

Explainable and Trustworthy AI for Fault Classification in the Tennessee Eastman Process: A Step Toward Industrial Autonomy

Jayanth Balaji Avanashilingam, and Bijuraj Pandiyath Velayudhan

AI Lab, Co-innovation Centre, Singapore Development Centre, Yokogawa Engineering Asia Pte. Ltd, Singapore

jayanth.avanashilingam@yokogawa.com

bijuraj.pandiyath@yokogawa.com

ABSTRACT

Achieving higher levels of Industrial Autonomy (IA) requires fault diagnostic systems that combine predictive accuracy with transparent decision-making. In safety-critical process industries like petroleum refinery, black-box AI models often face adoption barriers due to limited interpretability. The work introduces a glass-box fault classification framework for the Tennessee Eastman Process, comparing a baseline direct-modeling approach with a novel dual-branch architecture. The proposed method decomposes process parameters into trend and cyclic components, trains dedicated classifiers on each and fuses their probabilistic outputs. The proposed design improves sensitivity to both gradual drifts and oscillatory anomalies. In the present work SHAP explainability is incorporated to provide global, local, and class-wise feature attribution, enabling operators to trace model reasoning and align diagnostics with process knowledge. Building on this, a strong industrial AI platform, purpose-built for domain engineers, emerges as essential for operationalizing such capabilities, empowering process experts to directly harness AI for decision-making. The present work serves as a steppingstone toward realizing such an Industrial AI platform, demonstrating how interpretable AI can bridge the gap between advanced analytics and domain expertise. The experimental evaluation of the proposed technique demonstrates that 35% of the fault classes achieved improved accuracy, with an average accuracy gain of 4.34% over the baseline, with pronounced gains in cyclic-dominated faults. The approach demonstrates a pathway toward Level 5 IA by delivering interpretable, high-performance fault diagnostics ready for real-time deployment.

1. INTRODUCTION AND MOTIVATION

Industrial process industries are undergoing a gradual but fundamental shift toward IA, where plants operate with minimal human intervention while ensuring safety, reliability, and efficiency. The IA transformation is fueled by competitive market pressures, an evolving workforce landscape, and the escalating complexity of interconnected process systems. In traditional IA Levels 1-3 as shown in Figure 1, operations rely on deterministic control loops and predefined logic, typically implemented in Distributed Control Systems (DCS) or Programmable Logic Controllers (PLC). In IA Levels 4-5, the systems must be adaptive and context-aware, capable of perceiving process states, reasoning about alternative actions, and executing control decisions autonomously under uncertainty (*Autonomous Operations for Process Industries?* | ARC Advisory, n.d.-a)

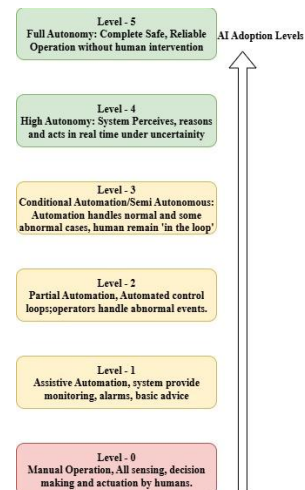


Figure 1. Levels of Industrial Autonomy.

However, reaching higher levels of autonomy in industrial environments poses significant challenges, as also

emphasized in (*Industrial-Grade Data Fabrics* | ARC Advisory Group, n.d.):

- Heterogeneous, high-dimensional sensor data from thousands of measurement points.
- Non-linear, strongly coupled process dynamics that defy simple models.
- Rare but high-impact fault events with limited historical data.
- Regulatory and operator constraints for autonomy.

Bridging these gaps requires AI systems that are both high-performing and explainable, enabling a transition from opaque black-box models to glass-box autonomy that integrates seamlessly into operator workflows.

