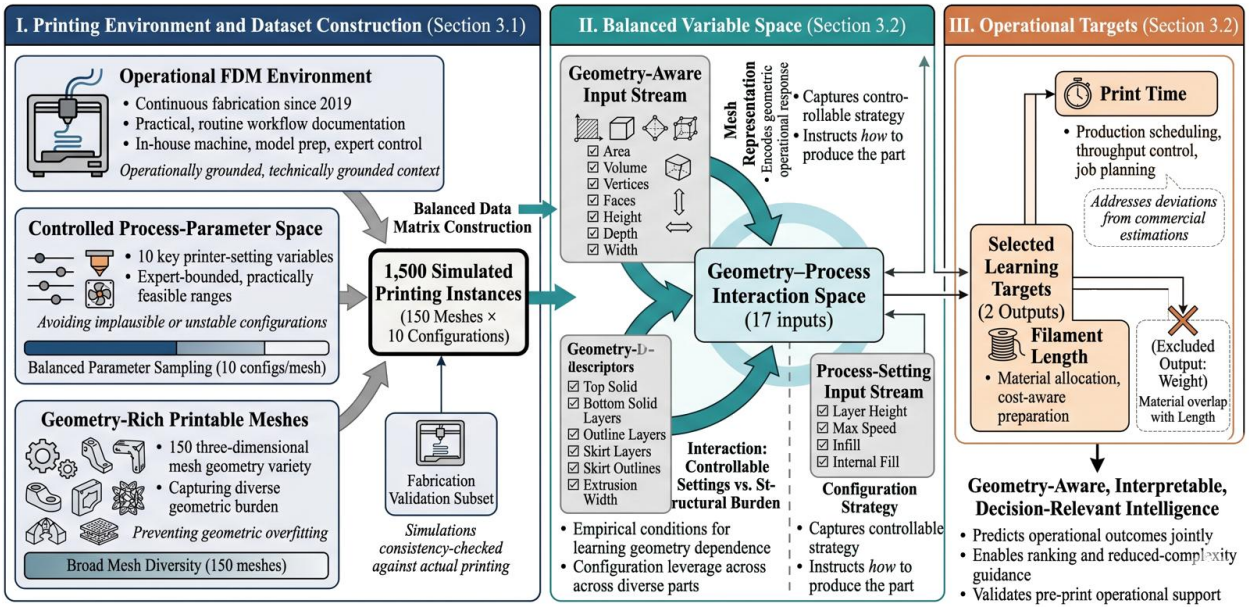


Figure 3. *Conceptual overview of the data foundation, variable structure, and operational targets used in the proposed geometry-aware interpretable learning framework*

The construction of the dataset followed a deliberately balanced design. The final data matrix comprised 1,500 simulated printing instances derived from 150 three-dimensional meshes, with each mesh evaluated under 10 distinct configuration settings. This structure is one of the most important strengths of the study because it captures two different but complementary sources of variation. On the one hand, the dataset preserves diversity across printed parts by using multiple meshes with different geometric characteristics. On the other hand, it preserves process variability by evaluating each mesh under multiple parameter combinations within expert-approved operating

ranges. The result is a dataset that does not collapse the problem into either a mesh-only or a settings-only formulation. Instead, it allows the model to learn how the same geometry responds to different configurations and how the same nominal process setting behaves differently across different geometries. This dual variation is central to the scientific logic of the paper, because the study argues that efficient FDM cannot be understood from process tuning alone; it must be interpreted through the interaction between part-dependent burden and controllable machine behavior.

The parameter space itself was not arbitrarily sampled. The original workflow relied on expert knowledge to identify printer settings that materially affect printing performance and that remain physically feasible in practice. Ten printer-setting variables were retained for explicit modeling, including layer-related, shell-related, flow-related, and fill-related controls. These parameter ranges were chosen to reflect realistic operating conditions rather than unrestricted combinatorial exploration. This design choice matters because it prevents the dataset from being distorted by implausible or unstable configurations that would have little practical value for real users. In parallel, the study retained broad mesh diversity so that the data would not be overfitted to a single geometric family. The resulting dataset therefore occupies a useful middle ground: it is sufficiently controlled to support meaningful learning and sufficiently diverse to represent the operational variability of FDM printing jobs. **Figure 3** highlights this balance by showing how expert-bounded settings and geometry-dependent part descriptions are merged into a unified data-generation workflow rather than treated as separate sources of evidence.

An additional feature of the dataset is its link to downstream fabrication reality. Although the main learning resource consisted of simulation-supported print instances, the original study also reports that only a subset of more than 3,000 available models was selected for physical printing because of the high cost and time burden associated with full-scale fabrication experiments. Through this process, the simulation setup was checked against actual printing behavior, and the authors explicitly note that the established parameters were consistent with the real-world process. This connection is important for the current manuscript because it strengthens the legitimacy of using simulation-supported operational data for predictive learning. The dataset is not merely synthetic in a detached sense; it is anchored in a real printing workflow and informed by real fabrication constraints. For a study whose main outputs are pre-print time and filament estimates, such an operationally grounded simulation framework is especially appropriate.

From a methodological standpoint, this data-construction strategy also helps explain why the proposed framework is able to support both prediction and interpretation. If the study had relied only on varied meshes under a single configuration, it would have learned geometry dependence but not configuration leverage. Conversely, if it had relied only on many settings for a fixed class of parts, it would have learned parameter sensitivity but not geometry-aware operational behavior. By combining multiple geometries and multiple feasible configurations per geometry, the dataset creates the empirical conditions necessary for identifying both the structural burden imposed by the part and the smaller control surface provided by the printer settings. This balance is the foundation on which the later interpretability results become meaningful.

3.2. Geometry–process variable space and operational targets

The final dataset was organized around a structured variable space containing 17 input variables and 3 recorded outputs, of which 2 outputs were retained as the principal targets for the current study. The 17 input variables were divided into two conceptually distinct groups. The first

group consisted of 7 geometry-derived descriptors: Area, Volume, Vertices, Faces, Height, Depth, and Width. These variables encode the physical characteristics of the input mesh and serve as a compact representation of geometric burden. In practical terms, they describe how large, how spatially extended, and how structurally intricate the printed part is likely to be. Such information is essential because it influences the amount of deposition, movement, and layer generation required before the part is complete. The second group consisted of 10 printer-setting variables: Top Solid Layers, Bottom Solid Layers, Outline Layers, Skirt Layers, Skirt Outlines, Extrusion Width, Layer Height, Max Speed, Infill, and Internal Fill Pattern. These variables capture the controllable strategy through which the printer executes fabrication. Collectively, they describe not what the part is, but how the machine is instructed to produce it.

This division between geometry-derived and process-setting variables is central to the scientific design of the present paper. The framework proposed in this study is explicitly geometry-aware, which means that the mesh is not treated as passive background information but as an active determinant of the operational response. At the same time, the model remains interpretable and decision-relevant because it also preserves the process settings that a user can actually manipulate. **Figure 3** synthesizes this internal structure by visually separating the geometry stream from the process-setting stream and then showing how both are fused into a common predictive space. This representation is important because it reinforces the conceptual argument of the paper: print time and filament consumption are not outputs of settings alone, but responses generated by the interaction between the physical structure of the mesh and the execution logic of the printer. Without the geometry block, the model would miss the dominant burden imposed by the part. Without the settings block, it would lose the operationally controllable component of the process.

The recorded outputs in the original dataset included Internal Fill Pattern, Time, Length, and Weight, but the operational learning task of the present study focuses on print time and filament length as the two principal targets. This target selection was deliberate and methodologically justified. Time is the key quantity for production scheduling, throughput control, and job planning, particularly when the printer must handle multiple queued tasks. Length, measured as the total filament consumed, is the key quantity for material allocation and cost-aware preparation. The original document also explains that estimated print time in commercial environments may deviate from actual behavior, which further increases the value of a dedicated predictive model. Weight, although recorded, was not retained as a primary target in the present formulation because it largely reflects material use and therefore overlaps substantially with filament consumption. Reducing the target space to print time and filament length sharpens the operational focus of the study and aligns the prediction framework with the two quantities most relevant before printing begins.

The data table itself combines numerical, ordinal, and categorical information. Geometry-related variables span broad ranges because the selected meshes differ considerably in size and complexity, whereas several printer-setting variables operate within narrow but practically meaningful expert-defined intervals. Internal fill pattern introduces a categorical dimension because it distinguishes between fill strategies such as grid, honeycomb, and rectilinear structures. This mixed-input composition is one of the reasons why the current study adopts a structured preprocessing and ensemble-learning pipeline rather than a simpler linear approach. The dataset is not only multivariate; it is heterogeneous in both physical meaning and numerical scale. That heterogeneity is scientifically useful because it reflects realistic FDM usage, but it also requires careful organization so that the resulting predictive model remains stable and interpretable.

The broader significance of this variable structure lies in its ability to support a more deployable form of FDM intelligence. Many prior studies in additive manufacturing either emphasize process settings while underrepresenting the part itself, or focus on predictive power without clarifying which variables matter most. The variable design of the present dataset avoids both limitations. It includes enough geometric information to model part-dependent operational burden and enough process information to reveal where users still retain meaningful control. This is why the Section 3 material does more than describe inputs and outputs. It establishes the core empirical premise of the paper: operational efficiency in fused deposition modeling is governed by a geometry–process interaction space, and any model that seeks to predict print time and filament consumption in a useful way must preserve that interaction explicitly. By structuring the dataset in this form, the study creates the conditions for later sections to show not only that the outputs can be predicted, but also that the dominant sources of variation can be ranked and converted into reduced-complexity configuration support.

4. Methodology

4.1. Geometry-aware problem formulation

The present study formulates fused deposition modeling as a geometry-aware multi-output regression problem in which the physical burden imposed by a mesh and the controllable settings of the printer are learned jointly to predict two operationally critical responses: print time and filament consumption. This formulation departs from conventional parameter-only optimization by recognizing that printer settings do not operate in a vacuum. Instead, the same nominal configuration can lead to substantially different outcomes when the geometry of the printed part changes in area, volume, spatial extent, or structural complexity. The original dataset structure supports this view directly: 150 meshes were evaluated under 10 feasible configurations each, producing 1,500 simulated printing instances with 17 input variables and 2 principal operational targets retained for the present manuscript.

Let the complete dataset be denoted by

$$\mathcal{D} = \left\{\left(\mathbf{x}_i^{(g)}, \mathbf{x}_i^{(p)}, \mathbf{y}_i\right)\right\}_{i=1}^{N}, \tag{1}$$

where $N = 1500$, $\mathbf{x}_i^{(g)} \in \mathbb{R}^7$ is the geometry descriptor vector of the i-th instance, $\mathbf{x}_i^{(p)} \in \mathbb{R}^{10}$ is the printer-setting vector, and $\mathbf{y}_i \in \mathbb{R}^2$ is the output vector. The two input groups are written as

$$\mathbf{x}_i^{(g)} = [A_i, V_i, Ve_i, Fa_i, H_i, D_i, W_i], \tag{2a}$$

$$\mathbf{x}_i^{(p)} = [TSL_i, BSL_i, OL_i, SL_i, SO_i, EW_i, LH_i, MS_i, IF_i, IFP_i], \tag{2b}$$

while the output vector is defined as

$$\mathbf{y}_i = \begin{bmatrix} t_i \\ \ell_i \end{bmatrix}, \tag{3}$$

where t_i and ℓ_i denote print time and filament length, respectively. The full predictive task is therefore represented by the nonlinear mapping

$$\hat{\mathbf{y}}_i = f\left(\mathbf{x}_i^{(g)}, \mathbf{x}_i^{(p)}\right) + \boldsymbol{\varepsilon}_i. \tag{4}$$

Equation (4) is the core methodological statement of the paper. It expresses that the model does not estimate the outputs from process variables alone, nor from geometry alone, but from their coupled state. This is a crucial distinction because the study is not concerned with generic property prediction; it is concerned with pre-print operational intelligence. As illustrated in **Figure 4**, the proposed workflow begins by separating the geometry stream and the process-setting stream, then fusing them into a common representation before moving into multi-output learning. The figure is important because it visually clarifies that geometry is treated as a first-class predictor rather than as passive contextual information. In turn, this framing allows the subsequent methodology to distinguish between variables that define printing burden and variables that provide operational leverage.

