Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A Multi-Task Learning Framework for Predictive Maintenance and Decision Support in Sustainable Manufacturing Systems

AMIR KARBASSI YAZDI

akarbassiy@academicos.uta.cl

Universidad de Tarapacá: Universidad de Tarapaca

Wajahat Ali

Yong Tan

Research Article

Keywords: decision support, predictive maintenance, smart manufacturing, production planning, supply chain resilience, robust analytics

Posted Date: November 4th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10973066/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

A Multi-Task Learning Framework for Predictive Maintenance and Decision Support in Sustainable Manufacturing Systems

Amir Karbassi Yazdi
Departamento de Ingeniería Industrial y de Sistemas
Universidad de Tarapacá, Arica, Chile
Email: akarbassiy@academicos.uta.cl

Wajahat Ali
Department of Statistics & Operations Research
Aligarh Muslim University, Aligarh, India
Email: gk2721@myamu.ac.in

Yong Tan
[e]School of Management
University of Bradford, Bradford, West Yorkshire, United Kingdom
Email: y.tan9@bradford.ac.uk

Abstract

This study develops a data-driven decision-support framework for sustainable manufacturing by integrating multi-task learning with sensitivity and robustness analyses in a unified environment. It addresses production forecasting, maintenance detection, and efficiency classification within a shared feature space, providing a comprehensive view of manufacturing performance and decisions. Evaluated on an industrial dataset with 1,000 observations and 36 features, results show efficiency differences across tasks. The classification method achieved nearly perfect accuracy (98.5%) in estimating production for each task, but explained only 71% of the variability in forecasting volume. It failed to detect maintenance events, yielding no F1 score, due to feature limitations in identifying rare events, leading to a performance score of 56.4/100. Challenges from class overlap and feature-space limitations highlight the need for improved imbalance-aware models. Sensitivity analysis showed stable performance with more data and moderate declines under noise, indicating reliable learning. Robustness diagnostics identified predictable responses and sensitivity to random initialization, emphasizing the importance of stability. The framework enhances transparency, interpretability, and managerial relevance in manufacturing analytics, with practical implications for planning, maintenance, efficiency, and resilient supply chains. It advances the decision sciences and smart manufacturing fields by offering a scalable, interpretable tool for robust industrial analytics and a foundation for future integration with prescriptive optimization and digital twin systems.

Keywords: decision support; predictive maintenance; smart manufacturing; production planning; supply chain resilience; robust analytics.

1. Introduction

Accelerating advances in high-tech manufacturing are redefining how industry operates as businesses transition into the next era of Industry 4.0 (Hughes et al., 2022). Industry 4.0, enabled by cyber-physical systems, intelligent automation, connected machining lines, and real-time data sharing, is a revolutionary approach that will transform the operational efficiency, sustainability, and resilience of the manufacturing ecosystem (Qudus et al., 2025). Global projections suggest the amount of connected industrial assets is set to grow beyond 50 billion by 2030, and data generated by production operations is growing at an annual rate of over 25%, demonstrating the necessity to orchestrate intelligent systems that can deliver actionable intelligence from noisy, fast-paced operational data (Edquist et al., 2021). In this context, the introduction of machine learning (ML) into manufacturing performance, including predictive modeling, anomaly detection, and decision optimization, has recently gained traction to achieve higher performance and maintain a competitive advantage (Fowowe & Adedapo, 2025). Among such prognostics-based maintenance approaches, PdM (predictive maintenance) has become prominent in practice, as unplanned equipment failures account for approximately 20–30% of annual manufacturing losses, with downtime costs estimated at $50–150 billion worldwide (Kremen, 2025). Typically, preventative maintenance relies on milliseconds or cycles but is largely unable to capture potential risks arising from a machine's real-time response behavior (Mohan et al., 2025). ML-based PdM offers a data-driven approach to predicting potential failures by learning patterns from historical and sensor data. Transitioning to this model will allow companies to prolong asset life, reduce downtime, prevent unnecessary maintenance, and increase safety for their personnel: all elements vital for operational excellence and reliability in today's industrial enterprises. (Chen et al., 2023).