In the present study an attempt is made to build an explainable and trustworthy fault classification model to address the gaps and to take a step towards IA journey. The paper is organized as follows: section 2 discusses the Challenge in the IA journey, followed by section 3 details out the need of AI for IA journey with fault classification as a cornerstone discussed in Section 4. Section 5 brings out the motivation for explainable AI and how the proposed solution fits in for IA journey. Section 6 discusses the proposed methodology and section 7 analyzes the results and section 8 concludes the present work with future directions of AI in IA journey.

2. CHALLENGES IN THE IA JOURNEY

The most common challenges of IA journey are based on data availability, latency requirement, transparency in decision and control integration. IA, particularly at Levels 4 and 5, is not just a technological leap from existing automation; it is an operational transformation that must overcome several deep-rooted challenges in process industries. The challenges span across data integrity, real-time decision-making, explainability, and system integration, which directly impacts the pace and safety of the autonomy transition. The following sections details out these challenges.

2.1. Data Quality and Variability

Accurate decision-making in autonomous systems depends on high-quality process data, yet real-world industrial data is often affected by:

- Sensor noise and bias, caused by instrument degradation or calibration drift.
- Unavailability of measurements.
- Latency in measurements, as with many quality parameters, is based on lab results.
- Missing data from network disruptions or sensor failures.

- Process variability from changes in raw materials, environment, or operating mode.

Poor data directly compromises model generalization and robustness. Data preprocessing, sensor validation, and adaptive model updating have become essential for sustained autonomous operation.

2.2. Fault Detection Latency

In process industries, time-to-detection is critical, slow detection of anomalies can escalate into:

- Safety incidents such as equipment failure, hazardous releases.
- Quality deviations like off-spec products, waste generation
- Production downtime like unplanned shutdowns, restart costs.

Traditional statistical process control (SPC) or model-based methods may detect deviations only after significant process drift, especially in non-stationary environments. In order to achieve Level 5 IA, the systems must integrate real-time analytics to detect and diagnose faults before they propagate.

2.3. Transparency and Trust

A core requirement for IA adoption is operator trust. Many operators, accustomed to deterministic control logic, view black-box AI models with doubt, particularly in high-stakes environments like petroleum refineries where accountability is paramount.

Glass-box AI approaches, which provide interpretable and justifiable outputs, are critical for bridging this gap. Explainability not only increases adoption but also assists in post-event root cause analysis and regulatory compliance (*IEC TR 63069:2019* / IEC, n.d.).

2.4. Control Integration

In order for autonomy to be operationally viable, AI-driven decisions must integrate seamlessly with:

- Existing safety instrumented systems (SIS).
- Plant standard operating procedures (SOPs).
- Real-time control platforms (DCS/PLC).

The integration is non-trivial, any conflict between AI recommendations and safety logic can cause system overrides, nullifying autonomy benefits.

In summary the above four challenge domains such as data quality, fault detection latency, trust and explainability, and control integration represent the gap between today's automation (Levels 1–3) and the envisioned self-optimizing, self-healing industrial plants of the future (Levels 4–5). Closing this gap requires trustworthy, explainable AI that can

operate safely, transparently, and in real time within the constraints of existing industrial ecosystems.

Bridging this gap is not solely a matter of incremental control logic improvement; it demands a paradigm shift in decision-making approaches. AI offers the ability to perceive, learn, and adapt in complex process environments attributes essential for achieving higher autonomy levels. The next section briefs the role of AI in IA.

3. ROLE OF AI AND FEATURE ENGINEERING IN INDUSTRIAL AUTONOMY

AI is pivotal for elevating automation from reactive control to adaptive, autonomous operations (Levels 4–5)(Ohara, 2020; Zhu et al., 2020). Its strengths such as pattern recognition, predictive reasoning, and real-time learning address the key limitations of legacy systems. Unlike fixed logic and deterministic control schemes, AI can learn from data, adapt to new conditions, and make informed decisions under uncertainty. Its contribution to IA can be understood through three core functional roles: perception and diagnosis, signal understanding and knowledge extraction, and transparent decision support.

