International Journal of Current Research and Applied Studies, (IJCRAS)

ISSN: 2583-6781

available at https://ijcras.com/

Volume 5 Issue 2 March-April 2026, Page 53-65

To cite this article: Dr. Srinivasa Rao Kosiganti, Karthikeya Jaghni and Sreehitha Kosiganti (2026). A SIMULATION-DRIVEN PREDICTIVE MAINTENANCE FRAMEWORK FOR PIPELINE AND WELL INFRASTRUCTURE USING SYNTHETIC SENSOR TIME-SERIES DATA AND AUTONOMOUS MAINTENANCE AGENTS, International Journal of Current Research and Applied Studies (IJCRAS) 5 (2): Article No. 144, Sub Id 253

A SIMULATION-DRIVEN PREDICTIVE MAINTENANCE FRAMEWORK FOR PIPELINE AND WELL INFRASTRUCTURE USING SYNTHETIC SENSOR TIME-SERIES DATA AND AUTONOMOUS MAINTENANCE AGENTS

Dr. Srinivasa Rao Kosiganti[1], Karthikeya Jaghni[2] and Sreehitha Kosiganti[3]

[1]Digixform Inc, 2-4-762/1, ROAD NO. 6B, NEW NAGOLE, HYDERABAD - 500035, INDIA, (Email: srinikosi@gmail.com).
[2]Undergraduate degree, University of Toledo in Ohio, Major: Computer Science Engineering Technology. Email: jaghnik@gmail.com
[3]Binghamton University, MS in Computer Science (AI & ML Track), 4400 Vestal Parkway East, Vestal, NY 13850, Currently pursuing MS in 3rd semester, Email: sreehithakosiganti@gmail.com; skosiganti@binghamton.edu

DOI : https://doi.org/10.61646/IJCRAS.vol.5.issue2.144

ABSTRACT

Predictive maintenance has emerged as a promising approach for improving the reliability and safety of pipeline and well infrastructure through continuous monitoring and data-driven decision-making. However, the development and evaluation of predictive maintenance systems in this domain are hindered by (i) the scarcity of labeled failure data, (ii) the nonstationary nature of sensor measurements, and (iii) the gap between predictive model outputs and actionable maintenance decisions.

This study presents a simulation-driven predictive maintenance framework that integrates synthetic multivariate sensor time-series data, machine learning based anomaly detection and remaining useful life (RUL) estimation, and an autonomous maintenance agent for risk- aware alert prioritization. The framework employs a configurable synthetic data generator to emulate realistic sensor dynamics and annotated failure events, enabling controlled and reproducible experimentation. Predictive models are trained to identify abnormal behavior and estimate degradation progression, while the autonomous agent supervises model outputs, monitors drift, schedules retraining, and translates predictive signals into prioritized maintenance alerts based on severity, urgency, and cost-consequence weights.

Simulation-based evaluation indicates that combining anomaly detection with RUL estimation yields complementary predictive evidence, and that agent supervision reduces false alarms while improving the relevance and prioritization of maintenance alerts. Overall, the study argues that a carefully designed, simulation-driven framework augmented by autonomous decision logic can support effective and interpretable maintenance decision- making for pipeline and well infrastructure under data-scarce and nonstationary conditions.

Keywords: Predictive Maintenance, Synthetic Sensor Time-Series Data, Anomaly Detection, Remaining Useful Life (RUL) Estimation, Autonomous Maintenance Agents

1. INTRODUCTION

Pipeline and well infrastructure constitutes a critical backbone of modern energy and resource distribution systems, supporting the transport of oil, gas, water, and other essential fluids across vast geographic regions. Because failures can lead to economic loss, environmental damage, and public safety risk, integrity management is a central requirement of pipeline and well operations. Concurrently, sensing and telemetry have expanded the volume and resolution of operational data available for monitoring, including pressure, temperature, flow rate, vibration, and spatial indicators.

