Available online at www.sciencedirect.com

ScienceDirect

Procedia CIRP 139 (2026) 319–324

www.elsevier.com/locate/procedia

13th CIRP Global Web Conference (CIRPe 2025)

An Explainable AI-Guided Feature Refinement Framework for Surface Roughness Prediction in Robotic Drilling

Dongpeng Li[a,c], Fang Hao[b], Zheng Ma[b], Yugi Ji[b], Pai Zheng*[a], Weihua Li[d], Liang Liu[e], Shuo Chen[e]

[a]*Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Hong Kong SAR*
[b]*COMAC Shanghai Aircraft Manufacturing Co., Ltd, Shanghai, China*
[c]*Shien-Ming Wu School of Intelligent Engineering, South China University of Technology, Guangzhou, China*
[d]*School of Mechanical and Automotive Engineering, South China University of Technology, Guangzhou China*
[e]*Hong Kong Indics International Information Technology Co. Ltd, Hong Kong SAR*

* Corresponding author. *E-mail address:* pai.zheng@polyu.edu.hk

Abstract

In aircraft assembly, the low structural stiffness of industrial robots complicates surface quality control, and the opacity of conventional machine learning models hinders their adoption for process optimization. To address this, this paper presents a systematic Explainable AI-Guided Feature Refinement Framework to develop a minimal, yet robust, and physically interpretable model for surface roughness prediction in robotic drilling. The framework utilizes SHapley Additive exPlanations (SHAP) as an active component in an iterative feature selection process to refine a Random Forest model. The experimental validation on an integrated industrial platform demonstrates that this approach successfully identifies a minimal set of critical features from a high-dimensional dataset, including process parameters and specific vibration characteristics. The resulting model achieves superior predictive performance and stability compared to conventional feature selection methods. Furthermore, the analysis uncovers key non-linear relationships and feature interactions, providing interpretable insights into how operational parameters and dynamic responses collectively influence surface quality, which facilitates process optimization in robotic aircraft assembly.

Keywords: Robotic Drilling; Surface Roughness; Explainable Artificial Intelligence; Feature Selection; SHAP; Process Optimization

1. Introduction

In commercial aircraft manufacturing, the assembly of large structural components, such as fuselage and wing panels, requires drilling numerous high-quality holes. The stringent demands for manufacturing quality and operational efficiency necessitate exceptional precision in these machining operations [1, 2]. Traditional Computer Numerical Control (CNC) machines, despite their high precision, are constrained in processing large, complex aircraft components by limited workspace, high equipment costs, and restricted flexibility [3].Therefore, industrial robots are being explored as a flexible, cost-effective alternative due to their large workspace; however, their inherent low structural stiffness leads to process vibrations and compromised surface quality [4].

Accurate prediction of surface roughness is essential for ensuring component quality. Although data-driven approaches, particularly machine learning (ML) models like Artificial Neural Networks (ANN) and Random Forests, have shown high predictive accuracy [5, 6], their "black-box" nature presents a barrier to industrial adoption [7]. Engineers are often hesitant to trust these predictions due to the lack of transparency regarding their underlying rationale and limited insight into adjusting process parameters for improvement.

The emergence of Explainable Artificial Intelligence (XAI) has begun to address the issue of transparency, with methods such as SHapley Additive exPlanations (SHAP) providing post-hoc insights into the behavior of machine learning models [8, 9]. However, current XAI applications in manufacturing primarily focus on passive model explanation rather than active optimization. Recent advances can be categorized into feature contribution analysis [10] and root cause analysis [11, 12], with successful applications across various manufacturing domains including milling [13], drilling optimization [14], and quality

10.1016/j.procir.2025.09.041

assessment. Despite these advances, a significant research gap persists: no systematic framework exists for utilizing these explanations to actively refine feature sets, simplify model architecture, and construct minimal yet accurate predictive models [15].