3.1. Perception and Diagnosis through Supervised Learning

The backbone of autonomy lies the ability to perceive process state and identify deviations before they escalate. Supervised learning models are trained on historical labeled process data which enables systems to detect abnormal conditions and classify underlying fault types. In complex chemical processes such as the petroleum refinery process, these models can capture subtle, multivariate signatures that conventional threshold-based detection would miss.

Accurate classification empowers the system to move beyond reactive alarms toward proactive interventions, which is essential for mitigating safety risks, avoiding quality deviations, and minimizing economic losses. The perception-and-diagnosis capability provides the situational awareness foundation necessary for higher autonomy levels.

However, the value of perception is directly linked to the system's ability to correctly interpret and process sensor signals which brings to the second role of AI in industrial autonomy: understanding and extracting meaningful knowledge from raw process data.

3.2. Signal Understanding and Knowledge Extraction by Feature Engineering

Industrial sensor data typically exhibits long-term drifts, short-term oscillations, and measurements with noise. Without careful treatment, these characteristics can obscure the patterns that AI needs to detect. Signal decomposition

techniques such as empirical mode decomposition, wavelet transforms, or trend–cycle separation allows AI models to treat slow degradation phenomena such as fouling, catalyst degradation separately from oscillatory phenomena like valve stiction, control loop instability.

The signal decomposition enhances model robustness and fault isolation, as certain fault modes manifest predominantly in one component. By extracting knowledge from the right signal layers, AI systems improve both their predictive performance and their diagnostic precision.

In the present work, the signal decomposition is extended by introducing a dual-branch modeling framework that independently learns from the trend and cyclic components of process signals, before fusing their probabilistic outputs for final fault classification. The proposed design exploits the distinct fault signatures present at different temporal scales, enhancing both detection accuracy and interpretability. In the context of IA, accurate predictions alone are insufficient, operators must also understand and trust these predictions which leads to the third and equally critical role of AI in autonomy: transparent decision support.

3.3. Transparent Decision Support via Explainable AI

In industrial autonomy, performance without transparency is insufficient for operational acceptance. Operators must be able to see why a model recommends a certain action, especially in safety-critical contexts. Explainable AI (XAI) methods, such as Shapley Additive Explanations (SHAP), Local Interpretable Model-agnostic Explanations (LIME), and Partial Dependence Plots (PDP), reveal the underlying drivers of AI decisions (Ohtani, 2020).

The representation of feature contributions in human-interpretable form, XAI transforms opaque “black-box” models into glass-box systems that can be audited, validated, and refined collaboratively with domain experts. Transparency is essential for regulatory compliance, fault investigation, and fostering operator confidence in autonomous decision-making and making the vision of IA reachable.

The proposed method further integrates class-wise explainability using SHAP-based visualizations for both the trend and cyclic branches, enabling operators to understand not only which features contributed to a prediction, but also from which temporal component the evidence originated. The combination of predictive accuracy and interpretability is central to closing the autonomy trust gap.

3.4. Integrated Impact on Autonomy

The three functional roles discussed above like perception & diagnosis, signal understanding & knowledge extraction by feature engineering, and transparent decision support together form the operational backbone of IA. When integrated into plant control ecosystems, they enable AI

systems to perceive, interpret, and act in a way that is adaptive, trustworthy, and aligned with safety and operational constraints.

The holistic approach of the integration of three functional roles allows AI to support not just isolated decision-making tasks, but the broader objective of sustained, self-optimizing plant operation and bringing the vision of Levels 4–5 industrial autonomy within operational reach.

Building on these foundational roles, the next section presents our proposed methodology that operationalizes these concepts in a dual-branch trend–cyclic classification framework with integrated class-wise explainability.

4. FAULT CLASSIFICATION AS A CORNERSTONE

Achieving higher levels of IA is impossible without accurate, timely, and reliable fault classification. In autonomous plant operation, classification does more than flag an abnormal condition, it enables the system to determine what kind of abnormality is occurring, so that responses can be targeted, effective, and safe.