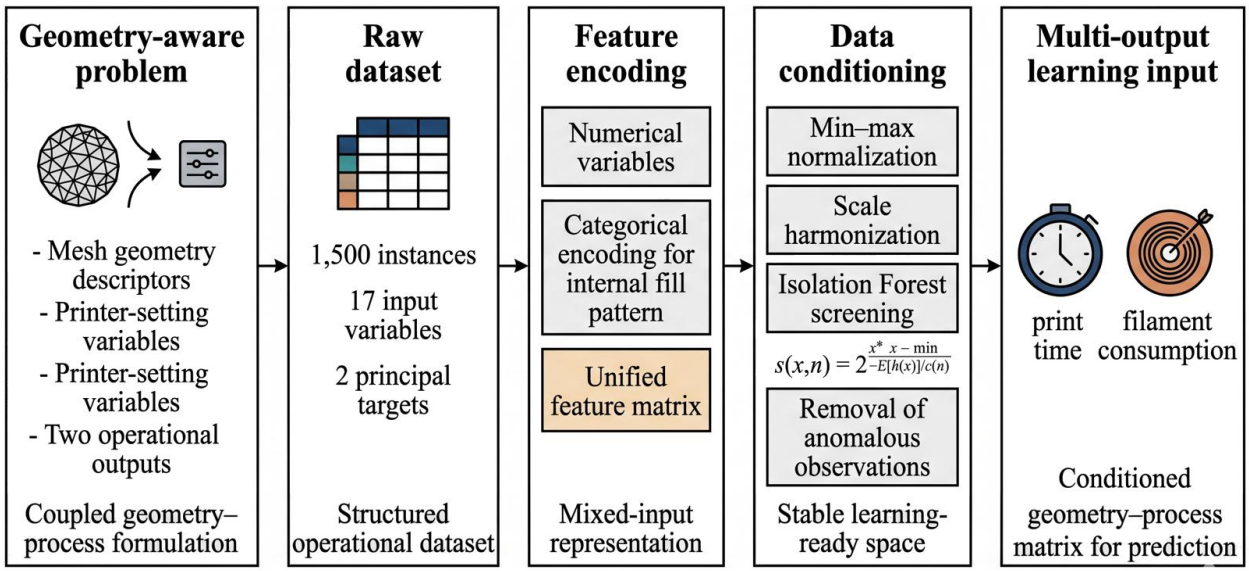


Figure 4. *End-to-end methodological pipeline from geometry–process representation to conditioned multi-output learning*

A further methodological decision follows from the target definition. Although the original dataset recorded time, length, and weight, the present framework prioritizes print time and filament length as the two principal outputs because they are the most directly actionable before a print begins. Weight was excluded as a primary target because it is strongly tied to material use and therefore contributes limited additional operational information once filament length is modeled. This target reduction sharpens the learning objective and aligns the model with the actual decision context of FDM users, who need reliable estimates of how long a job will take and how much material it will require.

4.2. Data conditioning and structured representation

Because the dataset contains numerical, ordinal, and categorical variables with markedly different ranges, the second methodological stage focused on transforming the raw records into a stable learning representation. The internal fill pattern variable is categorical, while the remaining variables are continuous or integer-valued. To place all predictors in a comparable input space, the

raw feature matrix $\mathbf{X} \in \mathbb{R}^{N \times d}$ was first encoded and then normalized using min–max scaling. For a given feature x_{ij}, the normalized value is written as

$$x_{ij}^{*} = \frac{x_{ij} - \min(x_j)}{\max(x_j) - \min(x_j)}. \tag{5}$$

Equation (5) ensures that variables with very large numerical magnitudes, such as volume or number of vertices, do not dominate variables with narrow operational ranges such as extrusion width or layer height. In a geometry-aware dataset, this is particularly important because the scale difference between geometric descriptors and process settings can be substantial.

The output matrix was preserved in continuous form:

$$\mathbf{Y} = \begin{bmatrix} t_1 & \ell_1 \\ t_2 & \ell_2 \\ \vdots & \vdots \\ t_N & \ell_N \end{bmatrix}. \tag{6}$$

Equation (3) makes explicit that the problem is treated as a multi-output rather than a sequence of isolated single-target regressions. This choice is methodologically justified because both outputs are generated by the same geometry–process state and therefore share part of their explanatory structure, even though their sensitivities are not identical.

A second conditioning step addressed data irregularities. The original study reported that some observations exhibited abnormal scaling behavior when the printed object became much larger or smaller than the average, which introduced noise and made learning more difficult. To screen these observations, the methodology adopted Isolation Forest. For an observation **x**, the anomaly score is defined as

$$s(\mathbf{x}, n) = 2^{-\frac{E[h(\mathbf{x})]}{c(n)}}, \tag{7}$$

where $E[h(\mathbf{x})]$ is the expected path length over random isolation trees and $c(n)$ is the average path length for a sample of size n, approximated by

$$c(n) = 2H(n-1) - \frac{2(n-1)}{n}. \tag{8}$$

Here, $H(\cdot)$ denotes the harmonic number. A larger score indicates that an instance is easier to isolate and is therefore more likely to be anomalous. In practical terms, this mechanism removes observations that distort the geometry–process relationship without discarding the wider diversity of the dataset.

To define the screening threshold, the original workflow further used a quartile-based rule:

$$\tau = q_1 - 1.5(q_3 - q_1), \tag{9}$$

where q_1 and q_3 are the first and third quartiles, respectively. Approximately 200 rows were removed through this procedure, leaving a more stable conditioned dataset for subsequent training.

As shown in **Figure 4**, data conditioning is not a peripheral preprocessing step but a structural component of the framework. It ensures that the geometry-aware representation used for learning reflects meaningful operational behavior rather than scale-driven anomalies. This is especially important for later interpretability analysis, because unstable observations can distort the ranking of influential variables and weaken the reliability of configuration guidance.

4.3. Multi-output ensemble learning architecture

The predictive core of the framework is based on ensemble learning, specifically Random Forest and Extreme Gradient Boosting, because the geometry–process relationship in FDM is strongly nonlinear and exhibits heterogeneous local sensitivities. Random Forest serves as the bagging-based learner. If B denotes the total number of trees and $T_b(\cdot)$ the b-th tree, the multi-output prediction is expressed as

$$\hat{\mathbf{y}}^{RF}(\mathbf{x}) = \frac{1}{B}\sum_{b=1}^{B} T_b(\mathbf{x}). \tag{10}$$

Equation (10) shows that the forest prediction is the average of multiple randomized tree estimators trained on bootstrap-resampled data. This structure reduces variance and improves robustness, making it suitable for tabular data that mix geometry and process inputs. The original training results also indicate that Random Forest achieved the lowest RMSLE during tuning, supporting its role as a strong base learner in the present problem.

Extreme Gradient Boosting complements Random Forest by adopting a sequential boosting strategy. The boosted predictor can be written as

$$\hat{\mathbf{y}}^{XGB}(\mathbf{x}) = \sum_{k=1}^{K} f_k(\mathbf{x}), f_k \in \mathcal{F}, \tag{11}$$

where K is the number of boosting rounds and $\mathcal{F}$ is the space of regression trees. The optimization objective takes the form

$$\mathcal{L} = \sum_{i=1}^{N} l(\mathbf{y}_i, \hat{\mathbf{y}}_i) + \sum_{k=1}^{K} \Omega(f_k), \tag{12}$$

with a regularization penalty

$$\Omega(f_k) = \gamma T_k + \frac{\lambda}{2} \| \mathbf{w}_k \|_2^2, \tag{13}$$

where T_k is the number of leaves, $\mathbf{w}_k$ is the vector of leaf weights, and γ, λ are regularization parameters. Equations (11)–(13) make clear that XGBoost learns by sequentially correcting residual structure while explicitly controlling model complexity. This is valuable in the current setting because the geometry–process interaction may contain local nonlinearities that are difficult to capture by averaging alone.

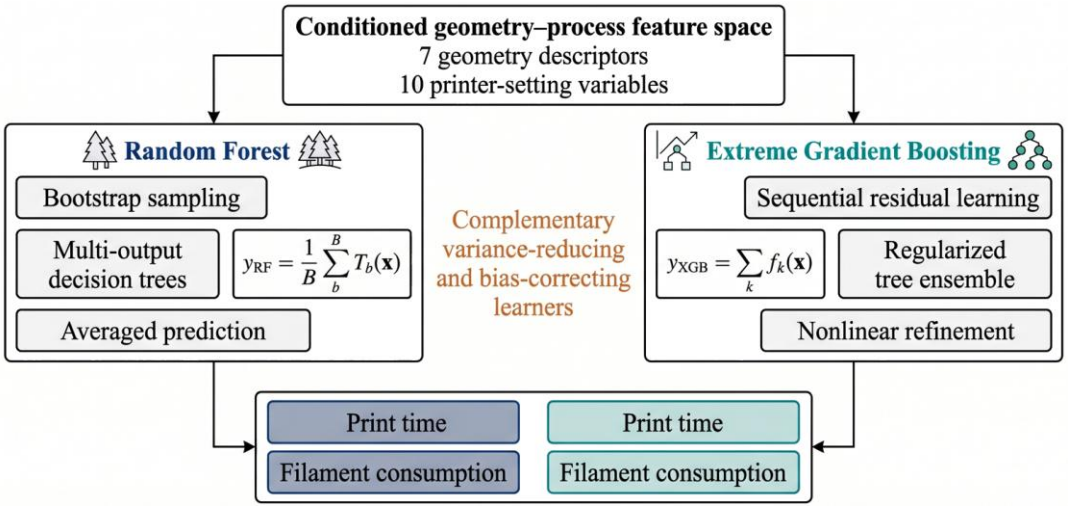


Figure 5. *Dual-branch ensemble learning architecture for geometry-aware multi-output prediction in fused deposition modeling.*

The complementarity of these two learners is central to the methodology. Random Forest provides stability and resistance to overfitting, while XGBoost offers greater sensitivity to structured residual patterns. **Figure 5** summarizes this dual-branch architecture by showing how the conditioned geometry–process representation enters both model families in parallel. The figure should not be read simply as a technical diagram of two algorithms. Its scientific role is to show that the framework intentionally combines two distinct perspectives on the same manufacturing problem: one variance-reducing and one bias-correcting. This complementarity later becomes the basis for stacking and for a more stable interpretation of the predictive space.

4.4. Stacking fusion and interpretable decision-support layer

Although each base learner produces an independent multi-output estimate, the final framework introduces a stacking layer to combine their strengths. Let

$$\hat{\mathbf{y}}_i^{RF} = \begin{bmatrix} \hat{t}_i^{RF} \\ \hat{\rho}_i^{RF} \end{bmatrix}, \hat{\mathbf{y}}_i^{XGB} = \begin{bmatrix} \hat{t}_i^{XGB} \\ \hat{\rho}_i^{XGB} \end{bmatrix}. \tag{14}$$

These are concatenated into a meta-feature vector

$$\mathbf{z}_i = \begin{bmatrix} \hat{t}_i^{RF} \\ \hat{\rho}_i^{RF} \\ \hat{t}_i^{XGB} \\ \hat{\rho}_i^{XGB} \end{bmatrix}. \tag{15}$$

The final stacked prediction is then generated through a linear meta-learner:

$$\hat{\mathbf{y}}_i^{Stack} = \mathbf{W}\mathbf{z}_i + \mathbf{b}, \tag{16}$$

where **W** is the coefficient matrix and **b** is the intercept vector. This formulation is consistent with the original study, in which a simple linear regression model was used as the meta-learner over the outputs of Random Forest and XGBoost.

The role of stacking here is not simply to pursue marginal numerical gain. It also provides a structured fusion mechanism in which the final prediction remains interpretable at the meta-level. Because the stacking layer is linear in the base-model outputs, it is possible to inspect how much each predictive stream contributes to the final estimate. This matters for a paper whose broader goal is not just prediction, but configuration support.

The final methodological step converts prediction into interpretability-driven guidance. Let $I_j^{(t)}$ and $I_j^{(\ell)}$ denote the normalized importance of feature j for print time and filament length, respectively. A joint operational relevance score can be defined as

$$I_j^{(joint)} = \alpha I_j^{(t)} + (1-\alpha) I_j^{(\ell)}, 0 \le \alpha \le 1. \tag{17}$$

Equation (17) allows the model to reconcile variable importance across the two outputs instead of ranking features in isolation. This is methodologically useful because some variables dominate both targets, while others influence only one. In turn, the recommendation objective for a given geometry $\mathbf{x}^{(g)}$ and candidate process vector $\mathbf{x}^{(p)}$ can be expressed as

$$\mathbf{x}^{(p)*} = \arg\min_{\mathbf{x}^{(p)} \in \mathcal{P}} J\left(\mathbf{x}^{(g)}, \mathbf{x}^{(p)}\right), \tag{18}$$

where $\mathcal{P}$ is the feasible configuration space and $J(\cdot)$ is an operational objective. For example,

$$J_{fast} = \hat{t}, J_{material} = \hat{\ell}, J_{balanced} = \beta\tilde{t} + (1-\beta)\tilde{\ell}. \tag{19}$$

These expressions formalize the three configuration modes implicit in the original study—fastest print, best print, and optimal print—but reframe them in a more rigorous decision-support language.