Machine learning has been increasingly applied to predictive maintenance, particularly for time-series anomaly detection and prognostics, due to its capacity to model nonlinear and multivariate temporal behavior [1], [15]. In pipeline contexts, research has emphasized leak detection and related integrity monitoring, including methods using pressure/flow signals and distributed sensing (e.g., fiber-optic systems) to capture spatially localized signatures [5]–[7]. Despite these advances, the systematic development and evaluation of predictive maintenance methods for safety-critical infrastructure remains constrained by limited labeled failures, shifting operating conditions, and the operational challenge of converting model outputs into prioritized maintenance actions.

This paper proposes a simulation-driven predictive maintenance framework for pipeline and well infrastructure that integrates synthetic multivariate sensor time-series data, predictive models for anomaly detection and RUL estimation, and an autonomous maintenance agent for drift-aware governance and risk-aware prioritization. Synthetic data has gained traction as a research strategy for predictive maintenance when real failure datasets are restricted or sparse, provided that simulation assumptions are explicit and conclusions are bound accordingly [4]. In parallel, recent work on concept drift and ML operations highlights the need for systematic monitoring, retraining triggers, and human-auditable decision processes in nonstationary settings [12], [13].

The central thesis of this study is that a carefully designed simulation-driven framework— augmented by autonomous supervision—enables reproducible investigation of predictive maintenance strategies and improves system-level decision relevance (alert burden, prioritization quality, and drift robustness) beyond model-centric accuracy measures. The remainder of the paper is organized as follows: Section 2 defines the problem; Section 3 states study objectives; Sections 4–5 provide motivation and related work; Sections 6–8 describe the architectural framework, data structure, and solution workflow; Sections 9–11 present a case study, evaluation, and discussion; Section 12 outlines future scope; and Section 13 concludes.

2. PROBLEM STATEMENT

Pipeline and well systems operate in complex, dynamic environments where failures are rare but potentially catastrophic. Although continuous sensing provides rich operational data streams, translating these data into reliable, actionable predictive maintenance decisions remains challenging.

Severe data scarcity is a primary barrier. Major integrity events occur infrequently, and historical failure records are often incomplete, inconsistent, or inaccessible due to proprietary and regulatory constraints. Consequently, supervised learning is limited, benchmarking is difficult, and methods developed on convenience datasets may not reflect the true rarity and heterogeneity of real failure modes.

Nonstationarity further complicates deployment. Changes in operating regimes, environmental conditions, asset aging, and sensor calibration drift induce distribution shifts that can degrade model performance over time. Surveyed time-series anomaly detection methods note that thresholding, calibration, and evaluation become especially difficult under nonstationary and rare-event conditions [1]. Without explicit drift monitoring and model governance, systems may exhibit escalating false alarm rates or missed detections.

Finally, predictive maintenance systems frequently optimize model metrics without formalizing downstream decision requirements. High detection accuracy does not guarantee operational value because maintenance decisions must balance uncertainty, competing priorities, limited resources, and asymmetric costs of false alarms versus missed detections. Static thresholding often yields excessive low-value alerts, and many approaches do not define retraining policies or decision auditability, despite the recognized importance of drift- aware ML operations [12], [13].

These challenges define a system-level problem: designing and evaluating predictive maintenance for pipeline and well infrastructure under limited labeled failures, nonstationary sensor behavior, and complex cost-risk trade-offs. Addressing this problem requires an integrated framework that supports controlled experimentation, combines complementary predictive signals, and embeds them within an interpretable and adaptive decision process.

3. OBJECTIVES

3.1 DEVELOP A SYNTHETIC MULTIVARIATE SENSOR TIME-SERIES GENERATOR

Design a configurable simulator for pipeline and well telemetry that emulates realistic baseline dynamics, sensor coupling, noise, drift, and explicitly annotated failure and degradation events.

3.2 DEVELOP PREDICTIVE MODELS FOR DETECTION AND PROGNOSIS

Train and evaluate time-series forecasting, anomaly detection, and remaining useful life estimation models under data scarcity and severe class imbalance, using temporally valid splits and rare-event metrics.