Fault classification supports four core operational benefits:

1. Targeted Corrective Actions, by identifying the specific root cause, the system can apply fault-specific recovery strategies for instance, adjusting control loops, initiating selective shutdown sequences.
2. Reduced Downtime and Production Losses, early and correct classification allows intervention before faults propagate to system-wide disruptions.
3. Improved Safety and Compliance, faults in safety-critical units can be addressed in a controlled manner, reducing the likelihood of hazardous events and regulatory violations.
4. Structured Input for Higher-Level Autonomy, at autonomy Levels 4–5, classification outputs serve as structured perception signals for higher-level decision-making agents, optimization algorithms, and scheduling systems.

The above benefits of position classification as a cornerstone capability in the IA stack by bridging perception and intelligent action(Ohtani, 2020). The next section discusses the data sets considered for the proposed work.

4.1. The Tennessee Eastman Process (TEP) as a Benchmark

The Tennessee Eastman Process (TEP)(*Additional Tennessee Eastman Process Simulation Data for Anomaly Detection Evaluation - Harvard Dataverse*, n.d.) is one of the most widely adopted benchmarks in process systems engineering for evaluating fault detection and diagnosis algorithms. TEP simulates a realistic petrochemical process involving multiple reactors, separators, and recycle streams. The TEP

simulates a realistic chemical process with five major unit operations, multiple interacting control loops, 53 variables with 41 measured variables, and 12 manipulated variables, reflecting typical plant instrumentation. The dataset contains two primary subsets: FaultFree (normal operation, fault number = 0) and Faulty (20 distinct fault types, fault numbers 1–20). Each record comprises faultNumber identifying process state, simulationRun indicating the simulation seed (1–500 for training, non-overlapping with testing), sample representing the time index (1–500 for training, 1–960 for testing) sampled every 3 minutes for total durations of 25 hours and 48 hours respectively, and columns xmeas_1 to xmeas_41 (measured variables) plus xmv_1 to xmv_12 (manipulated variables). Faults are introduced one hour into training runs and eight hours into testing runs, allowing models to learn both pre-fault and fault evolution patterns. Because of its realistic process complexity, public availability, and clear definition of multiple fault classes.

5. MOTIVATION FOR EXPLAINABILITY IN IA AND ITS PLACEMENT IN THE AUTONOMY HIERARCHY

High predictive accuracy in fault classification is essential for IA, but alone it is insufficient for operational acceptance in safety-critical environments. In these contexts, plant operators and engineers require more than a categorical fault label. They must be able to interpret the reasoning behind the classification, identify which process variables exerted the greatest influence, and understand how variations in those variables might affect the outcome.

Several operational imperatives drive the demand for explainability:

- Transparency: Model outputs must be accompanied by interpretable justifications, enabling operators to verify decisions against process knowledge (Ohtani, 2020).
- Accountability: Safety and regulatory compliance depend on the ability to audit and trace decision-making processes, particularly in cases of abnormal plant behavior(*AI Policy | Yokogawa Electric Corporation*, n.d.-a)
- Trust: Operators are more likely to adopt AI-driven recommendations when the underlying rationale aligns with plant physics, control logic, and operational experience(Ohtani, 2020).

Glass-Box AI methods address these needs by making the internal reasoning of machine learning models visible. Techniques such as SHAP, LIME, and PDP reveal variable-level contributions to model predictions, turning opaque “black-box” algorithms into auditable and collaborative decision-support tools.

