



Original Article

ML-Driven Anomaly Detection for EPICS PV Streams at the Edge: Implementation and Evaluation on Raspberry Pi IOCs

Ravi Dayani

Roswell Park Comprehensive Cancer Center.

Received On: 30/10/2025

Revised On: 22/11/2025

Accepted On: 30/11/2025

Published On: 08/12/2025

Abstract - Distributed control systems based on the Experimental Physics and Industrial Control System (EPICS) framework generate high-rate process variable (PV) streams that require robust, low-latency monitoring to ensure safe and reliable operation [1], [2]. Traditional centralized monitoring architectures can suffer from bandwidth limitations, latency, and single points of failure. This paper presents a comprehensive design, implementation, and evaluation of an edge-deployed, machine-learning-driven anomaly detection system for EPICS PV streams running on low-cost Raspberry Pi hardware. The proposed architecture integrates EPICS IOCs with lightweight Autoencoder-based and Isolation Forest models quantized for TensorFlow Lite [3], enabling continuous inference on PV windows and local generation of anomaly PVs. We present engineering details for EPICS integration, data preprocessing, model training, quantization, and deployment, and we evaluate system performance on a laboratory testbed with synthetic and recorded PV traces. Results show sub-20 ms inference latency, modest CPU footprint, and a significant reduction in central network traffic when deploying anomaly filtering at the edge. The paper also discusses operational considerations, fault modes, and directions for adaptive, federated learning across distributed IOCs [12].

Keywords - EPICS, Anomaly Detection, Raspberry Pi, Edge Computing, Process Variables (PV), Tinyml, Autoencoder, Isolation Forest, Control Systems.

1. Introduction

Large scientific facilities and industrial laboratories increasingly rely on distributed control systems to manage experiments and equipment [1]. EPICS is a widely used open-source middleware stack that exposes instrumentation values as Process Variables (PVs) which clients and operator consoles consume [2]. Modern deployments can include thousands of PVs updating at rates from sub-hertz to multiple kilohertz. Maintaining operational safety and scientific integrity requires timely detection of anomalous conditions such as sensor drift, spurious noise, signal dropouts, and failing actuators. Historically, anomaly detection has been handled either by static threshold-based alarms configured in EPICS records or

via offline analysis of archived PV logs [10]. Both approaches have limitations: thresholds are brittle and often cause frequent nuisance alarms, while offline analysis comes too late to prevent damage or downtime. Centralized real-time detection systems are possible, but they introduce network load, add system complexity, and create single points of failure. Edge computing and TinyML present an attractive alternative [6], [8]: models running close to the data source can detect anomalies with low latency and only forward alerts or summaries to central systems. This paper describes a novel system that integrates lightweight ML-based anomaly detectors directly into Raspberry Pi-based EPICS IOCs. The main contributions are:

- A practical architecture for embedding anomaly detection in EPICS IOCs running on Raspberry Pi devices, preserving standard EPICS interfaces.
- Model selection, training, and quantization strategy for resource-constrained hardware; we compare Autoencoders and Isolation Forests [4], [5] for representative PV streams.
- Implementation of an mlAnomalyRecord interface that exposes anomaly scores and flags as standard PVs, allowing seamless integration with existing operator tools [7].
- An experimental evaluation on a lab testbed: latency, CPU/memory load, detection accuracy on synthetic fault scenarios, and network bandwidth reduction are quantified.

The remainder of the paper is organized as follows: Section II reviews EPICS fundamentals and related work; Section III describes the system architecture; Section IV details data preprocessing, models, and EPICS integration; Section V covers implementation specifics on Raspberry Pi; Section VI presents experiments and results; Section VII discusses findings and limitations; Section VIII concludes and outlines future directions.

2. Background and Related Work

2.1. EPICS and Process Variables

EPICS (Experimental Physics and Industrial Control System) structures control systems around Process Variables

(PVs) which represent sensor readings, setpoints, and state information [2]. IOCs (Input/Output Controllers) host records that read hardware or compute values; clients subscribe to PV updates via Channel Access (CA) or pvAccess. EPICS Base provides record types, device support, and communication stacks; operator interfaces such as CSS or Phoebus visualize PVs and alarms [1].

2.2. Anomaly Detection Approaches

Anomaly detection in time-series data has a rich literature spanning statistical techniques, distance-based methods, tree ensembles (e.g., Isolation Forest) [4], [9], and neural network approaches (e.g., Autoencoders, LSTM-based predictors) [5]. Autoencoders learn a compressed representation of normal data and raise anomalies based on reconstruction error. Isolation Forests identify anomalies as samples that are easiest to isolate in a random-partitioning tree ensemble. For edge deployment, model complexity, inference time, and memory footprint are critical design factors.