3.3 INTEGRATE COMPLEMENTARY PREDICTIVE SIGNALS

Combine anomaly scores and RUL estimates into a unified evidence representation that captures both abnormality and urgency, enabling more informative maintenance reasoning than any single signal.

3.4 IMPLEMENT AN AUTONOMOUS MAINTENANCE AGENT FOR GOVERNANCE

Implement an agent that monitors model outputs and health, detects drift, schedules retraining or recalibration, and produces prioritized alerts with confidence and rationale for auditability.

3.5 EVALUATE SYSTEM-LEVEL OPERATIONAL EFFECTIVENESS

Assess alert burden, detection delay, false alarm reduction, prioritization quality, and drift robustness to evaluate decision relevance, rather than reporting model accuracy alone.

3.6 ANALYZE LIMITATIONS AND GENERALIZABILITY

Identify constraints of simulation-based evidence, discuss assumptions affecting transfer to real infrastructure, and outline practical and research extensions for deployment- oriented work.

4. MOTIVATION

Safety and consequence asymmetry: Pipeline and well failures are rare but high-impact events. Effective integrity management therefore benefits from early warning and risk- sensitive prioritization that reduces missed detections without creating excessive false alarms.

Data and access limitations: Labeled failure data are difficult to obtain at scale, and publicly available pipeline/well failure telemetry is limited. Synthetic data generation provides a reproducible mechanism to study model behavior and decision workflows when real failure datasets are unavailable, provided assumptions are explicitly stated [4].

Nonstationary operating conditions: Real infrastructure operates under changing regimes and environmental conditions. Drift-aware monitoring and retraining are increasingly recognized as essential elements of reliable ML systems, yet are often absent from predictive maintenance prototypes [12], [13].

From prediction to action: Maintenance decisions require more than anomaly scores. They require prioritization by urgency, severity, uncertainty, and cost. Autonomous agents offer a structured approach to supervising model lifecycles and translating predictive signals into operationally meaningful alerts [9]–[11].

5. RELATED WORK

5.1 Predictive maintenance in pipeline and well monitoring

Pipeline integrity research has emphasized leak detection and localization using pressure/flow signatures, as well as distributed sensing modalities capable of capturing spatially localized anomalies [5]–[7]. Comprehensive reviews catalog detection standards, method taxonomies, and comparative limitations across approaches [5]. These studies provide important foundations but often assume access to historical failures or focus on single failure modes.

5.2 Deep time-series anomaly detection

Deep learning approaches to time-series anomaly detection span reconstruction-based methods (e.g., autoencoders), forecasting-based methods, self-supervised representation learning, and hybrids. A recent ACM Computing Surveys paper provides a detailed taxonomy and discusses recurring challenges in thresholding, non-stationarity, and evaluation under severe class imbalance [1]. These findings motivate using continuous scores, temporal persistence logic, and rare-event metrics (e.g., AUPRC).

5.3 Prognostics and remaining useful life estimation

Prognostics research synthesizes pipelines from data acquisition to degradation modeling and RUL prediction, emphasizing uncertainty and the effects of operating condition changes [14], [15]. Recent reviews also highlight deep learning-based RUL approaches and caution that absolute RUL precision is often limited; instead, relative urgency ranking and trend stability are frequently more actionable [8].

5.4 Synthetic data for predictive maintenance

Synthetic data generation for predictive maintenance has been surveyed across augmentation, generative models, physics-based simulation, and hybrid approaches. Systematic reviews emphasize that synthetic data is useful for benchmarking and stress testing when modeling assumptions and limitations are explicit [4]. For data-driven generation, TimeGAN is a canonical approach aimed at preserving temporal dynamics and feature correlations [3], and survey work summarizes evaluation methods and privacy considerations for time-series GANs [2].

5.5 Autonomous agents and drift-aware governance

Agent-based approaches, including reinforcement learning, have been applied to maintenance policy optimization under degradation and repair assumptions [10], [11].