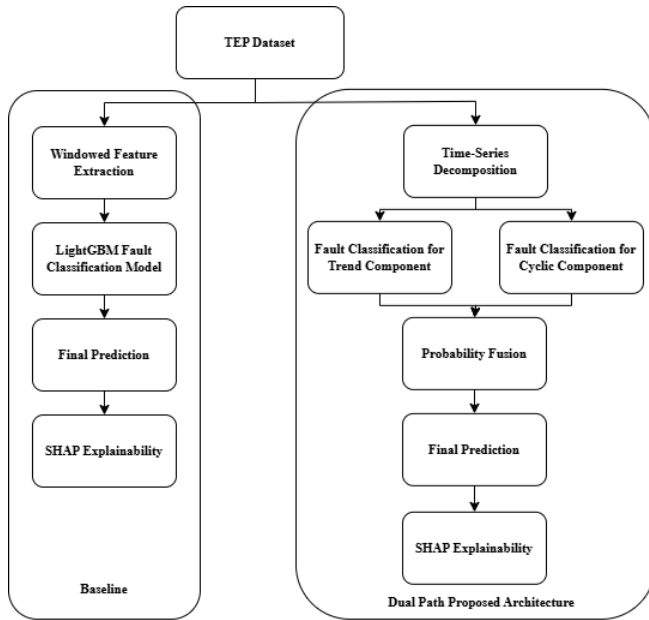


Figure 2. Proposed Methodology.

5.1. Positioning within the ISA/IEC Industrial Autonomy Hierarchy

The ISA/IEC framework describes industrial operational maturity as a progression through five levels (*ISA-95 Standard: Enterprise-Control System Integration*, n.d.) as shown in Figure 1

The fault classification methodology under consideration directly addresses Level 4 and partly in Level 3. It provides adaptive decision-support capabilities that integrate explainable AI into supervisory and optimization layers, enabling operators to make informed interventions with confidence. The decision support capability serves as a transitional step toward Level 5 autonomy, where the AI would execute similar decisions autonomously.

Achieving explainable, high-performance fault classification requires a methodological design that balances predictive accuracy with interpretability. The following section outlines a two-stage approach: beginning with a baseline supervised learning framework applied to the TEP dataset, and extending to a dual-path architecture that decomposes process variables into trend and cyclic components. By training specialized classifiers on each component and fusing their probabilistic outputs, the design improves classification robustness while maintaining transparent, variable-level interpretability, addressing both operational and trust requirements of IA.

6. PROPOSED METHODOLOGY

The current study evaluates two distinct fault classification strategies for advancing IA: a baseline supervised learning pipeline applied directly to raw process data with windowed

feature extraction, and a proposed dual-path architecture that decomposes signals into trend and cyclic components before classification as shown in Figure 2. In the present work Hodrick–Prescott filter is employed to decompose the time-series data into Trend and Cyclic components. Leakage-safe evaluation was ensured by partitioning datasets strictly according to simulationRun identifiers, thereby preventing cross-run contamination between training and testing sets. Furthermore, K-fold cross-validation with $K = 5$ was conducted within the training data to assess model generalization and identify potential overfitting or leakage artifacts.

The baseline pipeline follows a conventional supervised learning workflow comprising data preprocessing, statistical aggregation of windowed time-series segments, model training with gradient boosting algorithms such as LightGBM using stratified cross-validation, and evaluation through macro F1-score, accuracy, and confusion matrices. The windowed feature extraction stage employed a sliding-window decomposition strategy to extract robust statistical representations of the dynamic process behavior, where a fixed window length of 60 samples with a 10-sample overlap was applied to each simulation run, ensuring sufficient temporal context while maintaining high sample diversity. While this approach achieves competitive accuracy, it processes all signal characteristics together, potentially masking fault-specific patterns that manifest differently in slow trends versus rapid oscillations. The LightGBM algorithm employed in this study was identified through an extensive, internally developed AutoML-style workflow applied to the baseline modeling approach. The detailed description of this framework is considered beyond the scope of the present work. The same LightGBM algorithm was then adopted for the proposed trend–cyclic modeling framework to ensure a consistent basis for performance comparison.