Meanwhile, supply chain incidents have become increasingly common and severe in recent years, driven by a range of causes, including political instability, environmental volatility, pandemic lockdowns, and resource limitations (Yang & Fu, 2025). Recent global studies have shown that more than 80% of manufacturing firms experienced one or more significant supply chain disruptions over the past five years, and that the costs associated with supply chain failures worldwide exceed $1.2 trillion (Khan & Emon, 2025). With more complex, interlocked, and digitally connected supply chains, the capability to predict discrepancies has become essential at all levels to adjust production schedules and resource allocation, and to maintain continuity, thereby ensuring long-term sustainability (Alfaqiyah et al., 2025). ML is helping to drive sustainable supply chain resilience through features such as adaptive forecasting, risk-aware inventory management, proactive supplier assessment, and demand–capacity synchronization, which reduce waste and encourage circularity, thereby benefiting environmental and economic performance and sustainable development (Karim & Saha, 2025).

The prevailing framework is a comprehensive, uniform umbrella for PdM and resilient supply chain operations in sustainable manufacturing under Industry 4.0 (Keshar, 2025). In reality, however, the two components in the existing literature are mostly treated in isolation and fail to account for their interactional patterns and underlying correlations (Shehun, 2025). ML solves this problem and has a clear future in modeling production outputs, unified detection of potentially faulty equipment, and prediction of operating efficiency modes, thereby enabling analysis of complete plant operations that influence procurement stability, inventory placement, and resource planning (Carbone et al., 2025).

Our research makes value-added contributions to the emerging digital manufacturing literature by proposing a hybrid ML framework for an Industry 4.0 scenario that simultaneously predicts production volume, identifies maintenance events, and classifies efficiency states. This research will enhance operational reliability and sustainable supply chain resilience through multitask learning, calibrated uncertainty quantification, and cost-sensitive decision optimization. By bridging the gap between PdM and data-driven supply chain decisions, this paper presents a holistic and scalable solution for contemporary manufacturing firms aiming to optimize productivity, mitigate risks, and achieve a sustainable competitive advantage.

The significance of this research lies in its attempt to address the growing complexity of modern production systems as they transition to Industry 4.0. Modern-day factories generate vast amounts of plant, maintenance, and supply chain data; however, organizations still struggle to derive actionable intelligence to enhance the world's reliability, sustainability, and resilience (Owusu-Berko, 2025). Inadvertent equipment downtime alone accounts for nearly a third of annual productivity losses worldwide. At the same time, global supply chain crises have resulted in costs exceeding trillions of dollars for industries over the past decade (Islam et al., 2025). These are challenges that enterprise-level analytical tools should address to predict machine behavior, stabilize production output, and improve supply chain resilience (Xu & Bo, 2024). By fusing ML and uncertainty quantification with a cost-based decision model, this study presents a swift and applicable resolution aligned with the strategic core interest areas of innovative, sustainable, data-driven manufacturing (Hossain et al., 2025).

However, several shortcomings still undermine the effectiveness of ML in industrial applications, despite significant advances in the field. The majority of existing literature on PdM and supply chain resilience has been conducted in isolation, failing to account for the significant interrelationships among equipment health, production efficiency, and material availability. Moreover, traditional analytics largely depend on single-task models that generate one prediction at a time, leaving the rich relations among different tasks unused. Similarly problematic is the overwhelming reliance among practicing hydrologists on point predictions that do not account for uncertainty, leading to poor decisions when runoff variability is significant. Previous studies also rarely account for cost asymmetries (the weight placed on the values of

missed failures compared to that of unnecessary maintenance) when determining decision thresholds. Moreover, few papers quantify robustness to noise, temporal drift, and model instability, even though industrial datasets are typically noisy, imbalanced, and non-stationary. Efforts to establish sustainable supply chains are insufficiently addressed in the literature, resulting in isolated analytics solutions that fail to provide comprehensive support for end-to-end operational decision-making.

It introduces a hybrid multitask machine learning framework that simultaneously predicts production volume, maintenance events, and efficiency regimes using a standard encoder with dedicated output heads. Such a unified structure can effectively learn cross-task relations and achieve significantly higher accuracy and stability than single-task approaches. The work also advances methodological rigor by employing dual-uncertainty quantification, which combines quantile regression with conformal prediction to produce calibrated predictive intervals for risk-aware planning. To make predictive insights actionable, we develop cost-sensitive decision procedures to determine maintenance thresholds and resource allocation under asymmetric penalties and arbitrary action/maintenance strategies, based on the full predictive distribution. In this paper, the robustness of random seeds, regularization parameters, synthetic noise amplitude, feature subsets, and temporal window structures is examined to assess the model's stability and prepare it for practical application. The system has the potential to help supply chain sustainability through resource optimization, equipment condition monitoring, and production prediction. The above methods are essential for the sustainable and effective operation of a manufacturing process, thereby greatly benefiting production, reducing costs, increasing efficiency, and reducing downtime.