To address this gap, this paper introduces an Explainable AI-Guided Feature Refinement Framework designed for automated feature selection and model optimization. By sythesizing cross-domain features and integrating the traditional ML prediction model with SHAP, the framework iteratively identifies and retains only the most influential process and vibration features. The objective is to automatically construct a physically interpretable prediction model with a minimal feature set, derived from high-dimensional sensor data. This ensures not only high predictive accuracy but also clear and actionable engineering relevance. The primary contributions of this work are as follows:

1. A novel Explainable AI-Guided Feature Refinement Framework is proposed that actively employs explainability for automated feature selection and model simplification to construct a minimal and physically interpretable model.
2. The framework is applied to identify and quantify the critical, non-linear effects of process parameters and vibration characteristics on surface roughness, providing deeper physical insight into the robotic drilling process.
3. The resulting interpretable, minimal-feature model provides actionable guidance for process optimization by facilitating the definition of stable operating windows and the identification of key dynamic features for quality control.

2. Proposed Framework

Figure 1 illustrates the proposed Explainable AI-guided feature refinement framework, which systematically optimizes models from raw data through three interconnected components: the cross-domain features from sensor signals and process parameters are extracted for model training and inference, then the SHAP analysis evaluates feature importance, offering interpretable insights into underlying physical mechanisms. Finally, an iterative refinement process uses these SHAP values to prune the least significant features, maintaining predictive performance while creating more lightweight model.

2.1. Cross-domain Feature Synthesis

The quality of the input features directly determines the performance ceiling of any data-driven model. This module aims to transform raw sensor signals into a meaningful and condensed set of numerical features. First, raw signals are preprocessed through band-pass filtering to remove noise and are segmented to isolate data corresponding to individual drilling operations. Following this, a comprehensive set of features is extracted from both the process parameters and the preprocessed

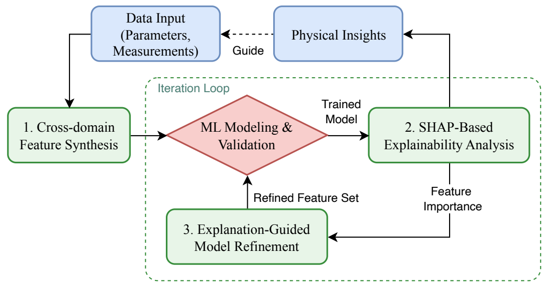

Fig. 1. The proposed Explainable AI-Guided Feature Refinement Framework for robotic machining analysis.

sensor signals to capture the full spectrum of process dynamics. As summarized in Table 1, these features are categorized into four main groups: direct Process Parameters (e.g., spindle speed, feed rate), Time-Domain Features to quantify statistical properties of the vibration signals, Frequency-Domain Features derived from the Fast Fourier Transform (FFT) to characterize dynamic responses, and Cross-Channel Features to capture system-level dynamics.

Table 1. Extracted features and corresponding notations

Category	Feature Description	Symbol
Process Parameters	Spindle Speed	n_s (rpm)
	Feed Rate	f (mm/min)
	Hole Number	N_{hole}
Time-Domain	Root Mean Square (RMS)	$\sigma_{RMS,c}$
	Kurtosis	κ_c
	Skewness	γ_c
	Peak Value	$x_{peak,c}$
	Crest Factor	CF_c
Frequency-Domain	Dominant Frequency	$f_{dom,c}$
	Spectral Centroid	$f_{c,c}$
	Spectral Rolloff	$f_{rolloff,c}$
	Band Energy (Low/Mid/High)	$E_{band,c}$
Cross-Channel	Correlation between channels	ρ_{ij}
	Coherence between channels	γ_{ij}^2

Note: Subscript c denotes the vibration channel (e.g., x, y, z). Subscripts i, j denote a pair of distinct channels.

2.2. SHAP-based Interpretability Analysis

This module serves as the explainability engine that opens the "black box" of machine learning models through SHAP (SHapley Additive exPlanations), a game-theoretic approach that decomposes model predictions into individual feature contributions. SHAP approximates model output with linear explanation model $g(x') = \phi_0 + \sum_{j=1}^{|F|} \phi_j x'_j$, where F is the feature set, x' is the binary input vector indicating feature presence, ϕ_j is the SHAP value for feature j, and ϕ_0 is the expected model output. The SHAP-based analysis provides insights essential for both feature selection and physical understanding. Feature importance is quantified by mean absolute SHAP value $I_j = \frac{1}{n}\sum_{k=1}^{n} |\phi_j^{(k)}|$, where $\phi_j^{(k)}$ is the SHAP value for feature j in sample k. This metric serves as the foundation for the iterative optimization algorithm by ranking features based on their actual contribution to model predictions rather than simple correlation measures. Global importance patterns are visualized through SHAP summary plots, while local interpretations emerge through SHAP dependence plots visualizing feature value x_j effects on corresponding SHAP value ϕ_j, revealing non-linear relationships. Feature interactions are assessed through SHAP interaction values ϕ_{ij}, quantifying synergistic or antagonistic effects between different process parameters and sensor measurements.