The proposed trend–cyclic dual-path architecture addresses this by applying time-series decomposition to separate long-term drift or degradation from short-term oscillatory dynamics. Independent classifiers are trained for each component, one capturing gradual fault evolution, the other targeting oscillatory fault behavior. The two models probabilistic outputs are fused, through averaging, allowing the model to leverage complementary dynamic information for improved fault discrimination and robustness.

In the present study, both approaches incorporate an explainability layer to meet transparency and trust requirements in industrial settings. SHAP explainability is adopted in the present work. SHAP quantifies the contribution of each process variable to classification outcomes, with class-wise analysis identifying the most influential drivers for each fault type. By enhancing both predictive performance and interpretability, this framework moves fault classification from opaque “black-box” models toward glass-box systems, positioning them as enablers for

Level 4 industrial autonomy and a steppingstone toward fully autonomous, self-optimizing operations.

7. RESULT ANALYSIS

The comparative evaluation between the baseline supervised learning pipeline and the proposed trend–cyclic dual-path architecture was conducted on the TEP dataset using the same stratified test split. Figure 3 provides the per-class accuracy comparison of precision, recall and F1 score of baseline model with proposed methodology.

Across all three primary evaluation metrics such as overall accuracy, macro-averaged F1-score, and macro-averaged area under the ROC curve (AUC), the dual-path architecture achieved consistent improvements over the baseline. Accuracy increased from 93.19% to 94.56%, and macro-F1 improved from 0.9283 to 0.9424, indicating enhanced balance between precision and recall across fault classes. Macro-AUC(ROC) and micro-AUC(ROC) both approached unity, with slight gains in the novel approach, confirming that both methods maintained near-perfect separability while the proposed method offered marginal but meaningful enhancements.

A class-level analysis reveals more granular insights. The dual-path approach improved classification accuracy for 7 out of 21 fault categories shown in Table 1.

Fault	% of Improvement in Accuracy in Proposed Method
8	7.81%
20	6.58%
19	5.26%
11	4.17%
16	3.57%
15	1.67%
18	1.32%


Table 1. Fault number and its % of accuracy improvement.

These improvements are particularly noteworthy because they correspond to fault modes with distinct temporal signatures in either trend or cyclic components. For example, Fault 8 exhibited strong oscillatory patterns superimposed on a slow process drift, which the baseline approach processing all features jointly could not isolate effectively. The decomposition in the dual-path method allowed each classifier to specialize, capturing subtle dynamic patterns and improving fault discrimination.


Performance remained unchanged in 12 classes, including several where baseline accuracy was already perfect (Faults

1, 2, 4, 14, 17), indicating that the proposed architecture preserved high performance where the baseline had no shortcomings. Small decreases in accuracy were observed in a few cases (Faults 4, 5, 10), with the largest drop being -1.79% for Fault 10. These declines are within acceptable margins and may reflect over-segmentation in the decomposition process when the underlying signal does not exhibit a strong separation between trend and cyclic components.

In operational perspective, the gains in fault classes with historically lower baseline performance are significant. In safety-critical process environments, even modest improvements in detection accuracy for specific fault types can translate into substantial risk reduction, especially when these fault types are precursors to major incidents or unplanned shutdowns. The per class comparison of precision, recall and F1 of baseline model and proposed model approaches are shown in Figure 3a and 3b respectively.



(a)



(b)

Figure 3. Per class metrics comparison of (a) Baseline with (b) Proposed Methodology.

From Figure 3, the per-class precision, recall, and F1-score comparison shows that while the baseline model delivers strong performance across most fault types, the proposed fused approach consistently improves detection for challenging classes such as Fault 8, 20, 19, 16, and 11, with notable recall gains and in some cases achieving perfect classification (e.g., Fault 8). Well-performing classes in the baseline (e.g., Fault 1, 2, 14) maintain their performance, and only minor drops are observed for a few classes (e.g., Fault 4, 5, 10). These results indicate that the trend–cyclic fusion

enhances robustness and fault discrimination, particularly for fault modes with distinct dynamic signatures, while preserving high accuracy in already well-classified cases.




