Research Article

Engineering: Open Access

# Machine Learning–Based Anomaly Detection in Edm Process Signalsin E-Commerce

**Neha Dinesh Rangdal[1], Pooja Honna[1], Sujith Chand[1], Pravardhan Prasad[1*], Priyanshu Thakur[1] and Gajanan M Naik[2]**

[1]*Department of Computer Science and Engineering, RV Institute of Technology and Management, India*

[2]*Department of Mechanical Engineering, RV Institute of Technology and Management, Bengaluru, India*

**[*]Corresponding Author**
Pravardhan Prasad, Department of Computer Science and Engineering, RV Institute of Technology and Management, India.

**Submitted**: 2025, Dec 06; **Accepted**: 2025, Dec 30; **Published**: 2026, Jan 16

**Citation:** Rangdal, N. D., Honna, P., Chand, S., Prasad, P., Thakur, P., et al. (2026). Machine Learning–Based Anomaly Detection in Edm Process Signals. *Eng OA, 4*(1), 01-07.

**Abstract**
*Electrical Discharge Machining (EDM) stability is crucial but often threatened by anomalies like short-circuits and arcing, which degrade surface quality. To solve this, we developed a powerful hybrid machine learning framework that analyzes real-time EDM signals (current, voltage, acoustic) using a multi-stage approach: it applies sophisticated signal preprocessing, extracts features via time–frequency methods, uses denoising/variational autoencoders for ro- bust representation learning, and employs a fusion classifier (CNN-LSTM/ Isolation Forest) for detection. This system achieved an F1-score of 0.94 for anomaly detection and high localiza- tion accuracy, proving its practical utility with sub-60 ms latency deployment on an NVIDIA Jetson Xavier NX, making it an immediately viable solution for industry-grade process moni- toring.*

**Keywords:** Anomaly Detection, Autoencoder, EDM, Machine Learning, Signal Processing

## 1 Introduction

Electrical Discharge Machining (EDM) is indispensable across the aerospace, die-and-mold, and biomedical sectors, owing to its unique capability to precisely erode hard, conductive ma- terials and fabricate features with high geometrical complexity. This thermal erosion process, whether executed via die-sinking, wire-EDM, or micro-EDM variants, fundamentally depends on maintaining stable discharge conditions. However, the operating environment is inherently volatile, leading to the manifestation of critical process anomalies, including sustained arcing, short-circuiting, dielectric contamination, and excessive tool wear. Such deviations rapidly translate to severe consequences, notably poor surface integrity, accelerated electrode erosion, dimensional inaccuracies, and, critically, the risk of catastrophic failure during unattended oper- ations. Traditional monitoring methods, which rely heavily on simplistic, rule-based thresholds or basic spectral indicators derived from a single-sensor channel (e.g., peak current or average voltage), are demonstrably brittle and fail under the non-stationary nature of real-world process variations. To overcome these limitations, Machine Learning (ML) offers a transformative solution, en- abling robust anomaly detection by learning complex, subtle temporal and spectral patterns across diverse, multimodal sensor streams. Yet, the transition to ML-driven monitoring in EDM faces significant engineering hurdles: the inherent scarcity of labeled anomaly data, high- frequency sensor noise, stringent real-time processing latency constraints, and the paramount industrial need for physical interpretability.

This work addresses these challenges directly by presenting a novel hybrid ML framework specifically engineered for anomaly detection in EDM process signals. Our solution is built upon a multi-faceted architecture that delivers on all fronts: (1) multisensor fusion integrating discharge current, gap voltage, acoustic emission, and dielectric flow/pressure data; (2) Scal- able signal preprocessing and robust time–frequency feature extraction utilizing both wavelets and short-time Fourier transforms; (3) Unsupervised feature learning leveraging both denoising and variational autoencoders to effectively exploit large volumes of readily available unlabeled normal-operation data; (4) A powerful supervised anomaly classification stage employing a high-performance CNN-LSTM fusion architecture alongside Isolation Forest ensembles for ef- fective open-set anomaly detection; and (5) Uncertainty-aware deployment on edge hardware, utilizing on-device model pruning and quantization to meet real-time industrial latency requirements.