2.3. Explanation-Guided Model Refinement

Based on the SHAP-based feature importance quantification, this module implements an iterative refinement strategy to construct parsimonious yet accurate predictive models. While the framework is model-agnostic, it is particularly effective with high-performance, non-linear ensemble models such as Random Forest, XGBoost, and LightGBM, which have demonstrated superior performance on tabular data compared to deep learning alternatives [16, 17] and are robust at capturing complex data interactions.

Then the SHAP values are used to guide feature selection through a systematic pruning process, as detailed in Algorithm 1. The algorithm iteratively prunes features based on their SHAP importance scores while balancing selection aggressiveness through the retention factor α, maintaining predictive capability via minimum feature count n_{min}, and ensuring computational efficiency through iteration limits i_{max}. Guided by SHAP-based importance scores, this procedure systematically prunes the least influential features while maintaining predictive performance above a specified threshold. This XAI-driven approach surpasses conventional correlation-based filtering by utilizing the model's learned insights to identify the most salient variables for the specific prediction task.

Algorithm 1 SHAP-based Iterative Feature Optimization

Input: F_{full} (full feature set), (X, y) (training data), α (retention factor), n_{min} (min. features), i_{max} (max. iterations), δ (performance threshold)

Output: F_{opt} (optimized feature set)

$M_0 \leftarrow \text{Train}(F_{full}, X, y)$; $P_{best} \leftarrow \text{Perf}(M_0)$
$F_{current} \leftarrow F_{full}$; $F_{opt} \leftarrow F_{full}$; $i \leftarrow 1$
while $|F_{current}| > n_{min}$ **and** $i \leq i_{max}$ **do**
  $\Phi_i \leftarrow \text{SHAP}(M_{i-1}, F_{current})$
  $F_{ranked} \leftarrow \text{Sort}(F_{current}, |\Phi_i|)$ ▷ By $\mathbb{E}[|\phi_j|]$ desc.
  $F_{new} \leftarrow F_{ranked}[1 : \alpha \cdot |F_{current}|]$ ▷ Keep top α fraction
  $M_i \leftarrow \text{Train}(F_{new}, X, y)$; $P_{new} \leftarrow \text{Perf}(M_i)$
  if $P_{new} \geq P_{best} - \delta$ **then**
    $P_{best} \leftarrow P_{new}$; $F_{opt} \leftarrow F_{new}$
  else
    break
  end if
  $F_{current} \leftarrow F_{new}$; $i \leftarrow i + 1$
end while
return F_{opt}

3. Experimental Studies

To validate the proposed Explainable AI-Guided Feature Refinement Framework, robotic drilling experiments were conducted on an integrated industrial platform in the laboratory. This section presents the experimental methodology, platform configuration, and results, demonstrating the framework's effectiveness in achieving high predictive accuracy and meaningful physical interpretability.

3.1. Experimental Setup and Data Acquisition

The experimental validation was conducted on a fully integrated robotic drilling system designed for aircraft manufacturing. The key system components and their specifications are summarized in Table 2.

Table 2. Experimental platform components and specifications.

Component	Model/Specification
Industrial Robot	KUKA KR300 R2700
Electric Spindle	NAKANISHI BMS-4040
Drilling Tool	DCB-D060L09-D100WL30
Workpiece	7075 aluminum alloy (4mm)
Accelerometer	YMC-141A100
Data Acquisition	CY-E100
3D Camera	SICK Visionary-S
Laser Scanner	SICK RULER 3004

The system architecture employs a distributed control scheme where the robot controller manages motion trajectories while the PLC coordinates spindle operations and data acquisition through Ethernet protocols. The accelerometer was mounted directly on the spindle housing to capture vibration signatures at 32 kHz sampling rate, providing high-fidelity process monitoring capabilities essential for the proposed Explainable AI-Guided Feature Refinement Framework validation.