Figure 4. Mean SHAP values for Baseline model.

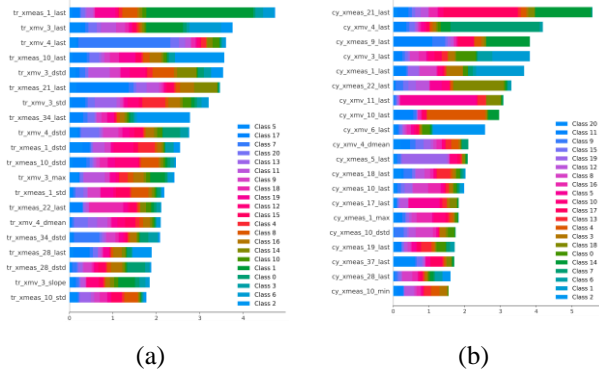


Figure 5. Mean SHAP Values for (a) Trend Component Model, (b) Cyclic Component Model.

The explainability layer, implemented using SHAP, further reinforced these quantitative gains by enabling class-wise identification of the most influential variables. The SHAP values are shown in Figure 4 and 5. While the nuanced implications of these influential variables are best interpreted by process or domain engineers, their identification is a critical enabler for autonomous operations. By translating complex model behavior into process-relevant insights, this step bridges the gap between data-driven decision-making and domain expertise, laying the groundwork for trustworthy, higher-level autonomy. Moreover, incorporating explainability aligns with emerging AI regulatory requirements, such as those outlined in the EU AI Act and industry standards, ensuring that model decisions remain transparent, auditable, and justifiable in safety-critical industrial contexts. The baseline model placed the highest emphasis on variables such as *xmv_10_last*, *xmeas_21_last*, and *xmv_4_last*, reflecting the dominance of certain actuator positions and sensor readings in fault

discrimination. In the cyclic pathway, variables like *cy_xmeas_21_last*, *cy_xmv_4_last*, and *cy_xmeas_9_last* were most influential, capturing fast-changing oscillatory behavior tied to control loop dynamics. The trend pathway highlighted slow-varying measurements such as *tr_xmeas_1_last*, *tr_xmv_3_last*, and *tr_xmv_4_last*, which are indicative of gradual performance degradation and long-term process imbalance.

The separation of variable importance profiles across pathways confirms that the trend and cyclic decompositions capture complementary fault-relevant dynamics. Class-specific color patterns in the SHAP plots further show that certain features are strongly tied to individual fault classes—for example, *tr_xmeas_1_last* with Fault 5 and *cy_xmeas_21_last* with Fault 20—facilitating targeted root cause analysis. This combination of improved predictive performance and transparent feature attribution transforms the architecture from an opaque “black box” into a “glass-box” decision support system, aligning with the trust, interpretability, and safety requirements of Level 4 IA.

In summary, the dual-path trend–cyclic decomposition architecture not only improved aggregate classification metrics but also delivered targeted enhancements in fault classes that stand to benefit most from dynamic signal separation. The combination of higher accuracy, class-specific performance gains, and enhanced interpretability positions this approach as a strong candidate for Level 4 industrial autonomy deployment, bridging the gap between predictive analytics and operational decision-making.

8. CONCLUSION AND FUTURE SCOPE

This study demonstrated that a trend–cyclic dual-path architecture, applied to the Tennessee Eastman Process benchmark, can enhance fault classification performance while preserving operational transparency through explainability. By decomposing process signals into slow-varying trends and fast oscillations, the proposed method improved overall classification accuracy and macro-F1 over a baseline supervised pipeline, with notable gains for specific faults that exhibit distinct temporal dynamics. The integration of SHAP-based explainability at both the pathway and fused levels enabled clear identification of the most influential process variables for each fault, providing actionable insights for operators and aligning with the transparency and safety requirements of Level 4 industrial autonomy. The complementary nature of the trend and cyclic pathways confirms the value of multi-timescale analysis in industrial AI fault diagnosis.