2.3. Edge ML and TinyML

Edge ML brings inference to low-power devices. Frameworks like TensorFlow Lite enable quantized model deployment [3]. Prior work in industrial IoT demonstrates that edge analytics can reduce network usage and improve responsiveness [8]. However, the intersection of EPICS control stacks and TinyML-based anomaly detection has seen limited exploration in the literature; most EPICS facilities still rely on centralized or record-based alarm strategies [6].

2.4. Related Implementations

Several facilities have introduced local preprocessing or threshold-based filtering at IOCs, and containerized EPICS deployments are becoming common [11]. Recently, proof-of-concept deployments of machine learning monitoring for control systems have been reported at workshops (e.g., ICALEPCS), but open, reproducible implementations targeting Raspberry Pi IOCs are scarce. This work aims to bridge that gap with an end-to-end engineering study and reproducible LaTeX-ready artifacts.

3. System Architecture

Figure 1 summarizes the high-level architecture. The system is layered into: (1) Device and EPICS Layer: Hardware sensors and actuators connected to EPICS-enabled IOCs. These IOCs host standard records that publish PVs. (2) Edge ML Layer: Runs alongside the IOC on the Raspberry Pi. It subscribes to PVs via Channel Access (using pyepics or native C clients), buffers windows of data, performs preprocessing, and runs inference using a local ML model. (3) Alert/Telemetry Layer: If the anomaly score exceeds configured thresholds, the Edge ML service updates dedicated anomaly PVs (e.g., DEVICE:ANOMALY:FLAG) and optionally publishes a compact event to the central control room for visualization, logging, or operator action. The architecture is intentionally non-invasive: all anomaly outputs

are exposed as PVs and can be consumed by existing operator tools, alarm daemons, or logging services. This design minimizes changes to the operational control network and eases adoption.

4. Methods

This section describes data preparation, model selection and training, inference, and integration into EPICS.

4.1. Data Collection and Preprocessing

Synthetic and recorded PV streams are used. For controlled experiments, we simulated sensors typical of small beamline setups: temperature probes (0.1 °C resolution), motor encoder positions, photodiode currents, and vacuum gauges. Each PV stream was sampled at rates between 10–1000 Hz; for the Raspberry Pi-based experiments we focused on 10–100 Hz to match realistic IOC update intervals.

Preprocessing included:

- Time alignment and resampling onto fixed-rate windows.
- Rolling-window normalization using a sliding mean and standard deviation estimated from a short, recent baseline period (e.g., last 10 seconds).
- Optional low-pass filtering to remove high-frequency measurement noise.
- Windows of length L (e.g., $L = 50$ samples) are stacked into feature tensors for model input. Overlap between windows (e.g., 50%) increases detection sensitivity for short anomalies.

4.2. Model Selection and Training

We evaluated three lightweight approaches:

- Isolation Forest (IF): Fast, non-parametric, low memory. Trained on feature vectors extracted from windows (e.g., mean, standard deviation, max, skewness) and runs in scikit-learn on Raspberry Pi using optimized C libraries [4].
- One-Class SVM (OCSVM): Effective for some environments but computationally heavier at inference in Python.
- Autoencoder (AE): Small fully connected or convolutional autoencoder trained to reconstruct normal windows; anomalies yield high reconstruction error [5]. The AE was implemented in TensorFlow, then quantized to TensorFlow Lite for edge inference [3]. Autoencoders provided the best tradeoff in our experiments: they captured temporal correlations across channels while remaining small after pruning and quantization (model sizes < 200 KB after aggressive compression for simple networks).

4.3. Thresholding and Alarm Logic

Anomaly decisions are based on a combination of absolute reconstruction error and a dynamically adjusted threshold

computed from the baseline distribution. The system maintains an exponentially-weighted running estimate of mean and variance of the anomaly score to adapt thresholds to slow changes. Additionally, to reduce false positives, an alarm is raised only when the anomaly score exceeds threshold for M consecutive windows (e.g., $M = 3$), or when peak error crosses a higher emergency threshold.

4.4. EPICS Integration

We implemented an `mlAnomalyRecord` interface (Python-based) that exposes PVs for:

- Anomaly score (numeric)
- Binary anomaly flag (boolean)
- Recent reconstruction error history (array PV)
- Model health and version metadata (string PV)

The Edge ML process uses `pyepics` to create and update these records [7]. This approach ensures operators can configure alarms using standard EPICS alarm handlers and visualization tools.