Parallel work in concept drift and ML operations stresses the importance of systematic monitoring, expert oversight, and retraining triggers to maintain reliability as data distributions shift [12], [13]. This study integrates these threads into a single framework oriented toward system-level decision relevance.

6. ARCHITECTURAL FRAMEWORK

6.1 Overview

The framework comprises four layers: (i) a synthetic data simulation layer, (ii) a predictive modeling layer, (iii) an autonomous maintenance agent layer, and (iv) a decision and reporting layer. This modular decomposition supports reproducibility and component-wise ablation.

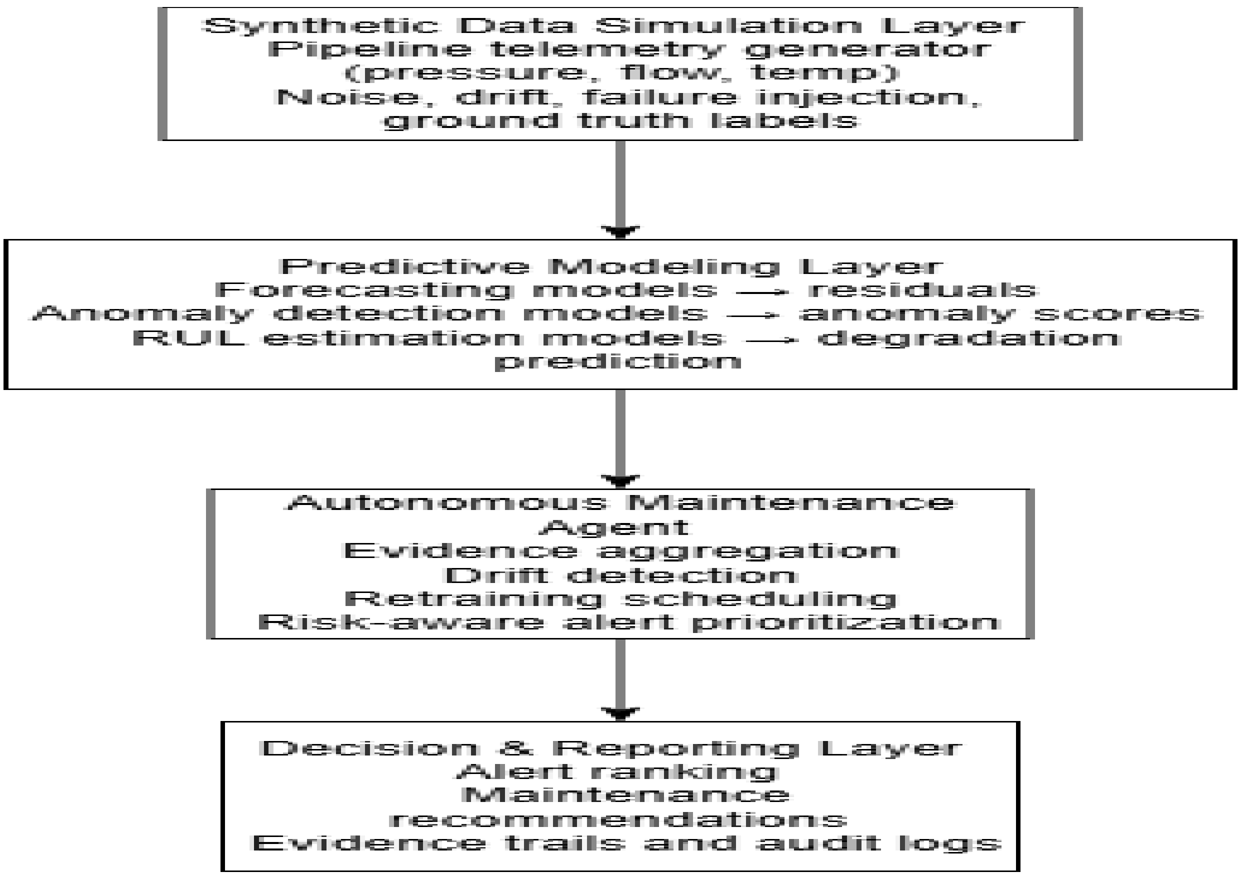

Figure No. 1

6.2 Synthetic data simulation layer

The simulator generates synchronized multivariate sensor streams (pressure, flow rate, temperature, vibration, and positional indicators) with configurable noise and drift.