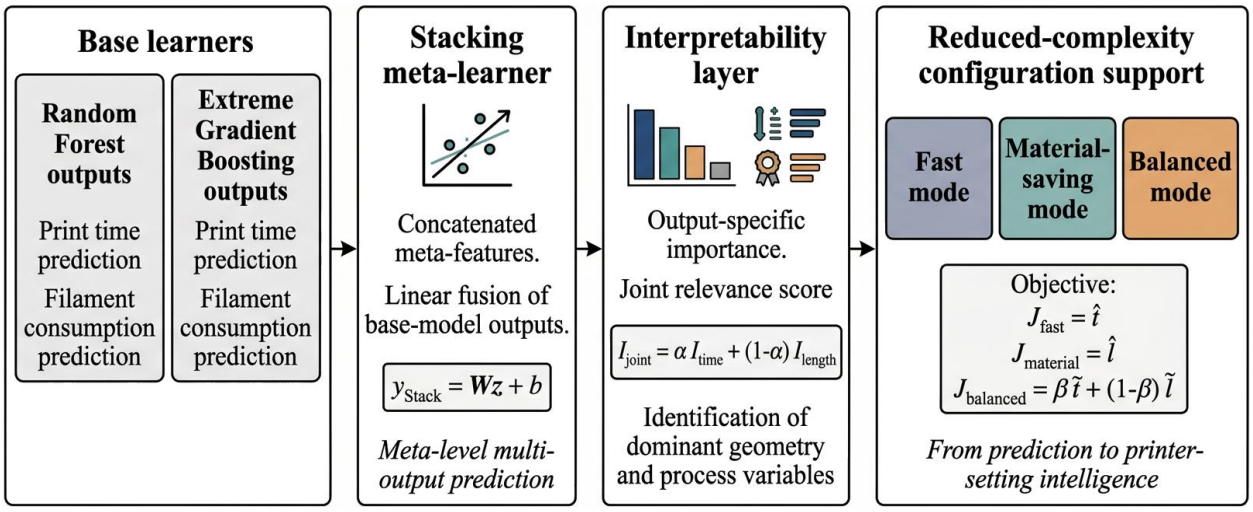


Figure 6. *Stacking-based fusion and interpretability-driven recommendation workflow for reduced-complexity printer configuration.*

Figure 6 captures this last transition particularly well. It shows that the framework does not end with model prediction; rather, it proceeds into an interpretability and recommendation layer that extracts dominant variables, ranks controllable settings, and supports simplified configuration guidance. This is one of the main scientific contributions of the paper. The methodology is therefore not merely a sequence of machine-learning steps. It is a coherent workflow that transforms mesh geometry and feasible printer settings into actionable operational intelligence for fused deposition modeling.

5. Experimental Design

5.1. Training protocol and hyperparameter optimization

The experimental design was constructed to test the proposed framework under conditions that reflect the real structure of the prediction problem rather than a simplified benchmark scenario. Because the dataset combines heterogeneous mesh geometries, mixed variable types, and two operational outputs, the training stage had to do more than fit a model to the entire dataset. It had to separate model development from model assessment, control the risk of information leakage, and ensure that hyperparameter selection was driven by stable validation behavior rather than by a favorable random split. For this reason, the experimental workflow was organized as a staged procedure consisting of dataset partitioning, hyperparameter optimization, cross-validation, full training with the selected settings, and final held-out evaluation. This structure is summarized in **Figure 7**, which shows that the study was designed around a disciplined train–validate–test sequence rather than a direct fit-and-report strategy. Such a design is especially important in geometry-aware learning, where the predictive space is shaped not only by process settings but also by broad variation in the geometry of the printed parts.

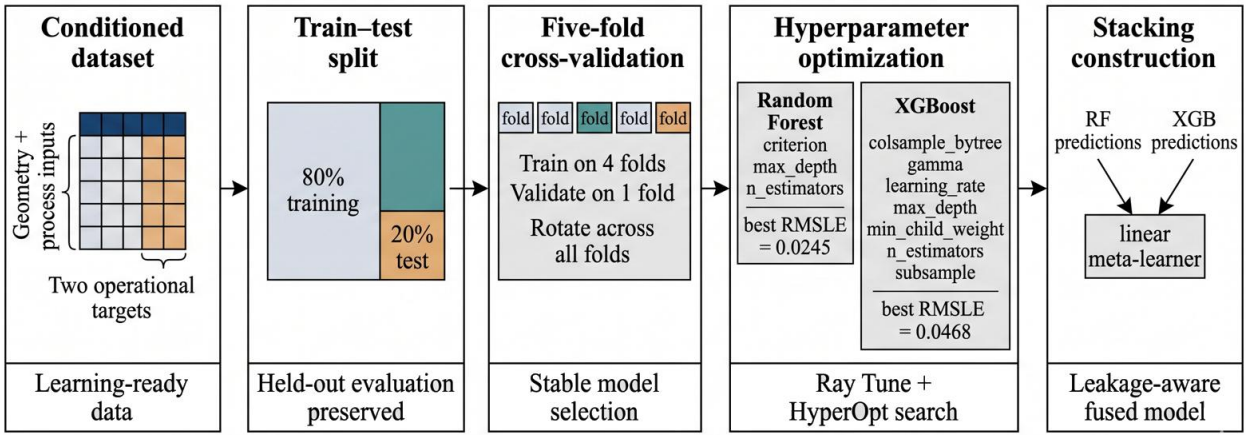


Figure 7. *Training workflow and hyperparameter-optimization protocol for the geometry-aware multi-output prediction framework*

The conditioned dataset was divided into training and test subsets using an 80:20 split, with the larger portion used for model fitting and hyperparameter selection and the smaller portion reserved for final evaluation. This choice follows the logic already embedded in the original study and provides a reasonable balance between learning capacity and out-of-sample verification. Within the training partition, model selection was conducted through five-fold cross-validation, so that each model was trained repeatedly on four folds and validated on the remaining fold before rotating across all partitions. This strategy serves two purposes. First, it reduces the dependence of model selection on a single split of the training data. Second, it yields a more stable estimate of

generalization behavior under heterogeneous mesh–process combinations. In the current context, where one geometry can be paired with multiple feasible configurations, cross-validation is not simply a routine statistical step; it is a way of testing whether the model preserves predictive consistency when the local geometry–process composition of the sample changes across folds. **Figure 7** makes this point visually by showing that validation is embedded inside the model-development loop rather than appended only at the end.

Hyperparameter tuning was carried out with Ray Tune and the HyperOpt search strategy, as reported in the source manuscript. This choice is methodologically appropriate because the search space includes model families with different complexity controls and response behavior. For Random Forest, tuning focused on tree criterion, tree depth, and ensemble size, resulting in the optimal configuration criterion = squared_error, max_depth = 40, n_estimators = 100, with a lowest tuning-stage RMSLE of 0.0245. For Extreme Gradient Boosting, tuning explored sampling and boosting controls, resulting in the configuration colsample_bytree = 0.8, gamma = 0.1, learning_rate = 0.2, max_depth = 15, min_child_weight = 5, n_estimators = 50, objective = reg:squarederror, subsample = 0.8, with a lowest RMSLE of 0.0468. These settings are not merely implementation details. They indicate that the Random Forest branch benefited from deeper hierarchical partitioning and a sufficiently large ensemble, whereas the XGBoost branch required a moderate learning rate and regularized tree growth to avoid over-amplifying local residual structure. The fact that both branches were tuned independently also reinforces the logic of the later stacking design: the study did not fuse arbitrary off-the-shelf models, but combined two individually optimized predictors.

A final and important feature of the training protocol concerns the construction of the stacking framework. Because the meta-learner must operate on the outputs of the base learners, the study used cross-validated base-model predictions to build the meta-level training representation. In practice, this means that the Random Forest and XGBoost branches were first validated across the training folds, their predictions were collected, and only then was the linear meta-learner fitted over that derived prediction space. This design prevents the stacking layer from benefiting from in-sample leakage and ensures that the final fused model reflects real predictive complementarity rather than memorized training behavior. As Figure 7 shows, the stacking procedure is therefore nested inside the broader training protocol rather than detached from it. This is one of the reasons why the final benchmark remains methodologically credible: each layer of the framework was developed under separation between tuning, validation, and final testing.

5.2. Comparative baselines and benchmark design

Once the training workflow had been defined, the second task of the experimental design was to establish a benchmark structure capable of evaluating both predictive performance and methodological value. The study was not intended as a simple competition between models with increasingly complex architectures. Instead, the comparative design was organized to answer a more specific question: does a geometry-aware and interpretable ensemble framework provide a stronger basis for predicting time and filament demand than single-model alternatives or the previously reported CNN–MLP reference? To address this, four predictive configurations were retained in the benchmark landscape: Random Forest, Extreme Gradient Boosting, the proposed stacking framework, and the previously developed CNN–MLP model used as a historical reference point for similar 3D-printing data. This comparative logic is synthesized in **Figure 8**, which places the four models within one evaluation ecosystem rather than presenting them as disconnected predictors.

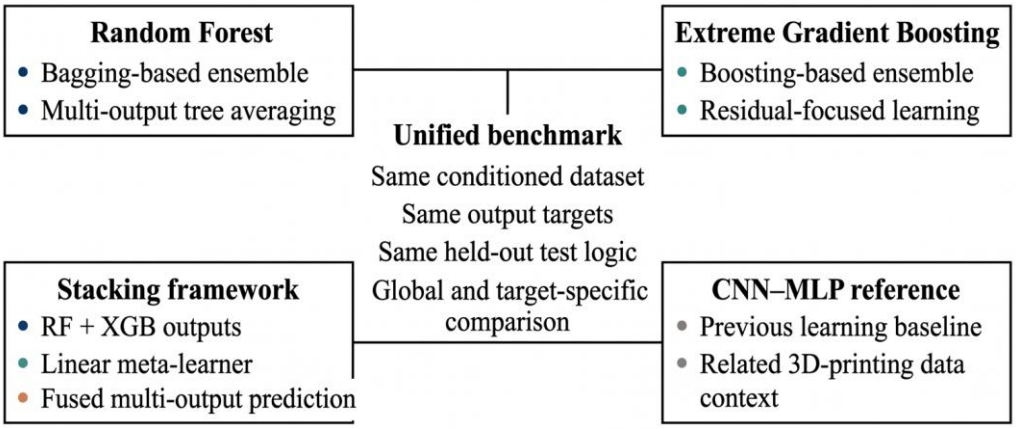


Figure 8. *Comparative benchmark design across ensemble learners, stacking fusion, and the CNN–MLP reference model*

The role of the two tree-based ensemble learners in this benchmark is complementary rather than redundant. Random Forest represents the bagging-oriented branch and is expected to offer stable prediction under heterogeneous tabular inputs, while XGBoost represents the boosting-oriented branch and is expected to capture more aggressive residual structure. The stacking framework, in turn, was introduced to test whether these two nonlinear perspectives could be combined in a controlled and interpretable way through a linear meta-learner. The CNN–MLP model serves a different purpose. It is not part of the new ensemble family, but a reference to the prior learning strategy used by the authors for a related geometry-aware printing problem. This makes it particularly useful in the present manuscript because it provides a historical baseline against which the benefits of the current structured ensemble approach can be assessed. Figure 8 highlights this clearly by showing three levels of comparison: base ensemble models, fused ensemble architecture, and legacy neural reference.

The benchmark design was also constructed to preserve the operational character of the task. Unlike many studies in additive manufacturing that focus on one mechanical property or one post-print quality indicator, the present evaluation targets two pre-print outputs simultaneously: print time and filament length. This means that benchmark superiority cannot be claimed on the basis of a single scalar score alone. A useful model must show strong global behavior while also maintaining balanced prediction across both outputs. For that reason, the comparative design was structured to report both overall performance and output-specific behavior. This dual evaluation strategy is central to the experimental logic of the paper. It allows the study to distinguish between models that are strong in aggregate but uneven across targets and models that remain consistently reliable for both time and filament prediction. In this regard, **Figure 8** plays a clarifying role because it presents the benchmark not simply as a list of models, but as a layered comparison framework aligned with the operational objectives of the study.

Another important feature of the benchmark design is that it preserves interpretability as part of the comparison rather than treating it as a post hoc add-on. The proposed stacking model is not only evaluated for accuracy, but also for its suitability as the basis of later variable-importance analysis and reduced-complexity configuration support. This means that the experimental benchmark is implicitly multi-criteria, even before the discussion section begins. A model is valuable not only if it predicts well, but also if it supports an explanation structure that can later be translated into printer-setting intelligence. The benchmark design therefore mirrors the broader

philosophy of the paper: prediction is necessary, but not sufficient. A model that performs well numerically but cannot support configuration guidance would only partially satisfy the goals of the study.