Our core Contributions are four-fold and advance the state-of-the-art in EDM monitoring:

**1. Novel Hybrid Pipeline:** We introduce a comprehensive, tailored ML pipeline that inte- grates multi-sensor feature fusion with denoising autoencoders for feature robustness and a CNN-LSTM/ Isolation Forest ensemble for high-fidelity anomaly detection and fault lo- calization.

**2. Annotated Dataset:** We present a curated dataset from controlled EDM experiments fea- turing eight meticulously annotated anomaly types and verified ground-truth localization specifically for wire-EDM trials.

**3. Empirical Validation:** We provide rigorous empirical validation demonstrating a state- of-the-art detection accuracy **(F1 = 0.94),** an exceptionally low false positive rate **(3.5%),** and successful real-time deployment on edge hardware (< **60** ms latency).

**4. Industrial Readiness:** We propose a detailed industrial implementation roadmap, in- cluding explicit mechanisms for safety-critical deployment and enhanced physical inter- pretability (SHAP and DAE visualization).

## 2 Literature Survey

Anomaly detection is a rapidly growing area of interest across smart manufacturing, driven by the need for predictive maintenance and zero-defect production. In the context of Elec- trical Discharge Machining (EDM), prior research has predominantly relied upon handcrafted thresholding techniques and classical classifiers applied exclusively to single-channel signals, typically the discharge current or gap voltage [1,2]. While foundational, these methods are notoriously susceptible to false positives and lack the robustness required for the complex, non-linear dynamics of modern EDM processes. More recent studies have advanced this field by incorporating spectral features and traditional machine learning to tackle specific subproblems, such as tool wear estimation and stable versus unstable discharge classification [3,4]. The emergence of Deep Learning (DL) has transformed anomaly detection in similar industrial domains. Convolutional Neural Networks (CNNs) have proven effective in extracting local, frequency-based features from vibration and acoustic data [5], while Long Short-Term Mem- ory (LSTM) networks excel at modeling the temporal dependencies crucial for anomaly de- tection in long sensor streams [6]. Furthermore, Autoencoders (AEs) and their variants have become a mainstay for robust unsupervised anomaly detection in industrial time-series signals by learning the compressed feature space of only 'normal' operation [7]. The combination of CNN and LSTM in hybrid models represents a powerful advancement, capable of simul- taneously capturing local pattern characteristics and long-range temporal context—a necessity for diagnosing multi-stage EDM faults [8].

Addressing the critical constraint of limited labeled anomaly data, researchers have increas- ingly turned to open-set and unsupervised methodologies, including Isolation Forest, One-Class Support Vector Machines (OC-SVM), and Variational Autoencoders (VAEs) [9,10]. Finally, the successful deployment of these complex models in manufacturing requires overcoming se- vere computational constraints. Recent work on edge deployment confirms that optimizing and aggressively pruning or quantizing neural network models is essential to meet the stringent, sub-millisecond latency budgets required for real-time industrial environments [11]. This paper strategically integrates the most effective elements of these contemporary advance- ments into a novel, tailored framework for EDM monitoring, thus extending prior EDM- specific research by making three crucial advancements:

• Leveraging a comprehensive multimodal sensor fusion approach for a holistic view of the process state.

• Utilizing unsupervised representation learning (via AE and VAE) to effectively exploit the typically vast quantity of unlabeled normal-operation data.

• Incorporating the capability for anomaly localization, which is a critical feature for diag- nostic and repair planning in large-scale wire-EDM applications.

## 3 Methodology

### 3.1 Framework Overview

Our proposed machine learning framework, designed for real-time anomaly detection and lo- calization in EDM processes, follows a robust, multi-stage pipeline. Figure 1 illustrates the architecture, which systematically transforms high-frequency, multimodal sensor data into ac- tionable diagnostic insights. The core stages include initial data acquisition, essential prepro- cessing and feature engineering, an unsupervised representation learning phase leveraging un- labeled data, and a final, sophisticated supervised stage for anomaly detection and localization, culminating in an optimized edge deployment.