Failures are injected with labeled onset time, type, progression profile, and severity, yielding ground-truth anomaly intervals and RUL targets.

6.3 Predictive modeling layer

Forecasting models learn baseline dynamics and produce residuals; anomaly detectors compute continuous anomaly scores from residual and reconstruction patterns [1].
Prognostic models estimate RUL (point and/or interval) to capture urgency [14], [15].

6.4 Autonomous maintenance agent layer

The agent aggregates model evidence over time, suppresses transient deviations, escalates persistent conditions, monitors drift and model health, schedules retraining, and ranks alerts using a composite risk score incorporating severity, urgency, confidence, and cost- consequence weights. This governance loop aligns with drift-aware ML operations [12], [13].

6.5 Decision and reporting layer

The system outputs prioritized alerts with supporting evidence (score trajectories, contributing sensors, and uncertainty indicators) and maintains an audit log to support safety-critical review and post-event analysis.

7. DATA AND SAMPLE STRUCTURE

7.1 Assets, scenarios, and sampling

The dataset comprises multiple simulated assets, each representing a pipeline segment or well. All sensor variables are sampled at uniform intervals and synchronized across channels.

7.2 Variables

Each asset includes pressure, flow rate, temperature, vibration, and positional indicators, reflecting common sensing modalities used in integrity monitoring and distributed sensing contexts [6], [7].

7.3 Baseline dynamics, noise, and drift

Baseline behavior is generated via autocorrelated stochastic processes with cross-sensor coupling. Measurement noise is additive with configurable variance. Nonstationarity is introduced via drift components that shift means and variances over time, representing regime changes and sensor calibration effects [1], [12].

7.4 Failure injection and labels

Failure modes (e.g., leak-like pressure loss, blockage-like flow restriction, vibration anomalies) are injected with onset, progression, and severity. Labels include anomaly intervals, failure-type identifiers (optional), and RUL targets computed deterministically from annotated failure time.

Failure Mode	Pressure Behavior	Flow Behavior	Vibration Behavior	Detection Difficulty
Leak	Gradual pressure drop	Slight decrease	Minor fluctuation	Medium
Blockage	Pressure increase	Flow decrease	Stable	Medium
Corrosion	Slow pressure decline	Increased variability	Growing vibration	High

Mechanical Fault	Sudden change	Irregular flow	Sharp vibration spikes	Low

Table No. 1

7.5 Temporal splits and rarity

Data are partitioned using time-aware train/validation/test splits to prevent leakage. Failures are rare relative to normal operation, preserving class imbalance consistent with real systems. All parameters and random seeds are recorded to ensure reproducibility and sensitivity analysis [4].

8. SOLUTION WORKFLOW

8.1 Data generation and annotation

Synthetic telemetry is generated under configurable operating regimes, with explicit injection and labeling of failures and degradations.