5.3. Metric hierarchy and robustness-oriented evaluation

The third component of the experimental design concerns the metric system used to judge performance. Because the problem is multi-output, heterogeneous, and operationally oriented, the study deliberately uses a hierarchy of error measures rather than relying on a single score. During model selection, the principal optimization metric was RMSLE, chosen because it reduces the disproportionate influence of very large target values and remains robust in the presence of residual noise. This choice is particularly appropriate in the present dataset, where the scale of time and material demand can vary substantially across different meshes. In a geometry-aware prediction problem, such magnitude variation is not noise in itself; it is part of the structure of the task. A logarithmic error measure helps preserve comparability across small and large jobs without allowing the latter to dominate the entire training objective. This rationale, already present in the source manuscript, is formalized in the current experimental design as a deliberate strategy for controlling scale effects during model selection.

For final evaluation on the test set, the design expands beyond RMSLE and reports MSE, MAE, MAD, RMSE, and R^2. Each of these metrics captures a different aspect of model behavior. MSE and RMSE place stronger emphasis on larger deviations and are therefore useful for identifying whether a model becomes unstable on certain mesh–configuration combinations. MAE gives a more intuitive measure of average absolute predictive error, while MAD provides an additional perspective on the dispersion of deviations around the median behavior. R^2, in turn, is included to quantify how much of the variation in the operational outputs is explained by the learned geometry–process mapping. The use of this metric set is summarized in **Figure 9**, where the metrics are arranged as an evaluation hierarchy from training-stage selection to final held-out assessment. The figure is particularly useful because it clarifies that different metrics are not competing alternatives in this study; they are complementary tools serving different roles in the evaluation chain.

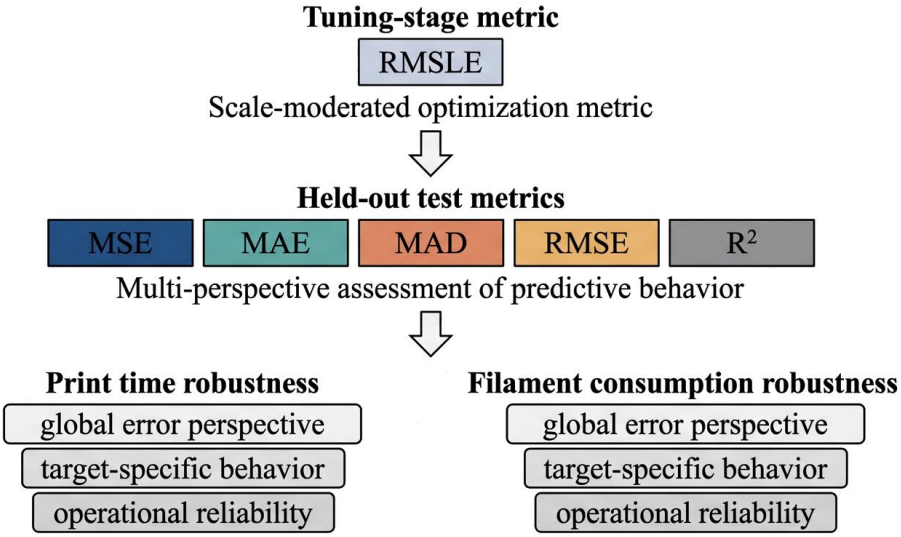


Figure 9. *Hierarchical evaluation-metric framework for model selection, held-out testing, and target-specific robustness analysis*.

The design also treats robustness as a target-specific property. Since print time and filament length are both operationally important, a useful model must perform well on both rather than excelling only in aggregate. Accordingly, the evaluation system was structured to produce both global multi-output results and separate results for each output. This allows the later results section to assess whether the same model behaves equally well for scheduling-related prediction and material-demand prediction, or whether trade-offs emerge between the two. Such a target-specific view is essential for the paper's practical relevance. In a real FDM environment, a model that predicts time accurately but misjudges filament demand could still lead to failed planning, just as a model that estimates filament well but mispredicts time could distort scheduling. Figure 9 reinforces this point by depicting the evaluation layer as bifurcating into two output-specific branches beneath the shared metric structure. This visually encodes a key methodological principle of the paper: robustness in multi-output FDM prediction must be examined at both the joint and target-specific levels.

A final strength of the metric design is that it aligns naturally with interpretability. Because the study later uses feature-importance ranking to build reduced-complexity configuration support, it is essential that the chosen performance criteria remain meaningful under the geometry–process decomposition of the dataset. The current hierarchy achieves this by separating scale-sensitive, scale-moderated, and explained-variance measures. In doing so, it ensures that later conclusions about dominant variables and useful control parameters are grounded in a model that has been validated from multiple quantitative perspectives rather than from one narrow error definition alone.

5.4. Practical validation and deployment-oriented test design

The final part of the experimental design concerns practical validation. Although the predictive framework was trained primarily on a simulation-supported dataset, the design of the study deliberately maintained a bridge to real fabrication behavior. This was achieved in two complementary ways. First, the original workflow had already established consistency between the simulated operational records and actual FDM printing conditions by verifying selected models physically. Second, the later validation experiments were structured to test whether the importance rankings and output predictions retained practical meaning when transferred back to representative real-print scenarios. In this sense, the practical validation stage was not an afterthought; it was part of the deployment logic of the framework. The goal was to confirm that the model not only performs well statistically, but also remains useful as a pre-print advisory tool in realistic configuration tasks. This logic is summarized in **Figure 10**, which presents the validation design as the final layer of the experimental pipeline and explicitly connects numerical prediction to physical print interpretation.

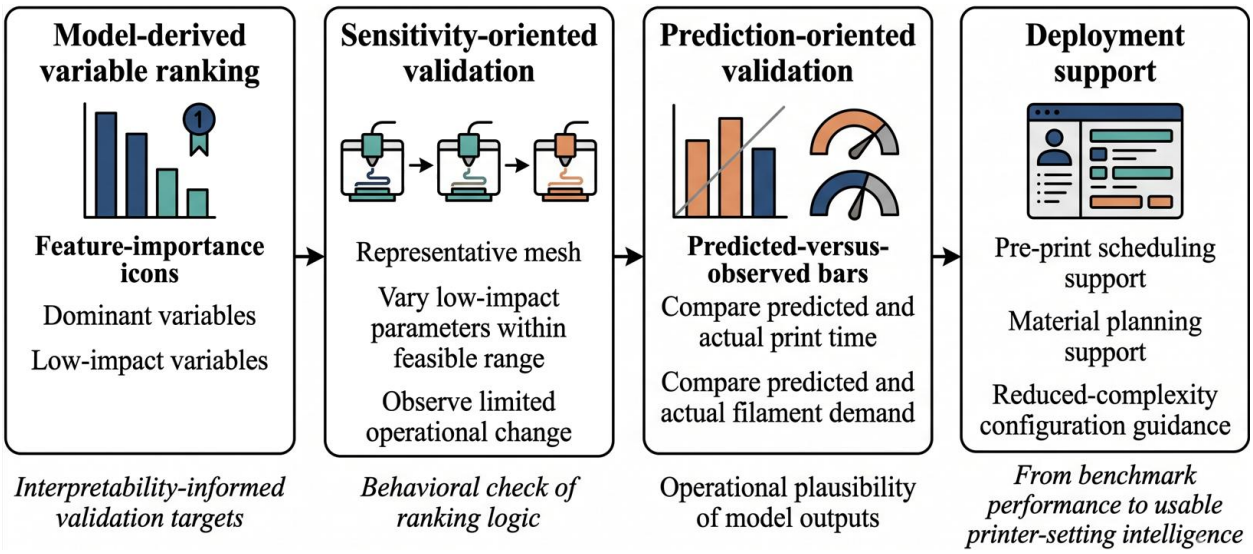


Figure 10. *Practical-validation and deployment-oriented evaluation design for geometry-aware printer-setting support*

The practical validation design followed two complementary lines. The first line evaluated low-impact variables. Representative meshes were printed multiple times while varying parameters that the model later ranked as weakly influential. This type of test is important because it checks whether the interpretability layer is behaviorally plausible: if low-ranked variables are changed within feasible ranges, the resulting operational response should remain comparatively stable. The second line evaluated direct predictive usefulness. Representative print cases were used to compare predicted versus observed time and filament demand, thereby testing whether the framework remains sufficiently accurate for planning under real fabrication conditions. These two forms of validation serve different but connected roles. The first validates the logic of variable ranking; the second validates the usefulness of the predicted outputs themselves. Figure 10 reflects this dual design by separating "sensitivity-oriented validation" from "prediction-oriented validation" within the same practical evaluation map.

This deployment-oriented testing logic is particularly important for the current study because the claimed contribution is not just a predictive model, but a pathway from mesh geometry to printer-setting intelligence. A purely statistical benchmark could show that one learner predicts well, but it would not by itself justify claims about simplified printer configuration or operational guidance. The practical validation stage fills that gap. It tests whether the model's internal structure—especially the distinction between dominant variables and marginal variables—corresponds to observable printing behavior. It also tests whether the numerical predictions remain close enough to reality to support scheduling and material planning. In effect, **Figure 10** demonstrates that the experimental design of the study closes the loop from data to model to deployment. This is one of the strongest aspects of the paper's design: the evaluation protocol is not restricted to abstract machine-learning performance, but extends to the conditions under which a real user might rely on the framework before launching a print job.

Taken together, the four components of the experimental design establish a coherent evaluation framework. The training protocol ensures disciplined model development, the benchmark structure allows fair comparison across ensemble and neural alternatives, the metric hierarchy supports multi-perspective performance interpretation, and the practical validation stage

connects the predictive framework back to real FDM operation. This layered design is necessary because the paper does not seek merely to report a high-performing regressor. It seeks to demonstrate that geometry-aware interpretable learning can serve as a scientifically grounded and practically deployable foundation for pre-print decision support in fused deposition modeling

7. Results

7.1. Global predictive benchmarking and comparative model behavior

The first objective of the results analysis was to determine whether the proposed geometry-aware learning framework could achieve sufficiently strong predictive performance to justify its use as an operational pre-print support system. This question is addressed at two levels: first, by comparing the three ensemble-based model configurations used in the present study, and second, by positioning the final stacking framework against the previously reported CNN–MLP model developed for related 3D-printing data. As summarized in **Table 1**, the benchmark among the three ensemble branches shows a clear pattern. Random Forest and the stacking framework occupied a closely matched low-error regime, with RMSLE values of 0.032052 and 0.032734, respectively, whereas Extreme Gradient Boosting produced a distinctly higher RMSLE of 0.057026. This result indicates that the prediction problem is indeed nonlinear, but not all nonlinear learners capture the geometry–process interaction equally well. In the present dataset, Random Forest provides the strongest base-level fit, while stacking preserves a nearly identical error level through model fusion.

Table 1. Global benchmark comparison of the principal predictive models

Model	RMSLE	MSE	MAE	MAD	RMSE	R^2
Random Forest	0.032052	—	—	—	—	—
Extreme Gradient Boosting	0.057026	—	—	—	—	—
Stacking framework	0.032734	0.0001929	0.0063829	0.0023750	0.01361	0.9619
CNN–MLP reference	—	0.0059000	0.0350000	0.0108700	—	—

This pattern becomes visually transparent in **Figure 11**, where the three models are arranged in a compact benchmarking plot centered on RMSLE. The figure is useful not because it merely repeats the numerical values of **Table 1**, but because it makes the comparative structure immediately legible. The close proximity of Random Forest and stacking confirms that the final fused framework does not sacrifice accuracy in order to gain interpretability and methodological structure. At the same time, the visibly larger separation between XGBoost and the other two models shows that the present problem does not simply reward boosting intensity; instead, it favors a more stable representation of the heterogeneous geometry–process feature space. In this sense, **Figure 11** does more than document numerical ranking. It clarifies the behavior of the predictive family and supports the later decision to use stacking as the central model for the interpretability-driven workflow.