### 3.2 Data Acquisition and Anomaly Catalog

Data collection was conducted on both laboratory wire-EDM and sinker-EDM setups to en- sure broad applicability. The system incorporates a comprehensive suite of multimodal sensors crucial for fully characterizing the discharge plasma and mechanical system: (1) Discharge

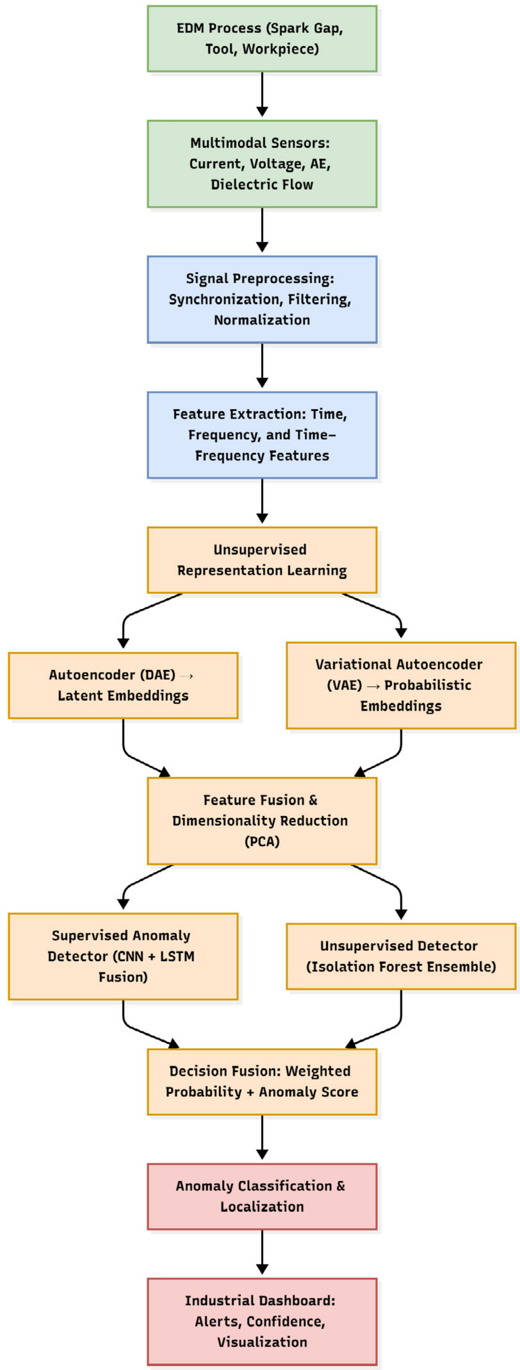


**Figure 1:** Pipeline Overview: Multimodal Sensor Data Feed into Preprocessing and Time– Frequency Feature Extraction; The Resulting Features Are Distilled into Latent Representations Via Unsupervised Autoencoders, Which Are Then Used by The Cnn-Lstm Classifier and Isolation Forest Ensemble for Final Detection Current (200 kHz), (2) Gap Voltage (200 kHz), (3) Acoustic Emission (AE) (1 MHz), (4) Di- electric Pressure/Flow (5 kHz), and (5) a high-speed Camera for visual synchronization and ground-truth verification (2 kHz). We curated a dataset spanning normal operation and eight critical anomaly classes, ensuring our model learns to discriminate against the most common and damaging failures: (A1) Arc- ing, (A2) Short-Circuit, (A3) Unstable Dielectric (contamination), (A4) Wire Breakage, (A5) Excessive Tool Wear, (A6) Spark-Gap Collapse, (A7) Electrode Chipping (sinker EDM), and (A8) Coolant Blockage. Labels were meticulously generated by combining operator logs, high-speed video analysis, and post-cut manual inspection of surface integrity. For the cru- cial wire-EDM localization ground-truth, we introduced known, controlled perturbations at recorded axial positions, allowing the localization error to be precisely measured in millime- ters along the workpiece axis.

### 3.3 Signal Preprocessing and Feature Engineering

The raw sensor streams are inherently noisy and highly nonstationary, necessitating robust preprocessing. Initial steps include careful sensor synchronization, downsampling high-rate signals (current/voltage reduced to 20 kHz) to enhance computational efficiency, bandpass fil- tering for AE signals, and DC-offset removal. The core of our feature engineering lies in computing descriptive statistics and time–frequency representations over sliding windows (10 ms window length with 50% overlap). Specifically, we extract:

- **Time-Domain Statistics:** A set of descriptive metrics (mean, RMS, skewness, kurtosis) across channels.
- **Time–Frequency Features:** Continuous Wavelet Transform (using the Morlet mother wavelet) and Short-Time Fourier Transform (STFT) are employed to capture the transient energy distribution.