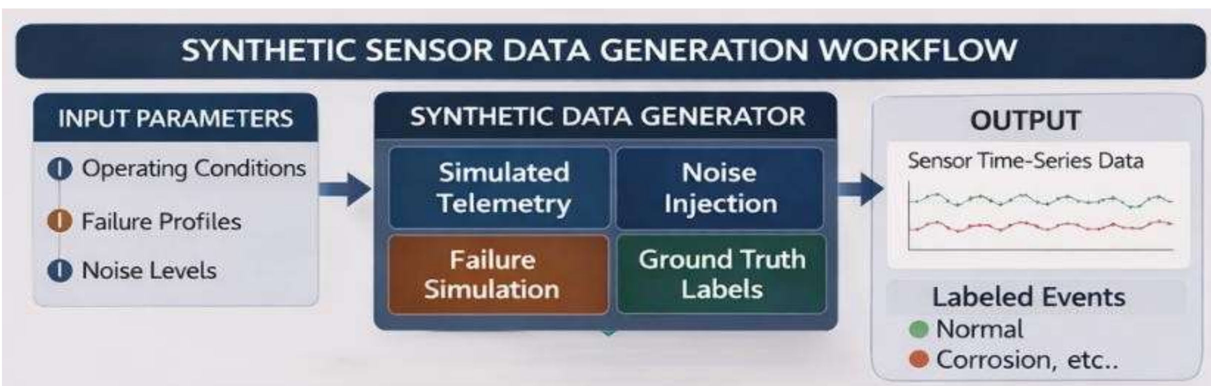

Figure No. 2

8.2 Model training

Forecasting models are trained to capture baseline temporal structure; anomaly detection models transform residual and reconstruction behavior into continuous anomaly scores

[1]. RUL models are trained to map degradation signatures to urgency indicators, consistent with prognostics pipelines [14], [15].

8.3 Streaming inference

Models consume rolling windows to produce anomaly scores, RUL estimates, and uncertainty indicators. Outputs are normalized and smoothed to reduce sensitivity to transient noise.

8.4 Autonomous supervision and governance

The maintenance agent aggregates evidence across time, incorporates persistence and confidence logic, and emits prioritized alerts. Drift indicators are monitored continuously; sustained shifts trigger retraining or recalibration policies consistent with drift-aware operations [12], [13].

8.5 Risk- and cost-aware prioritization

Alerts are ranked using a composite priority score integrating anomaly severity, urgency, confidence, and cost-consequence weights, yielding actionable maintenance queues for constrained operational resources.

Component	Description	Weight
Anomaly Severity	Magnitude of anomaly score	0.35
Urgency	Remaining Useful Life estimate	0.30
Confidence	Model certainty or ensemble agreement	0.20
Cost Consequence	Infrastructure impact risk	0.15

Table No. 2

9. CASE STUDY

A representative case study illustrates the framework on a simulated pipeline segment instrumented with pressure, flow, temperature, vibration, and positional sensing. The scenario models progressive internal corrosion leading to leak-like behavior, selected because early-stage signatures are subtle and multivariate.

9.1 Normal operation phase

The asset operates within nominal ranges with stochastic variability and low-amplitude drift.

9.2 Degradation onset and progression

At a predefined onset time, pressure exhibits gradual decline, flow shows increasing variance, and vibration signals display emerging instability, while individual sensor readings may remain within conventional thresholds.

9.3 Agent-mediated response

Anomaly scores rise gradually and exceed alert thresholds only when evidence persists. RUL estimates decrease as degradation accelerates. The agent suppresses transient deviations, issues a medium-priority alert when multivariate evidence becomes sustained, and escalates to high priority as urgency increases. In the simulated setting, escalation occurs before the annotated failure time, illustrating a balance between early warning and false-alarm control.

10. EVALUATION AND RESULTS (SIMULATION-BASED)

Evaluation is conducted across multiple synthetic scenarios with varying noise, drift intensity, and failure modes. Temporal splits are used to emulate deployment. Metrics reflect rare-event monitoring: precision/recall and area under the precision recall curve (AUPRC) for anomaly detection, detection delay, and false alarm rate for alerting, and mean absolute error (MAE) for RUL estimation [1], [15].

Scenario	AUPRC	Detection Delay	False Alarm Rate	RUL MAE
Baseline	0.82	6 min	3%	9.4 h
High Noise	0.74	11 min	7%	12.6 h
Drift Scenario	0.77	10 min	5%	11.8 h
Agent Supervision	0.84	7 min	2%	9.1 h

Table No. 3

10.1 Detection performance

Multivariate anomaly detection identifies injected failures effectively when failures manifest as cross-sensor inconsistencies. Abrupt failures yield shorter detection delays than gradual degradations. False alarms occur primarily during high-noise or regime- change periods, motivating downstream evidence aggregation.