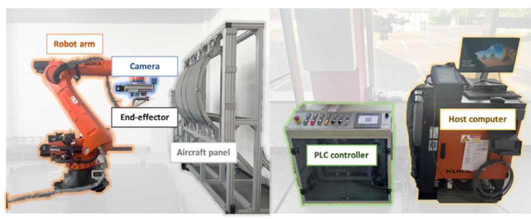

Fig. 2. Laboratory setup of the robotic drilling experimental platform showing the integrated system configuration.

The experimental parameters matrix covered spindle speeds from 2500 to 2800 rpm and feed rates from 100 to 160 mm/min, with 32 to 64 replications per parameter combination. Surface roughness (Ra) measurements were acquired using a TR300 profilometer, with final values calculated as the average of 4 measurements at different peripheral positions of the drilled hole wall following ISO 21920 standards. The description of the collected dataset are presented in Table 3, where only 1 tool is used for preliminary validation.

Table 3. Description of experiment dataset.

Name	Values
Train/Test samples (after augmentation)	345/87 (1035/261)
Spindle speed (rpm)	2,500, 2,600, 2,700, 2,800
Feed rate (mm/s)	1.0, 1.1, 1.2, 1.6
Surface Roughness (Ra, μm)	1.140–3.861
Mean ± Std of Ra (μm)	2.501 ± 0.681

3.2. Model Selection and Performance Validation

The collected dataset was randomly split into training and testing sets with a ratio of 80:20. Within the training set, 20% of the data was reserved for cross-validation to optimize hyperparameters and select the most suitable base model for subsequent interpretability analysis.

For the selection of the optimal base model for subsequent interpretability analysis, three tree-based algorithms particularly suitable for tabular data were evaluated: random forest, an ensemble method designed to reduce overfitting and improve generalization; XGBoost, a gradient boosting framework that sequentially builds models by correcting previous errors; and LightGBM, another gradient boosting framework optimized for efficiency and accuracy. To ensure optimal performance, the hyperparameters for each model were tuned using random search with 3-fold cross-validation on the training set, which can be done by scikit-learn's built-in functionalities.

3.3. Performance Evaluation Metrics

The predictive performance was evaluated using four key metrics: Mean Absolute Error (MAE), Root Mean Square Error (RMSE), Coefficient of Determination (R^2), and Mean Absolute Percentage Error (MAPE). These are defined as:

$$\mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{1a}$$

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{1b}$$

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2} \tag{1c}$$

$$\mathrm{MAPE} = \frac{100\%}{n}\sum_{i=1}^{n}\left|\frac{y_i - \hat{y}_i}{y_i}\right| \tag{1d}$$

where y_i represents the actual values, $\hat{y}_i$ represents the predicted values, $\bar{y}$ represents the mean of actual values, and n is the number of samples.

3.3.1. Model Performance Results

The performance comparison results are presented in Table 4 and Figure 3. The evaluation reveals that random forest achieves the best overall performance across all metrics, making it the optimal choice for the subsequent interpretability analysis.

Table 4. Performance comparison of tree-based models on the test set.

Model	MAE	RMSE	R^2	MAPE (%)
Random Forest	0.090	0.172	0.933	3.53
XGBoost	0.100	0.176	0.930	4.34
LightGBM	0.096	0.230	0.881	3.59

As demonstrated in both the tabular results and visual analysis, random forest outperforms the other models, exhibiting the lowest prediction errors and the highest R^2. The performance gaps reveal that XGBoost maintains relatively consistent accuracy with moderate increases in prediction errors, whereas LightGBM displays significantly degraded performance, especially in terms of prediction stability.

3.4. Validation and Explainability Analysis

3.4.1. Feature Selection Strategy Comparison

To comprehensively evaluate the proposed framework's effectiveness, three feature selection approaches were systematically compared across all three base models: random forest, XGBoost, and LightGBM. The strategies evaluated were: a baseline method using all features; a conventional Pearson correlation-based method; and the proposed iterative SHAP-based process for identifying a minimal-critical feature set. The evaluation protocol employed 10 different random seeds, which effectively creates diverse training-testing splits and sampling variations in the dataset. The comprehensive performance comparison is presented in Table 5.

Table 5. Performance comparison across models and feature selection approaches.