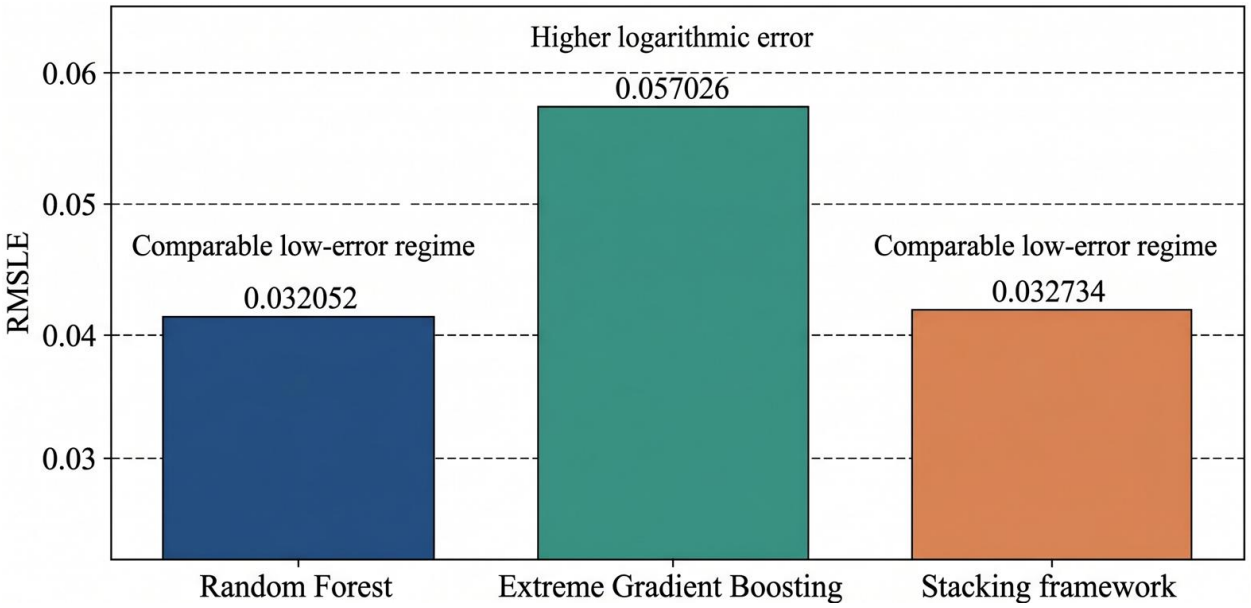


Figure 11. *RMSLE comparison of Random Forest, Extreme Gradient Boosting, and the stacking framework*

The second layer of the benchmark is even more important for the paper's narrative. As shown in **Figure 12** and listed in **Table 1**, the stacking framework substantially outperformed the CNN–MLP reference on every shared reported error measure. The multi-output MSE decreased from 0.0059 to 0.0001929, the MAE dropped from 0.0350 to 0.0063829, and the MAD fell from 0.01087 to 0.002375. The stacking framework also achieved an RMSE of 0.01361 and an $R^2$of 0.9619, two metrics that were not reported for the earlier CNN–MLP baseline. These differences are too large to be interpreted as marginal tuning effects. They indicate a substantive improvement in predictive fidelity, and, more importantly, they suggest that the explicit geometry–process representation used in the present framework is better aligned with the structure of the operational prediction problem than the earlier neural approach.

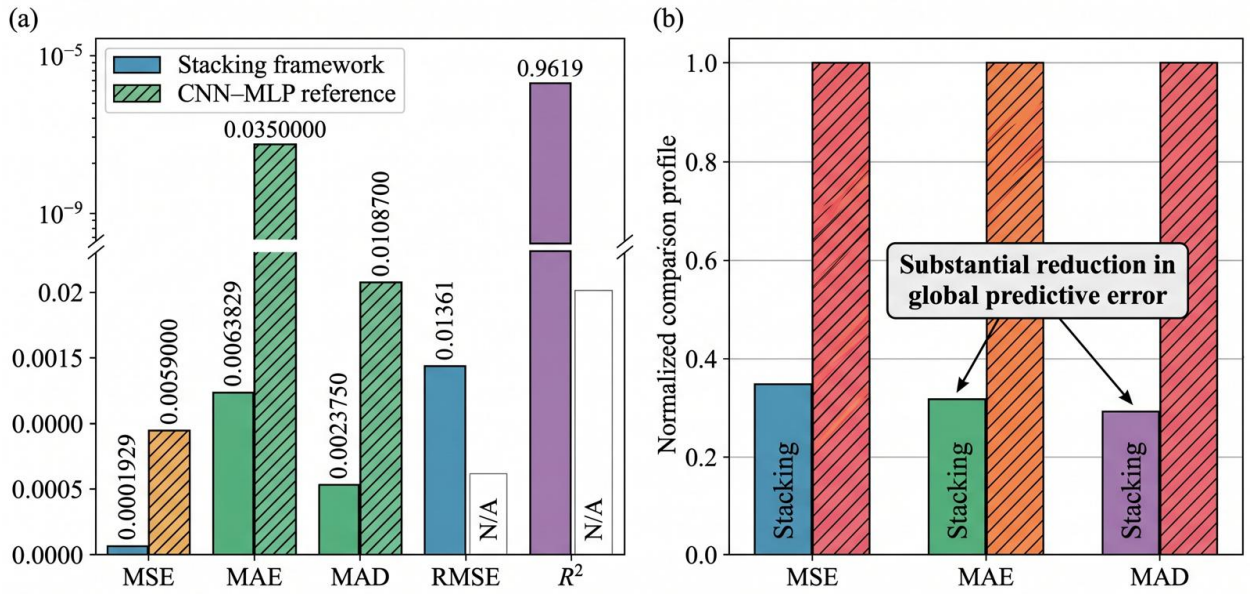


Figure 12. *Global error-profile comparison between the stacking framework and the CNN–MLP reference model*.

From an engineering standpoint, the significance of these results lies not only in the numerical gap itself, but in what that gap represents. The proposed framework does not rely on a more opaque or more computationally exotic learning design. Instead, it combines structured data conditioning, a geometry-aware feature space, ensemble learning, and stacking-based fusion. The fact that such a framework can outperform the earlier CNN–MLP model while remaining more interpretable directly supports the central claim of the manuscript: in fused deposition modeling, prediction becomes more useful when it is grounded in a physically meaningful representation of geometry–process interaction. Accordingly, the present section establishes a first key conclusion of the results: the proposed framework is not merely competitive; it is sufficiently accurate and sufficiently structured to serve as the predictive core of an interpretable printer-setting intelligence system.

7.2. Output-specific fidelity and balance across print time and filament prediction

While the global benchmark establishes the overall superiority of the proposed framework, the practical utility of the model depends on whether this performance remains balanced across the two operational targets. In fused deposition modeling, print time and filament consumption support different but equally important planning decisions. A model that is accurate only for one of these outputs would still be operationally incomplete. For this reason, the next layer of the results focuses on output-specific predictive behavior, and the corresponding evidence is condensed in **Table 2** and visualized in **Figures 13** and **14**. The numbers show that the stacking framework maintained strong prediction quality for both targets, but not in exactly the same way. For print time, the model achieved MSE = 0.00025, MAE = 0.00780, MAD = 0.00291, RMSE = 0.0158, and $R^2 = 0.9645$. For filament length, the corresponding values were MSE = 0.000169, MAE = 0.00571, MAD = 0.00125, RMSE = 0.0130, and $R^2 = 0.9449$.

Table 2. Output-specific predictive performance of the stacking framework and the CNN–MLP reference

Output	Model	MSE	MAE	MAD	RMSE	R^2
Print time	Stacking framework	0.000250	0.00780	0.00291	0.0158	0.9645
Print time	CNN–MLP reference	0.004880	0.02596	0.00710	—	—
Filament length	Stacking framework	0.000169	0.00571	0.00125	0.0130	0.9449
Filament length	CNN–MLP reference	0.007400	0.03900	0.01160	—	—

These values indicate that the framework does not merely perform well in aggregate; it remains stable and accurate when the two outputs are examined independently. **Figure 13** is particularly valuable here because it places the two output profiles side by side within a common metric space. The figure shows that the model explains slightly more variance for print time than for filament length, as evidenced by the higher R^2, while the absolute error burden is marginally lower for filament length on RMSE and MSE. This subtle asymmetry is scientifically meaningful. It suggests that the geometry–process relationship is highly learnable for both targets, but the two

outputs retain different sensitivity structures. Print time appears slightly more regular at the explained-variance level, whereas filament length benefits from a more concentrated absolute error profile. Rather than indicating inconsistency, this pattern confirms that the two targets should indeed be treated as distinct operational responses within a shared multi-output architecture.

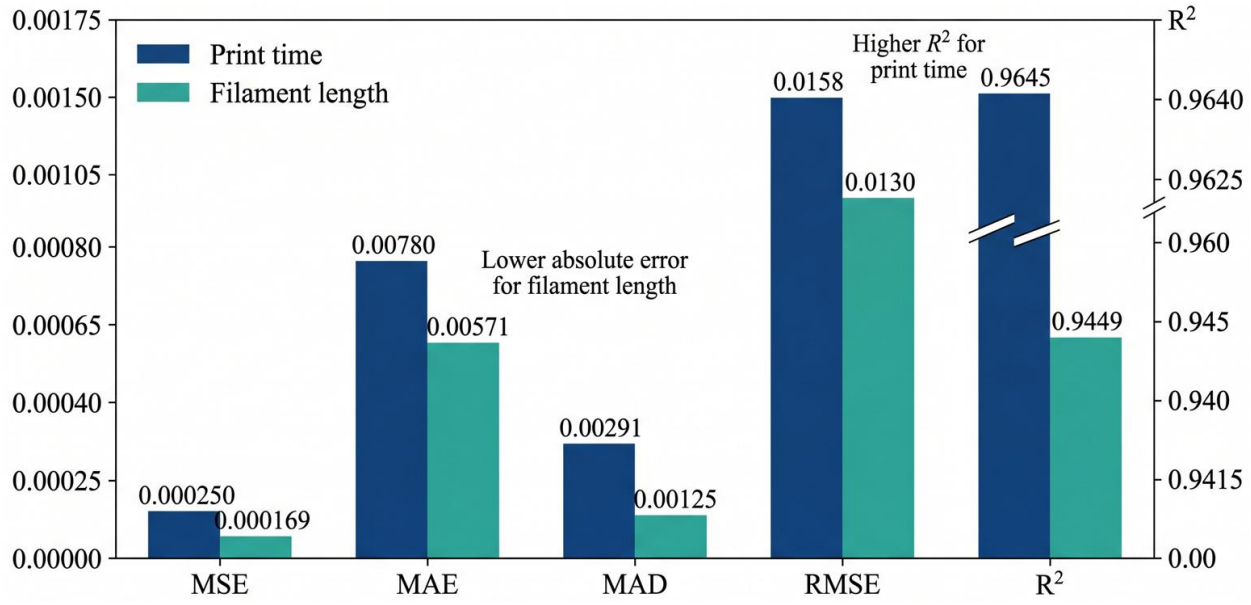


Figure 13. *Output-specific error comparison for print time and filament length under the stacking framework.*

The comparative significance of this result becomes clearer when the stacked framework is again placed against the CNN–MLP reference. As shown in **Figure 14** and listed in **Table 2**, the performance improvement holds for each target independently. For print time, the stacking model reduced MSE from 0.00488 to 0.00025 and MAE from 0.02596 to 0.00780. For filament length, the reductions were even more pronounced, with MSE dropping from 0.00740 to 0.000169 and MAE from 0.03900 to 0.00571. This target-specific comparison is crucial because it shows that the advantage of the present framework is not driven by only one output or by an averaging artifact. Instead, the gain is systematic across the two most operationally relevant responses retained in the study.

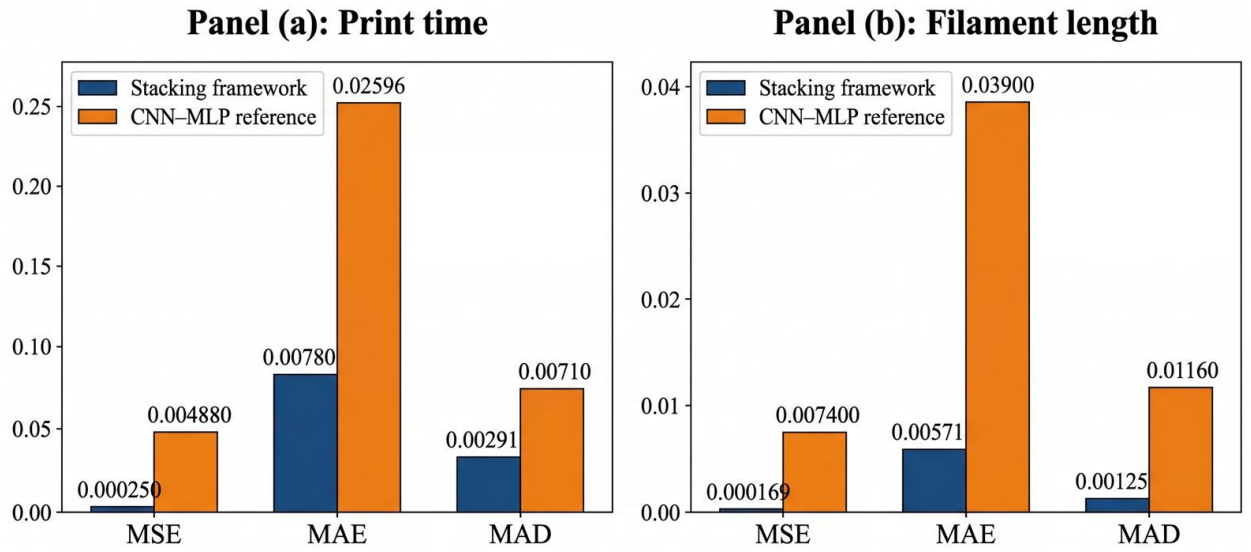


Figure 14. *Per-output comparison between the stacking framework and the CNN–MLP reference model.*

From the standpoint of deployment, this section of the results carries a direct implication: the model is sufficiently balanced to support both scheduling and material planning within the same predictive interface. This is a major strength of the manuscript's methodological direction. Many additive-manufacturing studies achieve good prediction for one task but do not establish whether the same framework can support multiple operational decisions simultaneously. Here, the target-specific results show that geometry-aware interpretable learning does not force a trade-off between the two main outputs of interest. Instead, it preserves high predictive quality for both, thereby justifying the later transition from model evaluation to configuration-support logic.