10.2 Prognostic performance

RUL estimation exhibits moderate absolute accuracy. However, the relative ranking of assets by urgency remains stable across scenarios, supporting the use of RUL as a prioritization signal rather than a deterministic forecast [8]. Error variance increases near end-of-life, consistent with prognostic uncertainty [14], [15].

10.3 System-level impact of agent supervision

Compared with anomaly-score-only alerting, the autonomous agent reduces false- positive alerts by suppressing transient anomalies and requiring sustained evidence before escalation. Risk-aware prioritization aligns alerts more closely with severity and urgency than single-score ranking. Drift-aware retraining triggers mitigate longer-term degradation under nonstationary conditions [12], [13].

Task	Metric	Description
Anomaly Detection	Precision / Recall	Detection accuracy for rare events
Detection Quality	AUPRC	Performance under severe class imbalance

Alerting System	False Alarm Rate	Frequency of incorrect alerts
Response Time	Detection Delay	Time between anomaly onset and alert
Prognostics	MAE (RUL)	Accuracy of remaining life estimates

Table No. 4

11. DISCUSSION

The simulation-based results indicate that carefully parameterized synthetic multivariate time series can support reproducible evaluation of predictive maintenance workflows under explicit assumptions. The variability in detectability across failure modes reinforces findings from anomaly detection surveys that temporal profile and cross-sensor coupling shape model behavior and evaluation outcomes [1].

RUL estimation provides complementary temporal context. Consistent with prognostics literature, uncertainty increases near failure, suggesting that RUL is best interpreted as a decision-support indicator rather than a precise failure-time forecast [8], [14], [15].

Consequently, integrating anomaly severity, urgency, and confidence yields more operationally meaningful prioritization than relying on any single signal.

The autonomous maintenance agent provides governance benefits at the system level. By aggregating evidence, suppressing transient deviations, and incorporating drift-aware retraining logic, the agent reduces alert fatigue and improves prioritization coherence. This aligns with recent work arguing for expert-in-the-loop drift monitoring and systematic governance mechanisms in deployed ML systems [12], [13].

Limitations include reliance on synthetic data, which cannot fully capture the complexity of real-world pipeline and well environments, and the use of rule- and utility-based agent logic rather than learned policies. While interpretability is strengthened by explicit logic, learning- based maintenance policies (e.g., reinforcement learning) may improve adaptation under richer feedback signals but introduce additional governance and transparency challenges [10], [11].

12. FUTURE SCOPE OF STUDY

Future work should investigate hybrid simulation that combines simplified physics-based models with stochastic and data-driven components to improve realism while preserving interpretability [4]. Extending the agent to learning-based policies (e.g., reinforcement learning) may optimize maintenance actions under more complex cost structures, though careful attention to transparency and data requirements is required [10], [11].

Scaling to portfolio-level planning across multiple assets, incorporating shared resource constraints, and modeling cascading risk across network components would increase practical relevance. Additional research should evaluate domain adaptation strategies using hybrid datasets that combine limited real sensor data with synthetic augmentation, enabling external validation and more credible generalization claims. Finally, improving human agent interaction and explanation interfaces can strengthen operator trust and support auditability in safety-critical deployment settings.

13. CONCLUSION

This study presented a simulation-driven predictive maintenance framework for pipeline and well infrastructure that integrates synthetic multivariate sensor time-series data, anomaly detection and RUL estimation models, and an autonomous maintenance agent for drift-aware governance and risk-aware alert prioritization. The framework addresses practical barriers— scarce labeled failures, nonstationary sensor behavior, and the prediction-to-action gap—by emphasizing reproducible experimentation and system-level evaluation.

Results suggest that anomaly detection and RUL estimation provide complementary evidence and that autonomous supervision reduces false alarms while improving prioritization coherence. Drift monitoring and retraining logic further support sustained performance under changing conditions, underscoring the importance of governance in predictive maintenance systems [12], [13]. While constrained by simulation assumptions, the proposed framework provides a methodologically grounded foundation for future research integrating hybrid simulation, real-world validation, and more advanced decision policies.