The remainder of this paper is organized as follows: Section 2 reviews the relevant literature on Industry 4.0, digital transformation in manufacturing, predictive maintenance, machine learning–based production forecasting, and supply chain sustainability and resilience. Section 3 presents the methods employed in this study. Section 4 describes the research methodology, including the dataset, variables and measurements, research goals and applications, data integrity and completeness considerations, and preliminary analyses of signals and dependencies. Section 5 presents the data analysis and discusses the empirical findings. Section 6 outlines the managerial implementation and practical implications of the proposed approach in industrial environments. Finally, Section 7 concludes the paper by summarizing the main findings, discussing the study’s limitations, and identifying directions for future research.

2. Literature review

Section 2 reviews the literature related to Industry 4.0, predictive maintenance, machine learning–based production forecasting, and supply chain resilience. It discusses the role of digital transformation technologies, including IoT(Internet of Things), cloud computing, big data analytics, and intelligent systems, in modern manufacturing environments. The section also examines recent advances in predictive maintenance and machine learning techniques for improving operational efficiency, production forecasting,

and decision-making. In addition, the literature highlights the growing importance of sustainability and resilience in Industry 4.0-enabled supply chains and identifies gaps in the integration of predictive maintenance, production forecasting, and supply chain optimization frameworks.

2.1 Industry 4.0 and the digital transformation of manufacturing

Realization of this paradigm in the manufacturing environment can improve sustainable and reliable operations (i.e., lower waste generation and minimizing outages). It might also help in designing more sustainable and reliable industrial production systems. As such, a fast-moving field has emerged, underpinned by the pace of technological progress and sustained research. For instance: autonomous intelligent systems, cloud computing, edge computing, big data analytics, and Internet of Things (Mishra et al., 2025). We need to move away from segmented data and processing to connected, active, contextual manufacturing systems. Numerous significant institutional transformations played out over time, collectively referred to as Industry 4.0. The idea of a new digital workplace was first conceptualized in Industry 4.0, where machines, sensors, and production assets interact openly and efficiently, enabling flexible operations and maximum intelligence through digitization (Dwivedy et al., 2025). There are already numerous digital technologies identified in the literature, including intelligent self-optimization of the production cycle, decision-making from multiple locations, and real-time process visibility (Li et al., 2025).

In this field of industrialization, digital transformation is a key aspect of Industry 4.0 alongside sustainable, productive, and efficient industries (Singh et al., 2025). Based on industry reports and certain academic works, interconnected factory shop floors have been reported to decrease operating costs by 10 to 20% and improve equipment efficiency by around 30% through better monitoring, predictive analysis, and automation (Lin et al., 2023). It is the same thing, and these results are consistent. Remote surveillance or data monitoring allows for service prediction, early detection, and business coordination. In this context, real-time data accessed from on-site installations or hardware-connected sensors is critical (Karim, 2025). It is a prime location for homes, shops, hotels, housing developments, and entertainment, and it benefits from environmental sustainability and sanitation. At the same time, there is potential for strength in these areas, but Industry 4.0 provides insights into and contributes to data on industrial production in terms of magnitude, diversity, and acceleration linked to complexity (Raptis et al., 2019). The emergence of new analytic models is crucial to the insights we can extract, as we have developed new production systems that use them at different times and intensities (Chang et al., 2025). Such nonlinearity and low failure rates of industrial systems cannot be considered intrinsic in a conventional statistical framework (Kafunah, 2024).

Industry 4.0 heavily depends on machine learning (ML), particularly for anomaly detection, demand forecasting for predictive maintenance (PdM), and supply chain risk assessment (Nezianya et al., 2024). (Kaur, 2025). found that integrating machine learning into Industry 4.0 is important for the transition from raw data to smart systems that automate decision-making. That, in itself, creates a future e-

manufacturing ecosystem that is "intrinsically more robust and sustainable." Sodhro et al. (2020) investigated the impact of IoT-based automation on production efficiency, finding higher throughput and lower energy consumption. Shamim (2025) demonstrated that cyber-physical integration facilitates machine-to-machine coordination, thereby reducing unplanned downtime. Ziari & Taleizadeh (2025) discussed how digital twins are used for real-time optimization and predictive scenario modeling to respond more effectively to a changing manufacturing system.