7.3. Interpretable feature hierarchy and the geometry–process contribution structure

The predictive strength of the framework becomes significantly more valuable when examined through the lens of interpretability. The original manuscript reports ranked feature-importance structures for both filament length and print time, and these rankings reveal the scientific core of the present study: operational behavior in fused deposition modeling is driven primarily by mesh geometry, while only a smaller subset of controllable settings provides meaningful leverage over the final outputs. This broader pattern is summarized concisely in **Table 3**, where the cumulative importance of geometry-derived descriptors and process-setting variables is aggregated for the two targets. For filament length, the seven geometry variables account for a cumulative importance of 0.7258, leaving 0.2742 for the process settings. For print time, the geometry contribution is even larger at 0.7987, while the process settings account for 0.2013. These values demonstrate that the proposed geometry-aware formulation is not merely conceptually reasonable; it is directly supported by the model's internal relevance structure.

Table 3. Group-level and top-feature contribution structure across the two operational targets

Output target	Geometry cumulative	Process cumulative	Top geometry drivers	Top controllable drivers

Filament length	0.7258	0.2742	Volume, Area, Vertices	Speed, Infill, Top solid layers
Print time	0.7987	0.2013	Area, Volume, Width	Skirt outlines, Layer height, Extrusion width

This group-level pattern is visualized in **Figure 15**, where the two input families are plotted against the two operational targets. The figure immediately shows that geometry dominates both responses, but it also makes clear that the process-setting branch remains non-negligible. That observation is important because it preserves the practical relevance of configuration support. If geometry explained nearly everything, then printer-setting guidance would contribute little beyond description. If process settings dominated, then the paper's emphasis on mesh descriptors would be weakened. The actual outcome is more informative than either extreme. Geometry defines the dominant structural envelope of the problem, while the process settings provide a smaller but still meaningful control surface through which users can influence operational efficiency. **Figure 15** therefore serves as a bridge between prediction and action: it shows why the model must be geometry-aware, but also why interpretability-driven configuration support remains feasible.

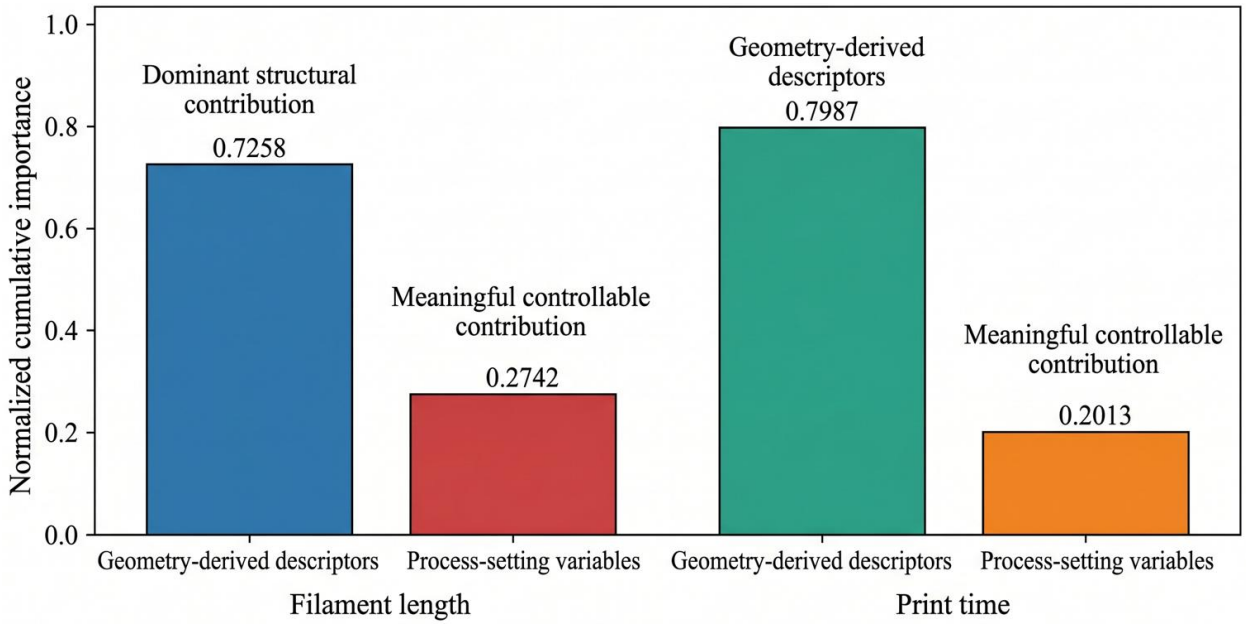


Figure 15. *Cumulative contribution of geometry-derived and process-setting feature groups to filament-length and print-time prediction*

The finer structure of this result is revealed in **Figure 16**, which ranks the individual variables for the two outputs. For filament length, the top three features are Volume (0.2805), Area (0.2411), and Speed (0.1669), followed by Vertices (0.0680) and Height (0.0390). This ranking is highly informative because it shows that filament demand is fundamentally governed by geometric burden, but still responds strongly to at least one controllable setting, namely speed. For print time, the leading variables are Area (0.5078), Volume (0.1125), and Skirt outlines (0.0816), followed by Width (0.0566) and Depth (0.0437). Here, the dominance of geometry is even more pronounced, especially through the very large contribution of area, but the presence of skirt outlines, layer

height, and extrusion width among the more influential controllable parameters reveals where meaningful timing control remains possible.

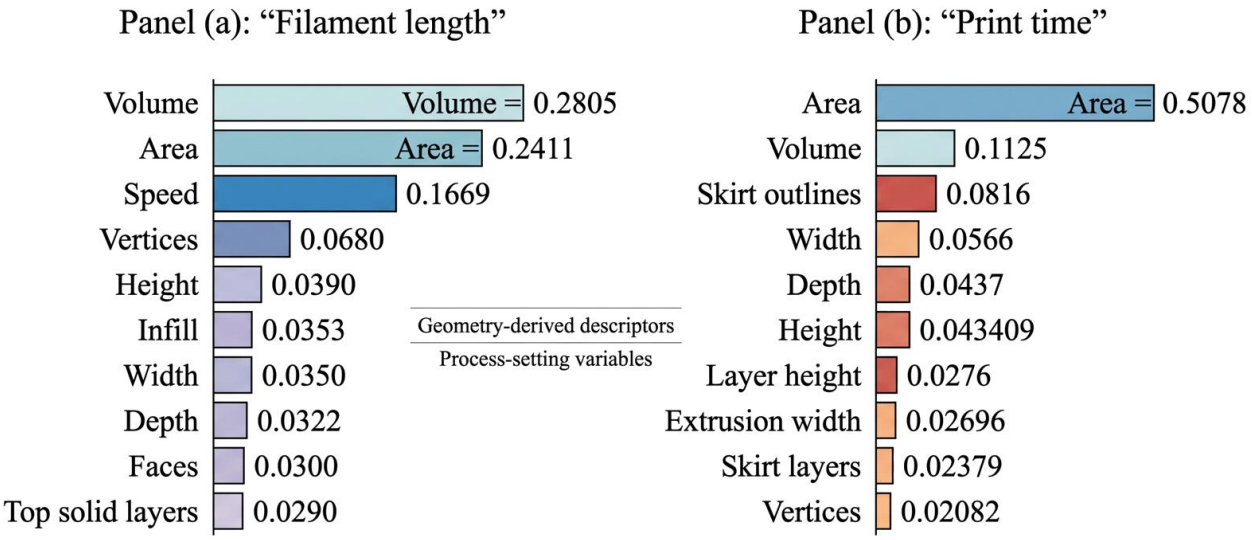


Figure 16. *Ranked feature-importance structure for filament-length and print-time prediction*.

These rankings have direct methodological and practical implications. Methodologically, they validate the paper’s decision to treat the problem as a fused geometry–process learning task rather than as parameter tuning alone. Practically, they justify a reduced-complexity configuration philosophy. The results indicate that users do not need to optimize every available printer setting to gain meaningful operational control. Instead, a smaller set of settings can be prioritized according to the predicted target of interest, while many low-impact variables can be left at default or expert-recommended values. This is precisely where the present study moves beyond conventional prediction. The feature-importance structure does not merely explain the model after the fact; it provides the empirical basis for a simplified, more deployable form of printer-setting intelligence.

7.4. Practical validation and deployment-oriented interpretation

The final and most application-oriented layer of the results concerns practical validation. A predictive framework intended for pre-print support must do more than produce low statistical error on held-out data; it must also remain behaviorally plausible when transferred to representative print scenarios. The source manuscript provides two complementary forms of practical evidence. The first is based on sensitivity-oriented experiments using low-impact variables. The second is based on direct predicted-versus-observed comparisons for representative cases of filament demand and printing time. These elements are summarized in **Table 4** and visualized in **Figures 17 and 18**. Together, they provide the final link between predictive modeling, feature ranking, and deployment-oriented printer-setting guidance.

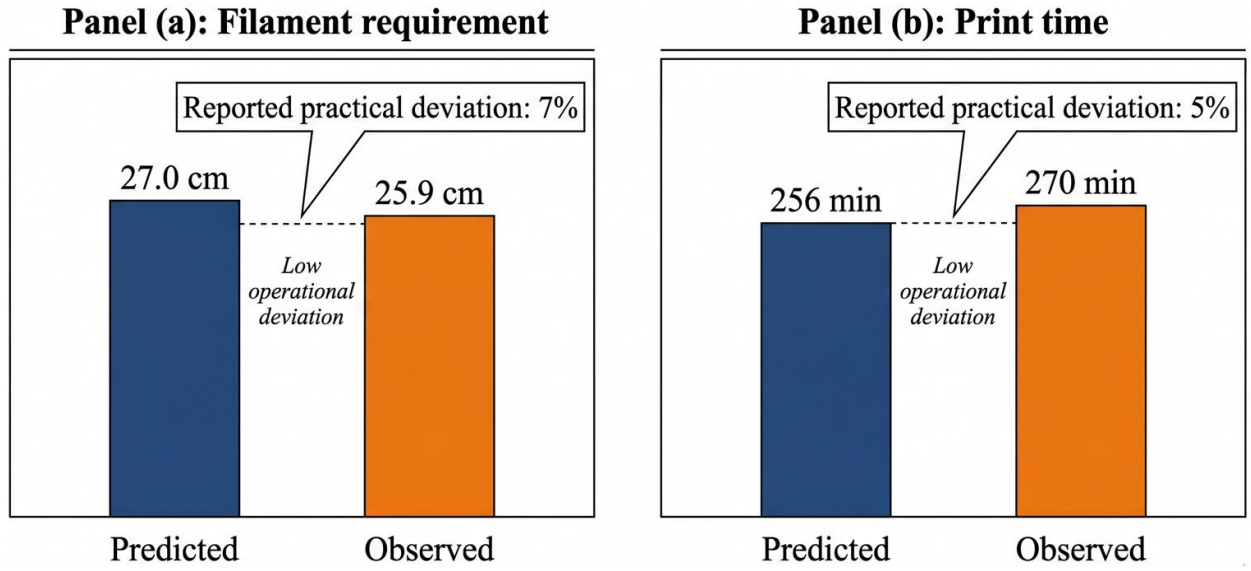


Figure 17. *Predicted versus observed values for representative practical validation cases of filament requirement and print time*

Table 4. Representative practical validation cases for operational interpretation

Validation type	Predicted value	Observed value	Reported deviation	Practical implication
Filament requirement	27.0 cm	25.9 cm	7%	Material demand remained close enough for pre-print allocation
Print time	256 min	270 min	5%	Scheduling estimate remained operationally reliable

The first part of this evidence concerns low-impact parameters. For filament length, the manuscript reports that when a randomly selected mesh was printed four times while varying layer height and skirt layers within permissible thresholds, the resulting differences in material usage and print quality were hardly discernible. Likewise, when another mesh was printed four times while holding those low-impact variables fixed and varying the remaining settings, the output quality remained similar, whereas thicker infill increased material consumption. For print time, an analogous experiment was conducted in which bottom solid layers and internal fill pattern were varied while other conditions were held fixed, and the printing time remained almost unchanged, consistent with their low importance ranking. These observations are depicted conceptually in **Figure 18**, which translates the feature-importance results into experimentally observed behavioral stability. The figure is not a substitute for the numerical tables; rather, it clarifies the practical meaning of low-ranked variables and shows why they can be treated as candidates for default-value control in a reduced-complexity configuration interface.