14. REFERENCES

[1] Z. Zamanzadeh Darban, G. I. Webb, S. Pan, C. C. Aggarwal, and M. Salehi, “Deep learning for time series anomaly detection: A survey,” ACM Computing Surveys, 2024. doi: 10.1145/3691338.

[2] E. Brophy, Z. Wang, Q. She, and T. Ward, “Generative adversarial networks in time series: A survey and taxonomy,” ACM Computing Surveys, 2023. doi: 10.1145/3559540.

[3] J. Yoon, D. Jarrett, and M. van der Schaar, “Time-series generative adversarial networks,” in Advances in Neural Information Processing Systems (NeurIPS), vol. 32, 2019, pp. 5509–5519.

[4] W. Nieminen, J. Lindström, and J. Lönnqvist, “Synthetic data for predictive maintenance: A systematic review and framework for Industry 4.0 applications,” Journal of Intelligent Manufacturing, 2026. doi: 10.1007/s10845-026-02795-6.

[5] J. Yuan, W. Mao, C. Hu, J. Zheng, D. Zheng, and Y. Yang, “Leak detection and localization techniques in oil and gas pipeline: A bibliometric and systematic review,” Engineering Failure Analysis, vol. 146, Art. 107060, 2023. doi: 10.1016/j.engfailanal.2023.107060.

[6] D. Wang, J. Li, and H. Sun, “Research on gas pipeline leakage model identification driven by digital twin,” Systems Science & Control Engineering, 2023. doi: 10.1080/21642583.2023.2180687.

[7] Z. Peng, J. Li, X. Zhang, and X. Bao, “Distributed fiber sensor and machine learning data analytics for pipeline protection against extrinsic intrusions and intrinsic corrosions,” Optics Express, vol. 28, no. 21, pp. 31538–31554, 2020. doi: 10.1364/OE.402838.

[8] Y. Pan, S. Kang, L. Kong, J. Wu, and Y. Yang, “Remaining useful life prediction methods of equipment components based on deep learning for sustainable manufacturing: A literature review,” AI EDAM, 2025. doi: 10.1017/S0890060424000271.

[9] F. Piccialli, D. Chiaro, S. Sarwar, and D. Cerciello, “AgentAI: A comprehensive survey on autonomous agents in distributed AI for Industry 4.0,” Expert Systems with Applications, 2025. doi: 10.1016/j.eswa.2025.128404.

[10] A. P. Marugán, J. M. Pinar-Pérez, and F. P. García Márquez, “A reinforcement learning agent for maintenance of deteriorating systems with increasingly imperfect repairs,” Reliability Engineering & System Safety, vol. 252, 2024. doi: 10.1016/j.ress.2024.110466.

[11] K. Varalakshmi and J. Kumar, “Optimized predictive maintenance for streaming data in industrial IoT networks using deep reinforcement learning and ensemble techniques,” Scientific Reports, 2025. doi: 10.1038/s41598-025-10268-8.

[12] J. Leest, C. Raibulet, I. Gerostathopoulos, and P. Lago, “Expert Monitoring: Human-Centered Concept Drift Detection in Machine Learning Operations,” in NIER@ICSE, 2024. doi: 10.1145/3639476.3639771.

[13] L. Serfilippi, A. Bujari, and A. Corradi, “Context-aware drift detection for quality-aware data-driven smart city digital twins,” in Proc. ACM GoodIT, 2025. doi: 10.1145/3748699.3749824.

[14] A. K. S. Jardine, D. Lin, and D. Banjevic, “A review on machinery diagnostics and prognostics implementing condition-based maintenance,” Mechanical Systems and Signal Processing, vol. 20, no. 7, pp. 1483–1510, 2006. doi: 10.1016/j.ymssp.2005.09.012.

[15] Y. Lei, N. Li, L. Guo, N. Li, T. Yan, and J. Lin, “Machinery health prognostics: A systematic review from data acquisition to RUL prediction,” Mechanical Systems and Signal Processing, vol. 104, pp. 799–834, 2018. doi: 10.1016/j.ymssp.2017.11.016.