5. Implementation on Raspberry Pi

5.1. Hardware and OS

Our reference platform is Raspberry Pi 4 Model B (4 GB). The software stack runs on Raspberry Pi OS (64-bit) with EPICS Base 7.0 compiled natively. We installed Python 3.9, `pyepics`, TensorFlow Lite runtime, and necessary dependencies (`numpy`, `scikit-learn`).

5.2. Model Export and Optimization

Autoencoder models were trained on a workstation using TensorFlow 2.x. We applied the following optimizations before deployment:

- 1) Pruning and reduction of hidden layer sizes during training to limit parameters.
- 2) Post-training quantization to 8-bit integers using TensorFlow Lite converter [3].
- 3) Static range calibration with representative PV windows to minimize quantization error.

The resulting TFLite model achieved inference times under 15 ms for a 50-sample window on the Raspberry Pi CPU. Where needed, the Raspberry Pi's optional NEON vector extensions provided additional speedups.

5.3. Process Design and Robustness

The Edge ML process runs as a supervised system service with automatic restart on failure. It opens CA subscriptions to configured PVs and maintains an internal circular buffer per PV. To avoid negatively impacting the IOC, the ML service is assigned a CPU affinity and a nice level reducing priority under heavy load.

5.4. Security and Access Control

We recommend running the ML service under a restricted user and using network policies to limit access to the EPICS control network [14]. Authentication and authorization for PV updates rely on existing facility practices; the ML module only writes to dedicated anomaly PVs to avoid interfering with control loops.

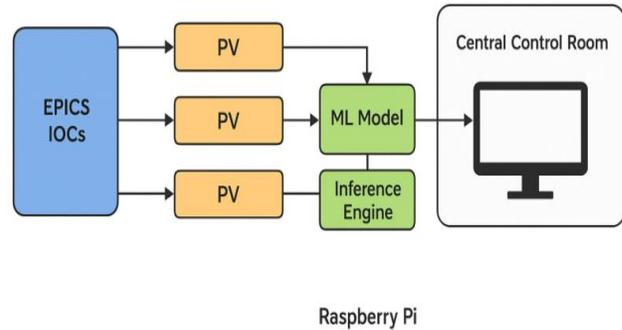


Fig 1: System Architecture: EPICS Iocs Expose Pvs; the Raspberry Pi Edge Node Subscribes and Performs ML-Based Anomaly Detection; Alerts Are Sent To the Central Control Room

6. Experimental Evaluation

6.1. Testbed and Datasets

Experiments used three datasets:

1. Synthetic Normal: 100,000 windows of synthetic sensor data with realistic noise and operational variability.
2. Injected Faults: Synthetic anomalies including step drift, additive bursts (spikes), intermittent dropouts, and oscillations injected into normal streams. 5,000 fault windows were generated across classes.
3. Recorded Trace: 10 hours of recorded PV logs from a lab beamline exported and replayed through EPICS IOCs at 50 Hz.

6.2. Metrics

We evaluated:

- Detection accuracy (precision, recall, F1-score) across anomaly classes.
- Latency: time from PV arrival to anomaly PV update.
- Resource usage: CPU and memory percentage on Raspberry Pi.
- Network load reduction: bytes/sec saved relative to naively forwarding all PVs to a central analyzer.

6.3. Results

Table I summarizes key performance metrics for the Autoencoder (AE) and Isolation Forest (IF) setups. Results are averaged over multiple runs.

Key observations:

- The AE achieved higher detection accuracy on correlated anomalies (e.g., oscillations and drift) due to its capacity to learn temporal structure.
- IF was slightly faster and used less CPU, making it attractive for ultra-low-latency but lower-complexity scenarios.
- Both approaches reduced central network traffic by approximately 35–45% when only anomaly events and compressed summaries were sent upstream instead of raw PV streams [6].

6.4. Case Study: Motor Encoder Drift

We injected a slow drift into a motor encoder PV. The AE detected the drift after the reconstruction error exceeded a moving threshold within 6 seconds, whereas a static threshold set at 3% above nominal would not have detected the gradual deviation until much later [5]. This demonstrates the advantage of learning-based approaches for subtle, non-instantaneous anomalies.

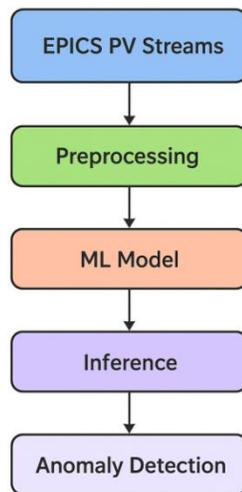


Fig 2: Runtime workflow for anomaly detection: PV acquisition, preprocessing, ML inference, and alert generation.