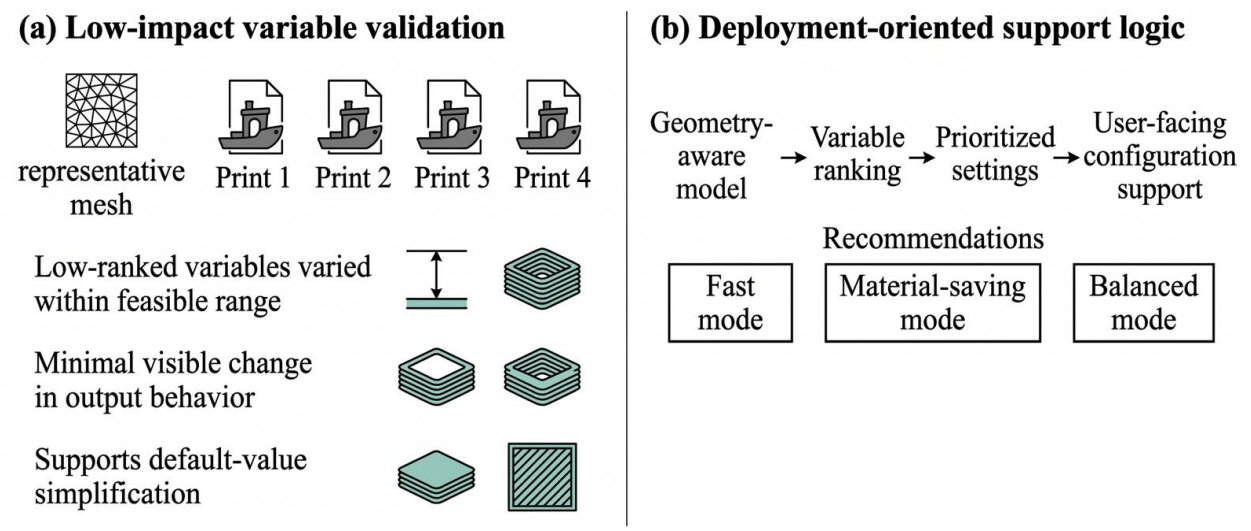


Figure 18. *Practical interpretation of feature-ranking validity and deployment-oriented printer-setting support.*

The second part of the validation is direct prediction realism. As shown in **Table 4** and plotted in **Figure 17**, the framework predicted a filament requirement of 27.0 cm for a case whose actual observed value was 25.9 cm, corresponding to a practical error reported in the original manuscript as 7%. For print time, the model predicted 256 min for a job that required 270 min in practice, corresponding to a reported error of 5%. Whether these deviations are read in relative or rounded practical terms, the key point is the same: the model remains close enough to observed fabrication behavior to support pre-print scheduling and material planning. This matters because the purpose of the framework is not exact physical emulation of every print event. Its purpose is to provide reliable operational intelligence before printing begins. **Figure 17** makes this point particularly effectively by showing the predicted and observed values side by side for each target, thereby making the magnitude of the gap immediately interpretable in practical terms.

Taken together, these practical results complete the logic of the study. The statistical results showed that the model predicts well; the feature-importance results showed that the geometry–process interaction can be interpreted; and the practical validation now shows that these two properties align with observable printing behavior. This alignment is what gives the present manuscript its strongest applied significance. The framework does not simply fit a structured dataset; it generates output estimates and variable rankings that remain plausible in real FDM operation. In that sense, the final result of the paper is not only an accurate predictor, but a credible basis for pre-print decision support. The combination of strong held-out performance, interpretable feature ranking, and practical validation suggests that geometry-aware interpretable learning can support more accessible and more efficient printer configuration, especially in settings where non-expert users need actionable guidance without having to navigate the full complexity of the FDM parameter space.

7. Discussion

7.1. Mesh geometry should be interpreted as the primary driver of operational printing burden

The most consequential outcome of this study is that mesh geometry emerged as the dominant determinant of both print time and filament consumption, while process settings acted

mainly as secondary modifiers of that burden. This result materially repositions how efficient fused deposition modeling should be interpreted. Much of the established FDM literature has focused on process-centric optimization, where variables such as layer height, print speed, extrusion temperature, and infill are treated as the main levers for improving manufacturing outcomes. That line of research remains valuable, especially for optimizing mechanical properties or surface quality under controlled specimen conditions, as shown by Dev and Srivastava [27], Feng et al. [47], and Rahmani Dabbagh et al. [48]. However, the present results indicate that for pre-print operational outcomes—specifically time and material demand—the structural burden imposed by the mesh itself governs most of the predictive space. In other words, the geometry of the part is not passive context; it is the primary source from which the operational response emerges.

This interpretation is consistent with the broader shift in the additive-manufacturing literature toward more integrated views of data, process, and design. Deka and Hall [33] showed that predictive performance for FDM parts improves when processing and environmental factors are modeled together, rather than treated through narrow local heuristics. More broadly, Nikooharf et al. [32] reviewed machine learning in polymer additive manufacturing as a domain in which design, process optimization, monitoring, and property control are increasingly interdependent, while Zhang et al. [23] similarly described machine-learning-driven 3D printing as an ecosystem that spans process optimization, monitoring, and motion planning rather than isolated tuning tasks. The present study extends that line of reasoning in an operational direction. It shows that when the target shifts from post-print properties to pre-print planning, geometry becomes even more central, because it defines the amount of path generation, deposition demand, and execution burden that the process settings subsequently modulate.

This also explains why parameter-only optimization studies, although methodologically strong, often do not translate directly into printer-setting support for diverse real-world jobs. Under fixed or narrowly varying geometries, the apparent influence of process variables can be magnified because the geometric burden is effectively held constant. Once geometry is allowed to vary across parts, that interpretation changes. The current results suggest that many settings remain relevant, but their influence must be read within a geometry-defined operational envelope rather than as universally dominant controls. This is one of the core scientific distinctions of the present work: it does not reject process optimization, but reframes it inside a geometry-aware structure that more faithfully represents how actual FDM print jobs behave.

7.2. The contribution of the framework lies in shifting from parameter tuning to operational inference

A second important implication of the results is that the study advances beyond conventional parameter tuning and toward operational inference. Prior work has already shown that machine learning can improve additive-manufacturing prediction and optimization. Nguyen et al. [18] proposed a data-driven learning approach for 3D-printing process optimization, while Wang et al. [28] demonstrated that machine learning can even infer parameter settings for desired mechanical properties. These studies are important because they move the field away from pure trial-and-error and toward learned relationships between inputs and manufacturing outcomes. Yet much of that literature remains centered on performance prediction for mechanical or quality-related endpoints. By contrast, the present study reorients the problem around two outputs that are often more immediately decisive in real printing environments: print time and filament consumption. This shift matters because scheduling and material planning are the first operational decisions that a user must make before fabrication begins.

The current results further show that this operational framing is not merely a change in target selection; it changes the structure of the learning problem itself. The strong target-specific performance of the stacking framework indicates that print time and filament demand are best learned as related but non-identical outputs of the same geometry–process state. This gives methodological support to the multi-output design of the framework. In many additive-manufacturing studies, outputs are modeled individually, which can obscure the fact that shared input structure underlies different operational responses. Here, the ability to maintain strong performance on both targets suggests that a joint formulation is more appropriate for pre-print decision support. The study therefore contributes not only a more accurate model than the earlier CNN–MLP reference, but also a more coherent representation of the problem: the model does not simply predict isolated quantities, it learns the operational behavior of a print job as a coupled system.

This distinction is also visible when the present study is placed against broader reviews of ML in additive manufacturing. Nikooharf et al. [32], Ukwaththa et al. [49], and Zhang et al. [23] all argue, in different ways, that the field is moving from narrow task-specific prediction toward integrated intelligent workflows that combine design, optimization, monitoring, and explanation. The present study fits that trajectory, but contributes a more focused and deployable interpretation: geometry-aware learning becomes most useful when it is linked directly to pre-print operational intelligence. That is why the present framework should not be read simply as another benchmark exercise between Random Forest, XGBoost, and stacking. Its stronger contribution is that it redefines the role of machine learning in FDM from "finding a better predictor" to "supporting earlier and better printing decisions."

7.3. Interpretability is the mechanism that turns prediction into deployable printer-setting support

The third and perhaps most distinctive implication of the study lies in interpretability. A large portion of recent additive-manufacturing research has sought better prediction, but fewer studies have translated model structure into simplified configuration logic. This is where the present work becomes especially relevant. The feature hierarchy reported in the results shows that once geometry is properly represented, only a limited subset of process settings consistently acts as meaningful operational levers. That finding has immediate practical value because it implies that users do not need to tune the full printer parameter space with equal attention. Instead, configuration support can be reduced to the subset of variables that materially changes time and filament demand. In other words, interpretability here is not a decorative explanatory layer; it is the mechanism through which the model becomes operationally deployable.

This conclusion aligns well with the recent explainable-AI literature in additive manufacturing. Kharate et al. [41] showed that explainable AI can clarify the relationship between process parameters and material properties in FDM-based biocomposites, while Ukwaththa et al. [49] reviewed the broader emergence of machine learning and explainable artificial intelligence across additive-manufacturing applications and emphasized that XAI remains comparatively underused despite its practical value. The current study extends that direction in a way that is particularly useful for FDM operation. Rather than using interpretability only to explain a black-box model after prediction, it uses feature relevance to identify the subset of controllable parameters that deserves active attention. This creates a bridge from model explanation to configuration simplification, which is still relatively uncommon in the literature.

Seen from that perspective, the most original aspect of the present manuscript is not the isolated use of Random Forest, XGBoost, or stacking—none of these models is novel by itself. The novelty lies in how they are embedded in a framework that combines geometry-aware representation, multi-output operational prediction, and interpretability-driven configuration support. This integration is what differentiates the study from earlier parameter-optimization papers, from property-prediction models that stop at numerical output, and from general XAI reviews that do not reach the level of deployment logic. The results suggest that the next generation of support tools for fused deposition modeling should not aim only for higher predictive accuracy. They should aim to make that accuracy usable by translating geometry and process data into a smaller, clearer, and more trustworthy printer-setting space. That, ultimately, is the deeper academic significance of the framework proposed in this paper.

8. Conclusions

This study developed a geometry-aware interpretable learning framework for predicting print time and filament consumption in fused deposition modeling and, in doing so, repositioned printer configuration as an operational inference problem rather than a purely parameter-tuning exercise. By integrating mesh-derived descriptors with controllable printer-setting variables, the framework captured the coupled geometry–process structure that governs pre-print behavior and translated that structure into a multi-output predictive system with direct planning value. Built on 1,500 printing instances generated from 150 meshes and evaluated through Random Forest, Extreme Gradient Boosting, and stacking-based fusion, the proposed approach achieved strong predictive performance across both operational targets. The final stacking framework attained an overall MSE of 0.0001929, MAE of 0.0063829, RMSE of 0.01361, and $R^2$of 0.9619, while the output-specific R^2values reached 0.9645 for print time and 0.9449 for filament length. These results substantially outperformed the previously reported CNN–MLP reference and confirmed that an explicitly structured geometry–process representation is more suitable for this problem than a less interpretable neural baseline.

Beyond predictive accuracy, the study showed that the dominant drivers of operational behavior are primarily geometric. For filament demand, volume and area emerged as the strongest predictors, followed by speed as the most influential controllable factor. For print time, area was the single most important variable, followed by volume and a smaller subset of process settings led by skirt outlines, layer height, and extrusion width. This ranking demonstrates that efficient FDM behavior is governed first by the burden imposed by the printed part and only secondarily by machine settings. At the same time, the results make clear that printer configuration remains meaningful because a limited subset of parameters still provides operational leverage once geometry is properly represented. This is one of the central contributions of the paper: it identifies a reduced and interpretable control surface for printer configuration rather than treating all settings as equally important.

The practical validation results further strengthened the relevance of the framework. Representative real-print checks showed that the model remained within the sub-10% deviation range for both filament requirement and print time, indicating that the learned geometry–process mapping is sufficiently reliable for pre-print scheduling and material planning. This supports the broader claim that the value of machine learning in fused deposition modeling lies not only in improving numerical prediction, but also in transforming printing data into actionable decision support.

Overall, the present work establishes that geometry-aware interpretable multi-output learning can provide a practical and scientifically grounded foundation for next-generation FDM support systems. By unifying prediction, explanation, and reduced-complexity configuration guidance within a single framework, the study advances fused deposition modeling from heuristic setup toward more transparent, transferable, and deployable operational intelligence.

Funding Declaration

The authors received no financial support for the research, authorship, and/or publication of this article

Conflict of Interest

The authors confirm that the intellectual content of this publication is the result of our own efforts, and that all outside aid or funds have been acknowledged. There are no conflicts of interest declared by the authors.

Declaration of Interest statement

This is to certify that to the best of authors' knowledge, the content of this manuscript is original. The paper has not been submitted elsewhere nor has been published anywhere.