Model	Approach	R^2 Score
Random Forest	Baseline	0.936±0.002
	Correlation	0.892±0.029
	Proposed	**0.956±0.010**
XGBoost	Baseline	0.938±0.003
	Correlation	0.875±0.021
	Proposed	0.926±0.011
LightGBM	Baseline	0.884±0.003
	Correlation	0.858±0.013
	Proposed	0.921±0.012

Table 5 reveals that the proposed SHAP-based approach consistently outperforms both baseline and correlation-based methods across all models. While all models benefit from the SHAP-based feature selection, random forest demonstrates superior performance and stability, making it the optimal choice for detailed analysis. Given random forest's superior performance, the detailed computational and efficiency analysis focuses on this model, with results shown in Figure 4. For random forest specifically, the SHAP-based method achieves the highest median R^2 score of approximately 0.956 with remarkably consistent performance across multiple runs, significantly outperforming the correlation-based method which exhibits substantial variability and the baseline method. However, this su-

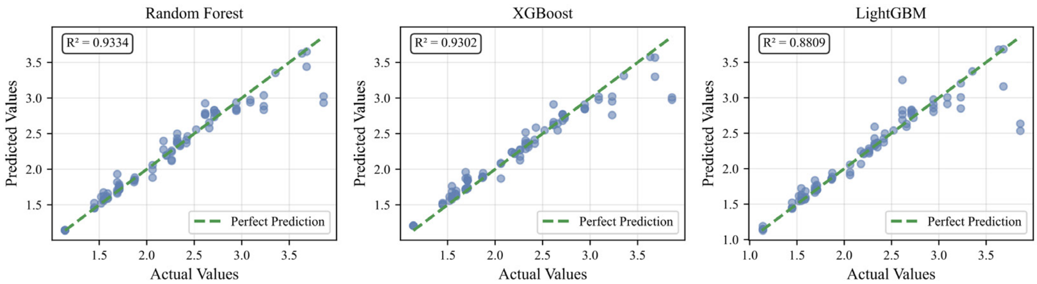

Fig. 3. Model performance comparison: scatter plots of predicted vs. actual values.

perior performance comes at computational cost, requiring approximately 135 seconds compared to the correlation-based approach and baseline method, representing a 7-fold and 19-fold increase respectively due to iterative SHAP value calculations—a critical consideration for real-time implementation scenarios.

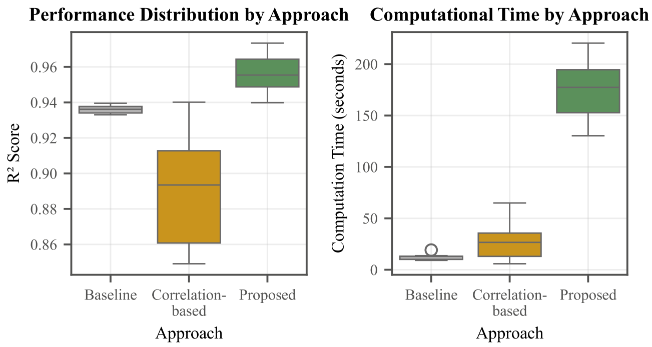

Fig. 4. Performance and computational comparison of feature selection approaches.

3.4.2. *Performance vs. Feature Efficiency Analysis*

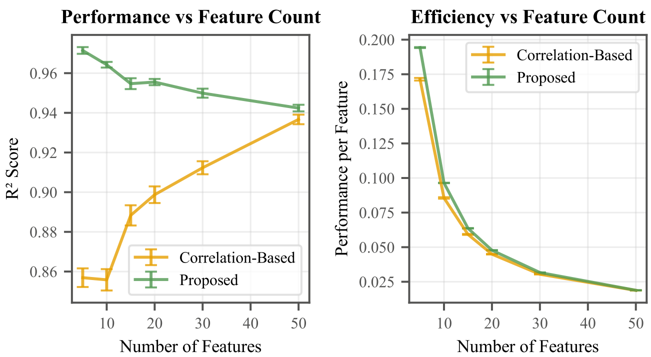

Fig. 5. Performance and efficiency analysis across varying feature counts.

Figure 5(a) highlights significant differences in how each method handles feature reduction for the random forest model. The proposed approach stands out, achieving an exceptional $R^2 > 0.96$ with just 5 features and maintaining consistently high performance ($R^2 > 0.94$) across all tested feature counts. This indicates its effectiveness in identifying critical features that capture underlying physical mechanisms. In contrast, the correlation-based method suffers noticeable performance drops at low feature counts, suggesting its difficulty in selecting impactful features when limited to smaller sets. Although the superior performance of our proposed method is most evident with the random forest model, similar patterns are observed with XGBoost and LightGBM models, as shown in Table 5, demonstrating the approach's model-agnostic nature.