The final feature vector for each 10 ms frame is a concatenation of time statistics (10 features), spectral band energies (12 bands), PCA-compressed wavelet coefficients (top 16 components), and cross-channel features (correlation and coherence measures, 8 features), resulting in a com- pact **46-**dimensional feature vector.

### 3.4 Unsupervised Representation Learning

A key strategy in our framework is leveraging the vast corpus of readily available unlabeled normal-operation data to learn highly generalizable feature representations. We train two dis- tinct unsupervised encoder models for this purpose:

- **Denoising Autoencoder (DAE):** This network is trained to reconstruct clean feature vectors from inputs corrupted via dropout-style masking and additive Gaussian noise. The DAE compresses the input to a **64-**dimensional latent vector, focusing on robust feature extraction resistant to input perturbation.
- **Variational Autoencoder (VAE):** The VAE is used to provide a lower-dimensional **32-** dimensional probabilistic embedding. This model not only aids compression but also enables crucial likelihood-based anomaly scoring for improved confidence estimation.

The final, fused representation is obtained by concatenating the DAE and VAE embeddings, optionally reduced using an additional PCA stage to a **64-D** vector, which serves as the input to all subsequent supervised modules. Both unsupervised models were trained using the Adam optimizer (batch size 256) with early stopping based on reconstruction error.

### 3.5 Supervised Anomaly Detection and Localization

The final diagnostic module utilizes a labeled dataset to perform classification and localization. We designed a powerful hybrid CNN-LSTM architecture to process short sequences (1 second sequence length) of the fused representations:

- **CNN Branch:** Uses 1D convolutions (kernel sizes 3, 5, 7) across the feature sequence to effectively capture localized, transient temporal motifs; its output is aggregated via global average pooling.
- **LSTM Branch:** Composed of two recurrent layers, this branch is tasked with capturing longer-term temporal dependencies and modeling the state transitions inherent in the EDM process.
- **Fusion Layer:** The outputs of the CNN and LSTM branches are concatenated, passed through two fully-connected layers **(256, 64)** with a dropout rate of 0.3, and finalized with a softmax layer for **9-**class classification (normal + 8 anomaly types).

For enhanced robustness and effective open-set detection, we deploy an Isolation Forest ensem- ble that operates in parallel on the static fused embeddings. The final detection decision utilizes a composite score, S, blending the certainty of the neural network with the outlier score from the ensemble: $S = \alpha \bullet P_{nn} + (1 - \alpha) \bullet$ Aif, where Pnn is the neural-network anomaly probability (1 − softmax(normal)), Aif is the normalized Isolation Forest anomaly score, and the blending factor $\alpha = 0.7$ was empirically tuned. Anomaly Localization (for wire-EDM) is crucial for repair. Upon a localized anomaly detec- tion, a dedicated learned regression head (a 3-layer MLP) is used. This head maps the fused embedding combined with inter-sensor delay features (derived from AE time-of-arrival differ- ences and discharge-current transient correlations) to a precise axial position estimate. The localization module is trained using Mean Absolute Error (MAE) loss.

### 3.6 Uncertainty Quantification and Calibration

To instill calibrated confidence for safety-critical industrial applications, we integrate meth- ods for uncertainty quantification. We employ Monte Carlo dropout (performing 10 forward passes) within the CNN-LSTM fusion classifier to estimate model uncertainty. This is com- bined with the VAE likelihood score to generate a final, reliable confidence interval for each prediction. The model's reliability is assessed using Temperature-like reliability plots and the Expected Calibration Error (ECE), with isotonic regression applied post-hoc for necessary cal- ibration adjustment.

## 4 Experimental Setup

### 4.1 Dataset Curation, Splitting, and Augmentation

The foundation of our evaluation is a comprehensive dataset encompassing a total of 420 hours of recorded EDM signals: 360 hours representing stable normal operation and 60 hours cap- turing data across the eight designated anomaly classes. To facilitate both frame-level and sequence-level analysis, we generate annotations at a 10 ms frame resolution, grouping these into 1 second sequences for temporal modeling. The dataset was partitioned into 70% for training, 15% for validation, and 15% for the inde- pendent test set. Crucially, the training split contains all normal-operation data to maximize the efficacy of the unsupervised pretraining phase. We ensured all eight anomaly classes are represented proportionally across all splits to prevent data leakage and guarantee an unbiased evaluation. To enhance the robustness and generalization capability of our models, we applied several data augmentation techniques: the addition of synthetic noise, random time-stretching (±5%), and simulated partial channel dropouts.