Authors confirm that the intellectual content of this paper is the original product of our work and all the assistance or funds from other sources have been acknowledged.

Originality Statement

I hereby declare that this submission is my own work and to the best of my knowledge it contains no materials previously published or written by another person, or substantial proportions of material. Any contribution made to the research by others, with whom I have worked in draft or elsewhere, is explicitly acknowledged in the draft.

Availability of Data and Materials

Data sharing is not applicable to this article as no datasets were generated or analysed during the current study.

References

1. Mohammad Hossein N, Mehdi S, Maziar A, et al. Machine learning in polymer additive manufacturing: a review. International Journal of Material Forming 2024;17(52, doi:10.1007/s12289-024-01854-8
2. Xia G, Shunxin Q, Xiao K, et al. Fused filament fabrication of polymer materials: A review of interlayer bond. Additive Manufacturing 2021;37(101658, doi:10.1016/j.addma.2020.101658

3. Yuwei T, Jiann Yang H, Jiaxin W, et al. A review on voids of 3D printed parts by fused filament fabrication. Journal of Materials Research and Technology 2021;15(4860-4879, doi:10.1016/j.jmrt.2021.10.108
4. Long W, Junhao J, Yingchao D, et al. Machine learning enabled 3D printing parameter settings for desired mechanical properties. Virtual and Physical Prototyping 2024;19(1):e2425825, doi:10.1080/17452759.2024.2425825
5. Dabbagh SR, Ozcan O, Tasoglu S. Machine learning-enabled optimization of extrusion-based 3D printing. Methods 2022;206(27-40, doi:10.1016/j.ymeth.2022.08.002
6. Arindam D, James H. Predictive modeling of mechanical properties for fused deposition modeling parts: A focus on processing and environmental parameters. Manufacturing Letters 2023;35(643-651, doi:10.1016/j.mfglet.2023.08.053
7. Daniela F, Daniela R, Raffaele C, Carola Esposito C. A Review of Polymer-Based Materials for Fused Filament Fabrication (FFF): Focus on Sustainability and Recycled Materials. Polymers 2022;14(3):465, doi:10.3390/polym14030465
8. Loganathan S, Wan Sharmiza Wan H, Ishak I, et al. A comprehensive review on fused deposition modelling of polylactic acid. Progress in Additive Manufacturing 2023;8(5):775-799, doi:10.1007/s40964-022-00356-w
9. Enemuoh EU, Asante-Okyere S. Impact of feature selection on neural network prediction of fused deposition modelling (FDM) print part properties. International Journal on Interactive Design and Manufacturing 2024;18(10):7413-7427, doi:10.1007/s12008-023-01598-w
10. Vicente CMS, Martins TS, Leite M, et al. Influence of fused deposition modeling parameters on the mechanical properties of ABS parts. Polymers for Advanced Technologies 2020;31(3):501-507, doi:10.1002/pat.4787
11. Popović M, Pjević M, Milovanović A, et al. Printing parameter optimization of PLA material concerning geometrical accuracy and tensile properties relative to FDM process productivity. Journal of Mechanical Science and Technology 2023;37(2):697-706, doi:10.1007/s12206-023-0113-6
12. Özen A, Auhl D, Völlmecke C, et al. Optimization of Manufacturing Parameters and Tensile Specimen Geometry for Fused Deposition Modeling (FDM) 3D-Printed PETG. Materials 2021;14(10):2556, doi:10.3390/ma14102556
13. Rahman MM, Sultana J, Rayhan SB, Ahmed A. Optimization of FDM manufacturing parameters for the compressive behavior of cubic lattice cores: an experimental approach by Taguchi method. The International Journal of Advanced Manufacturing Technology 2023;129(1329-1343, doi:10.1007/s00170-023-12342-9
14. Tura AD, Lemu HG, Mamo HB, Santhosh AJ. Prediction of tensile strength in fused deposition modeling process using artificial neural network and fuzzy logic. Progress in Additive Manufacturing 2023;8(3):529-539, doi:10.1007/s40964-022-00346-y
15. Schmidt C, Funk J, Griesbaum R, et al. Predicting Linear Dimensional Accuracy of Material Extrusion Parts in Dependence of Process Parameters Using Neural Networks Optimized by an Evolutionary Algorithm. 3D Printing and Additive Manufacturing 2025;12(1), doi:10.1089/3dp.2023.0192
16. Ziadia A, Habibi M, Kelouwani S. Machine Learning Study of the Effect of Process Parameters on Tensile Strength of FFF PLA and PLA-CF. Eng 2023;4(4):2741-2763, doi:10.3390/eng4040156

17. Nikzad MH, Heidari-Rarani M, Rasti R, Sareh P. Machine learning-driven prediction of tensile strength in 3D-printed PLA parts. Expert Systems with Applications 2025;264(125836, doi:10.1016/j.eswa.2024.125836
18. Nguyen PD, Nguyen TQ, Tao QB, et al. A data-driven machine learning approach for the 3D printing process optimisation. Virtual and Physical Prototyping 2022;17(4):768-786, doi:10.1080/17452759.2022.2068446
19. Xijun Z, Dianming C, Xinyue Z, et al. Machine learning-driven 3D printing: A review. Applied Materials Today 2024;39(102306, doi:10.1016/j.apmt.2024.102306
20. Oliveira RMA, Sant'Anna ÂMO, da Silva PHF. Explainable machine learning models for defects detection in industrial processes. Computers & Industrial Engineering 2024;192(110214, doi:10.1016/j.cie.2024.110214
21. Moosavi S, Farajzadeh-Zanjani M, Razavi-Far R, et al. Explainable AI in Manufacturing and Industrial Cyber-Physical Systems: A Survey. Electronics 2024;13(17):3497, doi:10.3390/electronics13173497
22. Chohan JS, Mittal N, Kumar R. Parametric optimization of fused deposition modeling using learning enthusiasm enabled teaching learning based algorithm. SN Applied Sciences 2020;2(12):1978, doi:10.1007/s42452-020-03818-4
23. Zhang X, Chu D, Zhao X, et al. Machine learning-driven 3D printing: A review. Applied Materials Today 2024;39(102306, doi:10.1016/j.apmt.2024.102306
24. Ogunsanya MA, Parupelli SK, Isichei JC. Predictive quality assurance of material extrusion process using artificial neural networks and ensemble models. The International Journal of Advanced Manufacturing Technology 2025;140(5-6):2691-2714, doi:10.1007/s00170-025-16359-0
25. Khusheef AS, Shahbazi M, Hashemi R. Predicting 3D printed plastic part properties: A deep learning approach with thermographic and vibration data fusion. Expert Systems with Applications 2024;255(PB):124605, doi:10.1016/j.eswa.2024.124605
26. Dereci U, Tuzkaya G. An explainable artificial intelligence model for predictive maintenance and spare parts optimization. Supply Chain Analytics 2024;8(100078, doi:10.1016/j.sca.2024.100078
27. Dev S, Srivastava R. Optimization of fused deposition modeling (FDM) process parameters for flexural strength. Materials Today: Proceedings 2021;44(3012-3016, doi:10.1016/j.matpr.2021.02.436
28. Wang L, Jiang J, Dong Y, et al. Machine learning enabled 3D printing parameter settings for desired mechanical properties. Virtual and Physical Prototyping 2024;19(1):e2425825, doi:10.1080/17452759.2024.2425825
29. Abeykoon C, Sri-Amphorn P, Fernando A. Optimization of fused deposition modeling parameters for improved PLA and ABS 3D printed structures. International Journal of Lightweight Materials and Manufacture 2020;3(3):284-297, doi:10.1016/j.ijlmm.2020.03.003
30. Chohan JS, Kumar R, Yadav A, et al. Optimization of FDM Printing Process Parameters on Surface Finish, Thickness, and Outer Dimension with ABS Polymer Specimens Using Taguchi Orthogonal Array and Genetic Algorithms. Mathematical Problems in Engineering 2022;2022(2698845, doi:10.1155/2022/2698845
31. Balasubramanian NK, Kothandaraman L, Sathish T, et al. Optimization of process parameters to minimize circularity error and surface roughness in fused deposition modelling (FDM) using Taguchi method for biomedical implant fabrication. Advanced Manufacturing: Polymer & Composites Science 2024;10(1):2406156, doi:10.1080/20550340.2024.2406156

32. Nikooharf MH, Shirinbayan M, Arabkoohi M, et al. Machine learning in polymer additive manufacturing: a review. International Journal of Material Forming 2024;17(52, doi:10.1007/s12289-024-01854-8
33. Deka A, Hall J. Predictive modeling of mechanical properties for fused deposition modeling parts: A focus on processing and environmental parameters. Manufacturing Letters 2023;35(643-651, doi:10.1016/j.mfglet.2023.08.053
34. Hsueh M-H, Lai C-J, Wang S-H, et al. Effect of Printing Parameters on the Thermal and Mechanical Properties of 3D-Printed PLA and PETG, Using Fused Deposition Modeling. Polymers 2021;13(11):1758, doi:10.3390/polym13111758
35. Kutnjak-Mravlinčić S, Godec D, Pilipović A, Sutlović A. Optimization of the FDM Processing Parameters on the Compressive Properties of ABS Objects for the Production of High-Heeled Shoes. Journal of Manufacturing and Materials Processing 2024;8(3):106, doi:10.3390/jmmp8030106
36. Liu Z, Mazzei Capote GA, Grubis E, et al. Predicting Properties of Fused Filament Fabrication Parts through Sensors and Machine Learning. Journal of Manufacturing and Materials Processing 2023;7(5):186, doi:10.3390/jmmp7050186
37. Patel KS, Trivedi N, Shah DB, Joshi SJ. Prediction of tensile strength using machine learning algorithms in fused deposition modeling. Proceedings of the Institution of Mechanical Engineers, Part E: Journal of Process Mechanical Engineering 2024, doi:10.1177/09544089241286428
38. Obaidat S, Tamimi MF, Mumani AA, Alkhaleel B. An artificial neural network-based predictive model for tensile behavior estimation under uncertainty for fused deposition modeling. Rapid Prototyping Journal 2024;30(10):2056-2070, doi:10.1108/RPJ-04-2024-0168
39. Tzionis G, Mouratidis P, Kougka G, et al. A review of explainable AI methods and their application in manufacturing systems. Discover Applied Sciences 2026;8(52, doi:10.1007/s42452-025-07908-z
40. Senoner J, Netland T, Feuerriegel S. Using Explainable Artificial Intelligence to Improve Process Quality: Evidence from Semiconductor Manufacturing. Management Science 2022;68(8):5704-5723, doi:10.1287/mnsc.2021.4190
41. Kharate N, Anerao P, Kulkarni A, Abdullah M. Explainable AI Techniques for Comprehensive Analysis of the Relationship between Process Parameters and Material Properties in FDM-Based 3D-Printed Biocomposites. Journal of Manufacturing and Materials Processing 2024;8(4):171, doi:10.3390/jmmp8040171
42. Di Bonito LP, Campanile L, Di Natale F, et al. eXplainable Artificial Intelligence in Process Engineering: Promises, Facts, and Current Limitations. Applied System Innovation 2024;7(6):121, doi:10.3390/asi7060121
43. Jeon JE, Hong SJ, Han S-S. Utilization of Machine Learning and Explainable Artificial Intelligence (XAI) for Fault Prediction and Diagnosis in Wafer Transfer Robot. Electronics 2024;13(22):4471, doi:10.3390/electronics13224471
44. Chen T-CT, Wang Y-C, Lin C-W. Enhancing the effectiveness of output projection in wafer fabrication using an Industry 4.0 and XAI approach. The International Journal of Advanced Manufacturing Technology 2024;134(1-2):113-125, doi:10.1007/s00170-024-14105-6
45. Lin C-Y, Gim J, Shotwell D, et al. Explainable artificial intelligence and multi-stage transfer learning for injection molding quality prediction. Journal of Intelligent Manufacturing 2025;36(5):3587-3606, doi:10.1007/s10845-024-02436-w

46. Gim J, Lin C-Y, Turng L-S. In-mold condition-centered and explainable artificial intelligence-based (IMC-XAI) process optimization for injection molding. Journal of Manufacturing Systems 2024;72(196-213, doi:10.1016/j.jmsy.2023.11.013
47. Feng Q, Maier W, Moehring H-C. Application of machine learning to optimize process parameters in fused deposition modeling of PEEK material. Procedia CIRP 2022;107(1-8, doi:10.1016/j.procir.2022.04.001
48. Rahmani Dabbagh S, Ozcan O, Tasoglu S. Machine learning-enabled optimization of extrusion-based 3D printing. Methods 2022;206(27-40, doi:10.1016/j.ymeth.2022.08.002
49. Ukwaththa J, Herath S, Meddage DPP. A review of machine learning (ML) and explainable artificial intelligence (XAI) methods in additive manufacturing (3D Printing). Materials Today Communications 2024;41(110294, doi:10.1016/j.mtcomm.2024.110294