Future work will focus on extending this framework in two key directions. First, real-time deployment will be explored by integrating the model with distributed control systems (DCS) and safety instrumented systems (SIS) to evaluate latency, robustness, and interoperability in live environments. Second, adaptive explainability mechanisms

will be developed to tailor model interpretations to different operator roles summarized alerts for shift operators, detailed root-cause paths for process engineers, and compliance-focused reporting for auditors. Collectively, these enhancements aim to transform the current interpretable model into a fully operational, human-centered decision-support layer within industrial AI ecosystems, strengthening trust, usability, and continuous learning toward Level 5 self-optimizing and self-healing operations.

REFERENCES

- Additional Tennessee Eastman Process Simulation Data for Anomaly Detection Evaluation - Harvard Dataverse.* (n.d.). Retrieved October 26, 2025, from <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/6C3JR1>
- AI Policy | Yokogawa Electric Corporation.* (n.d.-a). Retrieved October 26, 2025, from [https://www.yokogawa.com/about/ai-policy/Autonomous Operations for Process Industries? / ARC Advisory](https://www.yokogawa.com/about/ai-policy/Autonomous%20Operations%20for%20Process%20Industries?%20ARC%20Advisory). (n.d.-a). Retrieved October 26, 2025, from <https://www.arcweb.com/blog/autonomous-operations-future-process-industry>
- Autonomous Operations for Process Industries? | ARC Advisory.* (n.d.-b). Retrieved August 11, 2025, from <https://www.arcweb.com/blog/autonomous-operations-future-process-industry>
- IEC TR 63069:2019 | IEC.* (n.d.). Retrieved August 11, 2025, from <https://webstore.iec.ch/en/publication/31421>
- Industrial-grade Data Fabrics | ARC Advisory Group.* (n.d.). Retrieved August 12, 2025, from <https://www.arcweb.com/industry-concepts/industrial-grade-data-fabrics>
- ISA-95 Standard: Enterprise-Control System Integration.* (n.d.). Retrieved October 26, 2025, from <https://www.isa.org/standards-and-publications/isa-standards/isa-95-standard>
- Ohara, K. (2020). Yokogawa's Commitment to AI Technology. *Commitment to AI Technology Yokogawa Technical Report English Edition*, 63(1).
- Ohtani, T. (2020). Application of AI to Oil Refineries and Petrochemical Plants. *Yokogawa Technical Report English Edition*, 63(1).
- Regulation - EU - 2024/1689 - EN - EUR-Lex.* (n.d.). Retrieved August 12, 2025, from <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- Zhu, L., Cui, Y., Takami, G., Kanokogi, H., & Matsubara, T. (2020). Scalable reinforcement learning for plant-wide control of vinyl acetate monomer process. *Control Engineering Practice*, 97. <https://doi.org/10.1016/j.conengprac.2020.104331>

BIOGRAPHIES

Jayanth Balaji Avanashilingam is an Executive Engineer at the AI Lab, Co-Innovation Centre, Singapore Development Centre, Yokogawa Engineering Asia Pte. Ltd., Singapore. He leads cross-functional initiatives to develop scalable AI platforms for industrial automation, acting as a key link between customers and technical teams to translate business requirements into practical solutions and advance product development.

Bijuraj Pandiyath Velayudhan currently serves as Head of the AI Lab at the Co-innovation Centre, Singapore Development Centre, Yokogawa Engineering Asia Pte. Ltd., Singapore. He spearheads the rapid commercialization of AI applications across both cloud and edge platforms. With extensive experience as a digitalization and automation consultant, he has a proven track record in industrial automation R&D. He leads co-innovation initiatives, new business incubation, and proof-of-concept developments in ML and AI, overseeing the process from business case formulation through to successful market launch.