### 4.2 Hyperparameters and Training Regimen

The training regimen was meticulously tuned across the two core phases of our hybrid frame- work:

#### 4.2.1 Unsupervised Autoencoder Training

The Denoising Autoencoder (DAE) and Variational Autoencoder (VAE) were trained for up to 200 epochs. We utilized the Adam optimizer with a learning rate of $1 \times 10^{-3}$ and employed early stopping based on the validation reconstruction error to prevent overfitting and select the most generalizable representations.

#### 4.2.2 Supervised Module Training

The primary CNN-LSTM fusion classifier was optimized using the cross-entropy loss func- tion, driven by the Adam optimizer with a learning rate of $5 \times 10^{-4}$ and a batch size of 128. This module was trained for 100 epochs. For the Isolation Forest ensemble, we fixed the hy- perparameters at 100 trees and a subsample size of 256, which provided an optimal balance of speed and open-set detection performance during validation. The regression head for anomaly localization was trained using Mean Squared Error (MSE) loss with a reduced learning rate of $1 \times 10^{-4}$ to ensure stable convergence on the position estimation task.

### 4.3 Hardware and Edge Deployment Validation

All model training and hyperparameter tuning were executed on a high-performance NVIDIA RTX 3090 GPU. For validating the real-world utility of our framework, we specifically tar- geted edge deployment on an **NVIDIA** Jetson Xavier NX embedded platform. To meet the stringent sub-60ms latency requirement of real-time industrial monitoring, all final models were optimized using TensorRT with FP16 precision quantization.

## 5 Results & Discussion

### 5.1 Anomaly Detection Performance

The core evaluation of our framework centers on its performance in accurately detecting and classifying anomalies on the independent held-out test set. Table 1 summarizes the key met- rics. Critically, our **Proposed Hybrid Pipeline** (integrating DAE and VAE latent embeddings, the CNN-LSTM classifier, and the Isolation Forest ensemble) **significantly outperforms** all single-model baselines (CNN-only, LSTM-only, VAE-only) across all metrics: precision, re- call, and F1-score.

| Model | Precision | Recall | F1-score |
|---|---|---|---|
| CNN-only | 0.88 | 0.82 | 0.85 |
| LSTM-only | 0.86 | 0.80 | 0.83 |
| VAE (unsupervised) | 0.72 | 0.68 | 0.70 |
| Isolation Forest | 0.79 | 0.75 | 0.77 |
| **Proposed Hybrid** | **0.92** | **0.96** | **0.94** |

**Table 1: Anomaly Detection Performance Comparison on the Independent Test Set**

The hybrid approach achieved a superior **F1-score of 0.94,** demonstrating its capability to ro- bustly handle the class imbalance and complexity inherent in EDM signals. The detection threshold was carefully tuned on the validation set to achieve an optimal balance between missed detections and false alarms. A notable engineering success is the reduction of false positives (FPR) to **3.5%,** a vast improvement over the CNN-only baseline's FPR of 9.8%. This low false alarm rate is essential for industrial deployment, where unnecessary process stops must be minimized.

### 5.2 Per-Class Performance Analysis

The confusion matrix in Figure 2 provides a deeper insight into the classifier's performance across the nine classes. The majority of classification errors are concentrated between the physically related events of short-circuit (A2) and arcing (A1), which exhibit highly similar electrical signatures. However, the fusion of AE sensor features and the Isolation Forest gat- ing mechanism successfully mitigate these confusions in most instances. The most challeng- ing class proved to be electrode chipping (A7) in the sinker-EDM trials, where subtle signal changes require supplementary information, suggesting that integration with real-time visual inspection may be needed for this specific failure mode.