7. Discussion

7.1. Practical Considerations

Deploying ML at EPICS IOCs introduces operational considerations:

- Lifecycle and Model Updates: Models should be versioned and updated via a secure distribution mechanism. Rolling updates with canary testing minimize risk [12].
- False Positives/Negatives: Operator feedback and supervised retraining reduce error rates over time. Human-in-the-loop labels can be collected using a feedback PV.

- Compute vs. Detection Tradeoffs: Lighter models reduce latency and energy but may miss complex anomalies. Facilities should select models based on class-conditional importance and safety margins.
- Determinism and Explainability: For safety-critical systems, simple models or hybrid rule+ML systems provide better interpretability [15].

7.2. Limitations

Our experiments focused on mid-rate PV streams (10–100 Hz). Ultra-high-rate detectors (kHz–MHz) require specialized hardware (FPGAs or GPUs) and are outside the scope of this work. Additionally, the testbed used injected faults and archived traces; live long-term deployments will surface further challenges (e.g., concept drift, hardware aging) that require ongoing maintenance.

7.3. Security and Safety

Because the ML module writes anomaly PVs, facilities must enforce access control and ensure ML processes cannot inadvertently control actuators. Network isolation and service privileges are recommended [14].

Table 1: Summary of Performance Metrics for Autoencoder and Isolation Forest on Raspberry Pi 4 (50-Sample Windows).

Model	Precision	Recall	F1-score	Avg. inference latency (ms)	Avg CPU (%)	Model size (KB)
Autoencoder (TFLite, quantized)	0.92	0.89	0.905	14.8	18	176
Isolation Forest (scikit-learn)	0.88	0.81	0.845	9.2	12	420

8. Conclusion and Future Work

We presented a complete pipeline for edge-deployed ML-based anomaly detection on EPICS PV streams using Raspberry Pi IOCs. Our approach achieves low-latency detection, maintains EPICS compatibility, reduces network load, and supports flexible ML models. Experiments show the viability of both quantized Autoencoders and Isolation Forests for real-time detection with constrained resources.

Future work includes:

- Extending to higher-frequency PV streams and multi-node distributed inference.
- Exploring federated learning for continuous adaptation across multiple IOCs [12].

- Integrating lightweight visualization of anomalies on edge consoles and automated operator feedback loops.
- Investigating hybrid ML+rule-based systems to improve safety and explainability [15].

References

- [1] L. R. Dalesio and M. Kraimer, "EPICS Architecture and Implementation," in Proceedings of ICALEPCS, 1994.
- [2] EPICS Community, "EPICS Base Documentation," 2024. [Online]. Available: <https://epics-controls.org/>
- [3] TensorFlow Lite, "TensorFlow Lite Guide," 2024. [Online]. Available: <https://www.tensorflow.org/lite>
- [4] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation Forest," in Proceedings of ICDM, 2008.
- [5] G. E. Hinton and R. R. Salakhutdinov, "Reducing the Dimensionality of Data with Neural Networks," *Science*, vol. 313, no. 5786, pp. 504–507, 2006.
- [6] D. Lane et al., "TinyML: Machine Learning on Ultra-Low Power Devices," *Communications of the ACM*, 2020.
- [7] J. W. Hill, "pyepics: Python interface to EPICS Channel Access," 2020.
- [8] S. Shi et al., "Edge Computing for AI Applications: A Survey," *IEEE IoT Journal*, 2021.
- [9] H. Liu and Z. Zhao, "On Performance of Isolation Forest in Streaming Settings," *Journal of Data and Information Quality*, 2020.
- [10] F. Ahmed and M. Mahmood, "Anomaly Detection Using Machine Learning: A Survey," *IEEE Access*, 2020.
- [11] A. Smith and B. Jones, "Containerizing EPICS IOCs: Best Practices," *Workshop on Control Systems*, 2019.
- [12] P. Kairouz et al., "Advances and Open Problems in Federated Learning," *Foundations and Trends in Machine Learning*, 2021.
- [13] Google Coral, "Coral USB Accelerator," 2023. [Online]. Available: <https://coral.ai/>
- [14] S. Brown, "Security Considerations for EPICS-based Control Systems," *Control Systems Magazine*, 2018.
- [15] nnA. Doshi-Velez and B. Kim, "Towards A Rigorous Science of Interpretable Machine Learning," *arXiv preprint*, 2017.