The efficiency analysis in Figure 5(b) quantifies performance per feature, revealing that the proposed method maintains substantially higher efficiency at lower feature counts. This efficiency advantage diminishes as feature count increases, with all methods converging to similar efficiency levels at 50 features, emphasizing the particular strength of the SHAP-based approach in creating highly parsimonious models.

3.4.3. *SHAP-Based Interpretability Analysis*

The proposed method identified 5 critical features for surface roughness prediction, with the SHAP summary plot in Figure 6 displaying their importance in descending order: feed rate, local hole number, spectral centroid of x channel, energy ratio of high frequency band, and Spindle Speed. The analysis reveals that feed rate exhibits the highest importance with SHAP values spanning the entire impact range, while the scattered distribution of colored dots across all features indicates highly non-linear relationships that cannot be captured through simple correlation analysis. Notably, the bidirectional impact of feed rate—where higher values (red points) increase predicted roughness while lower values (blue points) reduce it—directly reflects the physical mechanism of increased chip load and cutting forces deteriorating surface quality. The spectral centroid of x channel reveals condition-dependent effects where high-frequency vibration signatures indicate chatter tendencies that compound with aggressive process parameters to severely degrade surface finish.

Further insights emerge from SHAP dependence plots in Figure 7, which reveal critical feature interactions beyond individual effects. The feed rate-spectral centroid interaction demonstrates a coupling where high-frequency vibrations (indicating chatter) significantly amplify the negative impact of aggressive feed rates, with red points in the upper right showing the most severe surface quality degradation. This analysis transforms the random forest from a black-box predictor into an interpretable tool that reveals how process parameters and

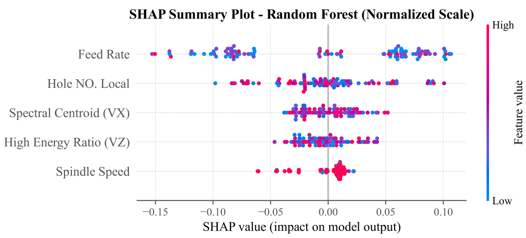

Fig. 6. SHAP summary plot for the proposed method showing global feature importance with only 5 features.

dynamic system responses jointly determine surface roughness outcomes.

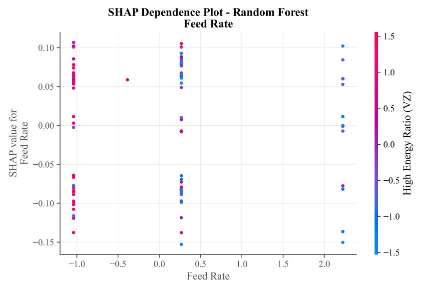

Fig. 7. Dependency analysis of feed rate v.s. spectral centroid (X channel).

Acknowledgment

This work was mainly supported by the funding support from COMAC International Collaborative Research Project (COMAC-SFGS-2023-3148), the National Natural Science Foundation of China (No. 52422514), the General Research Fund (GRF) (Project No. PolyU 15210222 and PolyU15206723) and the Collaborative Research Fund (CRF) (Project No. C6044-23GF) from the Research Grants Council (RGC), Hong Kong.

4. Conclusion

This paper proposed and validated an explainable AI-guided feature refinement framework to address the challenges of model prediction accuracy and interpretability in robotic drilling for surface roughness assessment. By integrating a base model with an iterative, SHAP-guided feature selection process, the proposed method moves beyond passive, post-hoc explanation toward an active feature refinement modeling approach. The experimental results indicate that the proposed framework can effectively distill a high-dimensional feature space into a minimal, critical set. The resulting models, not only achieved higher and more consistent predictive accuracy compared to baseline and correlation-based methods but also revealed clear patterns between process parameters (e.g., feed rate) and vibration characteristics (e.g., spectral centroid and energy ratio). Future validation could include systematic comparisons with advanced feature selection alternatives such as recursive feature elimination to further establish the superiority of the SHAP-based approach. The lightweight, minimal-feature model is a strong candidate for deployment in online monitoring and predictive quality control systems.