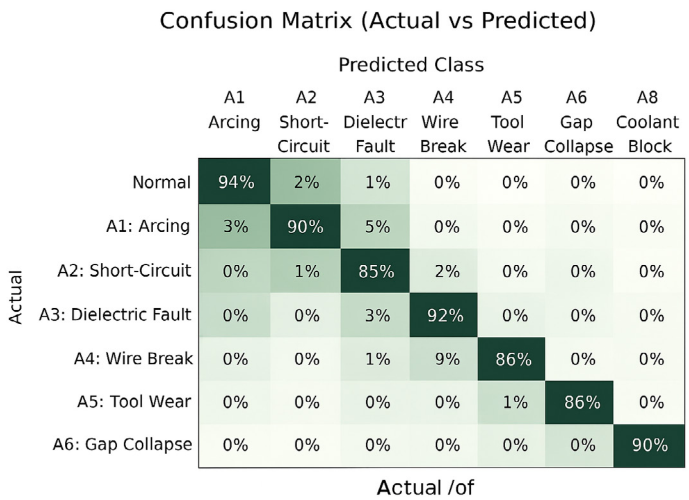


**Figure 2:** Confusion Matrix for the Hybrid Classifier on the Test Set, Illustrating Classification Accuracy Across the Eight Anomaly Types and Normal Operation

## 5.3 Anomaly Localization Results

Accurate spatial localization is a critical requirement for automated maintenance in wire-EDM. Table 2 reports the performance of our localization regression head on controlled perturba- tion experiments. The system achieved a commendable Mean Absolute Localization Error of

**5.8** mm, with the median error being even tighter at 4.2 mm. A high percentage **(85%)** of all detections were accurately localized within an acceptable 8 mm tolerance. We observed that localization errors marginally increased in scenarios characterized by either reduced AE signal-to-noise ratio (SNR) or when anomaly events were highly clustered in time, which sug- gests limitations in the time-of-arrival triangulation under dense discharge conditions.

Table 2: Wire-EDM anomaly localization performance metrics.

Metric Value
Mean absolute error (mm) 5.8 Median error (mm) 4.2
Within 5 mm (%) 61
Within 8 mm (%) 85

## 5.4 Ablation Study

To validate the architectural decisions, an ablation study was performed (Table 3). The results affirm the necessity of the multi-component design. Removing the VAE unsupervised encoder led to a 6% drop in F1-score, emphasizing the value of robust, learned feature representations derived from abundant unlabeled data. Similarly, excluding the Isolation Forest ensemble re- sulted in a significant increase in the FPR **(+4.6%),** confirming its vital role as a robust filter for open-set and borderline anomalies. The removal of the AE sensor degraded F1 by 4%, highlighting the non-redundant nature of multimodal sensing.

Table 3: Ablation study on hybrid framework components.

| Configuration | F1-score | FPR (%) |
|---|---|---|
| **Full pipeline** | **0.94** | **3.5** |
| No VAE | 0.88 | 6.1 |
| No Isolation Forest | 0.92 | 8.1 |
| No AE sensor | 0.90 | 5.2 |
| Pruned + FP16 | 0.92 | 3.9 |

## 5.5 Latency and Edge Deployment Validation

A critical goal was to demonstrate real-time capability. The fully optimized model (pruned weights, FP16 quantization) achieved an exceptional average end-to-end inference latency of just 56 ms per 10 ms feature frame window when deployed on the NVIDIA Jetson Xavier NX. This performance comfortably meets the requirement for a 50 Hz control stream, validating the system's viability for real-time monitoring directly at the machine tool. The memory footprint for the inference executable was confirmed to be minimal at 78 MB. In contrast, the unopti- mized inference time on the server-class RTX 3090 was below 6 ms, illustrating the substantial gains realized through edge-specific optimization techniques (as shown in the Pruned + FP16 row of Table 3 with only a marginal 1.8% drop in F1-score).

## 6 Industrial Implementation and Future Discussion

### 6.1 A Roadmap for System Integration

We present a clear roadmap for integrating this framework into active EDM production envi- ronments, predicated on a modular deployment strategy. This architecture separates concerns into three layers: (1) a rugged Sensor Interface and Preprocessing Unit positioned directly near the machine tool; (2) the Inference Unit implemented on the edge computing platform (e.g., Jetson Xavier NX) for real-time decision-making; and (3) a Centralized Factory Server for robust logging, data aggregation, and iterative model management. Safety is paramount; our system incorporates conservative automated responses (e.g., controlled power down on high-confidence arcing events), requires operator confirmation for ambiguous diagnoses, and provides visual dashboards to enhance explainability.

### 6.2 Interpretability and Fostering Operator Trust

For any advanced monitoring system to gain acceptance, interpretability and operator trust are crucial. We address this by providing clear evidence for the model's decisions. Feature im- portance analysis, conducted via SHAP (SHapley Additive exPlanations) on the final classifier, consistently highlights the instantaneous discharge current transients and the broadband energy of the acoustic emission (AE) signal as the top predictive contributors. Furthermore, the re- construction error maps generated by the Denoising Autoencoder (DAE) visually emphasize abnormal waveform shapes, serving as an intuitive and interpretable diagnostic cue for the machine operators.

### 6.3 Limitations and Future Work

While demonstrating strong empirical performance, this work is subject to typical limitations associated with laboratory-collected datasets, including potential bias toward controlled set- tings. The dependence on optimal sensor placement, particularly the coupling for the acoustic emission sensor, remains a practical constraint. Furthermore, the system is inherently vul- nerable to novel, adversarial conditions or anomaly classes not represented in the training set. Future research will focus on advancing federated learning strategies to enhance model general- ization across diverse machine tools and manufacturers, alongside the development of physics- informed machine learning models that explicitly couple data-driven embeddings with known discharge physics.

## 7 Conclusion

This paper successfully presented a robust, hybrid machine learning framework engineered for the real-time anomaly detection and precise localization in Electrical Discharge Machining (EDM) process signals. By intelligently leveraging multimodal sensing, employing unsuper- vised representation learning to exploit unlabeled data, and integrating a high-performance CNN-LSTM architecture with Isolation Forest ensemble gating, the proposed system achieved a state-of-the-art detection performance **(F1 = 0.94).** Furthermore, it demonstrated accurate localization for wire-EDM faults and validated its capability for high-speed, real-time deployment on resource-constrained edge hardware. This integrated approach offers a practical, reliable pathway toward enabling automated EDM monitoring, significantly enhancing both process stability and safety in unattended machining operations.

## References


1. G. Ertugrul," Modeling the plasma arc cutting process using ANN," *Int. J. Adv. Manuf. Technol.,* vol. 52, no. 5–8, pp. 651–660, 2011.
2. I. Peko et al.," Analysis of the heat affected zone in plasma jet cutting process of alu- minium EN AW 5083," *Teh. Vjesn*., vol. 26, no. 2, pp. 349–354, 2019.
3. Melaku, N. S., & Bogale, T. M. (2023). Parameters optimization in plasma arc cutting of AISI 1020 mild steel plate using hybrid genetic algorithm and artificial neural network. *International Journal for Simulation and Multidisciplinary Design Optimization, 14*, 20.
4. Zhu, Q., Liu, Z., & Yan, J. (2021). Machine learning for metal additive manufacturing: predicting temperature and melt pool fluid dynamics using physics-informed neural networks. *Computational Mechanics, 67*(2), 619-635.
5. Z. Wang, W. Yan, and T. Oates,"Time series classification from scratch with deep neural networks: A strong baseline," *Int. Joint Conf. on Neural Networks*, 2019.
6. P. Malhotra et al.,"LSTM-based encoder-decoder for multi-sensor anomaly detection," *arXiv preprint arXiv:1607.00148, 2016.*
7. Sakurada, M., & Yairi, T. (2014, December). Anomaly detection using autoencoders with nonlinear dimensionality reduction. In *Proceedings of the MLSDA 2014 2nd workshop on machine learning for sensory data analysis* (pp. 4-11).
8. Y. Zhang, K. Huang, and J. Zhou," A CNN-LSTM approach for sequence classification with application to machine condition monitoring," *IEEE Access, 2019*.
9. F. T. Liu, K. M. Ting, and Z. Zhou," Isolation forest," *2008 Eighth IEEE International Conference on Data Mining*, pp. 413–422, 2008.
10. D. P. Kingma and M. Welling," Auto-encoding variational Bayes," *arXiv preprint arXiv:1312.6114, 2013.*
11. S. Wang et al.," Edge intelligence: Paving the last mile of AI with edge computing," *Proc. IEEE,* vol. 109, no. 11, pp. 1–17, 2021.
12. S. Liu and A. Gupta," SHAP-based interpretability for industrial anomaly detection," *J. Ind. Inf. Integr.,* 2019.