References

[1] A. Frommknecht, J. Kuehnle, I. Effenberger, S. Pidan, Multi-sensor measurement system for robotic drilling, Robotics and Computer-Integrated Manufacturing 47 (2017) 4–10, sI: FAIM 2015. doi:https://doi.org/10.1016/j.rcim.2017.01.002.

[2] P. Zheng, C. Li, J. Fan, L. Wang, A vision-language-guided and deep reinforcement learning-enabled approach for unstructured human-robot collaborative manufacturing task fulfilment, CIRP annals 73 (1) (2024) 341–344.

[3] S. Elhami, M. Razfar, M. Farahnakian, A. Rasti, Application of gonns to predict constrained optimum surface roughness in face milling of high-silicon austenitic stainless steel, The International Journal of Advanced Manufacturing Technology 66 (5) (2013) 975–986.

[4] D. Li, H. Du, W. S. Yip, Y. M. Tang, S. To, Online chatter detection for single-point diamond turning based on multidimensional cutting force fusion, Mechanical Systems and Signal Processing 206 (2024) 110850.

[5] D. Y. Pimenov, A. Bustillo, T. Mikolajczyk, Artificial intelligence for automatic prediction of required surface roughness by monitoring wear on face mill teeth, Journal of Intelligent Manufacturing 29 (5) (2018) 1045–1061.

[6] M. Hadad, S. Attarsharghi, M. Dehghanpour Abyaneh, P. Narimani, J. Makarian, A. Saberi, A. Alinaghizadeh, Exploring new parameters to advance surface roughness prediction in grinding processes for the enhancement of automated machining, Journal of Manufacturing and Materials Processing 8 (1) (2024) 41.

[7] L. Xia, P. Zheng, M. Herrera, Y. Liang, X. Li, L. Gao, Graph embedding-based bayesian network for fault isolation in complex equipment, IEEE Transactions on Reliability (2024) 1–14doi:10.1109/TR.2024.3416064.

[8] Z. Alexander, D. H. Chau, C. Saldaña, An interrogative survey of explainable ai in manufacturing, IEEE Transactions on Industrial Informatics 20 (5) (2024) 7069–7081.

[9] D. Li, S. Liu, B. Wang, C. Yu, P. Zheng, W. Li, Trustworthy AI for human-centric smart manufacturing: A survey, Journal of Manufacturing Systems 78 (2025) 308–327. doi:https://doi.org/10.1016/j.jmsy.2024.11.020.

[10] G. Weihong, V. Gawade, B. Zhang, Y. Guo, Explainable ai for layer-wise emission prediction in laser fusion, CIRP Annals 72 (1) (2023) 437–440.

[11] M. Calaon, T. Chen, G. Tosello, Integration of multimodal data and explainable artificial intelligence for root cause analysis in manufacturing processes, CIRP Annals 73 (1) (2024) 365–368.

[12] M. Calaon, H.-P. Yeh, S. Shan, J. H. Hattel, Y. Zhang, H. N. Hansen, Advancing quality prediction in polymer pbf-lb: a hybrid ai and physics-guided approach, CIRP Annals (2025).

[13] S. Zeng, D. Pi, Milling surface roughness prediction based on physics-informed machine learning, Sensors 23 (10) (2023) 4969.

[14] J. Qian, L. Deng, X. Zhang, S. Pi, Z. Song, X. Zhang, Xaip: An explainable ai-based pipeline for identifying key factors of surface defects in strip steel, steel research international 96 (3) (2025) 2400499.

[15] A. Puthanveettil Madathil, X. Luo, Q. Liu, C. Walker, R. Madarkar, Y. Qin, A review of explainable artificial intelligence in smart manufacturing, International Journal of Production Research (2025) 1–44.

[16] R. Shwartz-Ziv, A. Armon, Tabular data: Deep learning is not all you need, Information Fusion 81 (2022) 84–90. doi:https://doi.org/10.1016/j.inffus.2021.11.011.
URL https://www.sciencedirect.com/science/article/pii/S1566253521002360

[17] Y. Gorishniy, I. Rubachev, V. Khrulkov, A. Babenko, Revisiting deep learning models for tabular data, in: Proceedings of the 35th International Conference on Neural Information Processing Systems, NIPS ’21, Curran Associates Inc., Red Hook, NY, USA, 2021.