According to Duan & Da Xu (2024), the introduction of big data analytics inherently enhances Industry 4.0 capabilities. As Zhou (2025) asserts, large-scale industrial systems provide the necessary support for analysis, including anomaly detection, process optimization, and early failure prediction. Trigka & Dritsas (2025) illustrated a cloud–edge hybrid analytics framework, demonstrating the capacity of decentralized processing to reduce response times for processes requiring immediate attention from industry actors. In his analysis of intelligent sensor networks, Lakhouit (2025) demonstrated that intelligent monitoring significantly reduces operational waste overall and enhances quality control. Other studies have emphasized the sustainability aspects of Industry 4.0. Niukkanen (2024) found that digital traceability could improve resource allocation and reduce work loss in manufacturing. Pallaprolu (2025) reported that AI-driven automated processes help reduce emissions by ensuring optimal machine operating conditions. Vetrivel et al. (2026) observed that real-time life-cycle embedded in smart factories promotes green manufacturing process planning, an overall reduction in policies and footprint due to insights.

Industry 4.0 also enhances strategic decision-making and the processes that shape companies' futures. Specifically, Belhadi et al. examine how this technology facilitates adaptive production strategies via intelligent automation. Fagbore et al. (2024) investigated the potential of industrial networks to enhance operational direction by developing a faster approach to information exchange. Pietsch et al. (2024) argued that digital ecosystems make root-cause and action-taking processes more accessible. Lăzăroiu et al. (2022) explore how it increases predictive capabilities. However, as we become more connected, it also becomes more complex to handle all the data. Businesses now have to deal with vast, multidimensional, and time-dependent data that exceeds the capabilities of traditional analytical methods (Data, 2025). Zhou et al. (2025) emphasized that traditional statistical methods are insufficient to capture nonlinear machine actions and intricate crash mechanisms. As a consequence, advanced ML models have become indispensable for analyzing industrial big data, thereby supporting PdM, quality control, demand forecasting, and even supply chain resilience (Obioha Val et al., 2024). As highlighted in leading industrial studies, intelligence-driven SQMS and ISQMS-based systems can reduce operational costs by 10–20% and increase asset utilization by up to 30%, thereby supporting the transformational impact of Industry 4.0-driven digitalization (Ghobakhloo et al., 2025).

2.2 Predictive maintenance in modern manufacturing

PdM has evolved into one of the main pillars of Industry 4.0, redefining how manufacturers monitor equipment health and maximize its operational lifetime (Baroud et al., 2024). Unlike conventional preventive maintenance (PM), which relies on predefined schedules and historical mean statistics, PdM uses ML, real-time sensor data, and prognostics to analyze patterns of equipment degradation and predict when a failure is likely to occur, thereby preventing it (Khan et al., 2025). Hardware and software are increasingly equipped with sensors, IoT Modules, and cyber-physical monitoring devices that enable predictive maintenance, enabling proactive intervention to avoid unscheduled downtime and extend the life of vital equipment in shops (Gomaa, 2025). The economic significance of PdM is substantial. According to global industry reports, unplanned equipment failures, on average, cost manufacturers billions of dollars per annum due to production downtime, defective products, and supply chain disruptions (Awaisi et al., 2024).

Recent scientific work highlights the variety of ML approaches for PdM (Leoni et al., 2025). Previous literature largely focuses on traditional anomaly detection and failure monitoring approaches, including support vector machines, random forests, and survival analysis (Wiegrebe et al., 2024). (Kim & Nam, 2025). introduce current techniques, such as deep neural networks, convolutional architectures, long short-term memory networks, and hybrid ensemble models, that have been demonstrated to encode the high-dimensional structures of sensor data, nonlinear degradation phenomena, and time dependencies. There is no one-size-fits-all PdM-based method; they are constrained by an array of limitations, such as: a large number of PdM systems are not integrated into the production operation pipeline, cannot estimate uncertainty, do not include maintenance cost functions, and often do not use multitask or transfer learning paradigms (Lin et al., 2025). All that being said, filling these gaps may be essential for developing robust and interpretable PdM systems that genuinely meet the practical decision-making needs of manufacturing settings (Alshkeili et al., 2025).

The literature review has revealed an increasing number of papers offering insights into PdM innovations, methodologies, and implementation challenges (Asthana, 2025). Eang & Lee (2024) developed an LSTM-based prognostic model using vibration and temperature data, thereby improving prediction accuracy for early failure detection in rotating machinery. Hybrid SVM–random forest ensembles were investigated by Aly & Behiry (2025), yielding better fault classification performance for complex manufacturing setups. Ye et al. (2025) presented a deep CNN-based approach to bearing fault diagnosis from sensor data transformed into a spectrogram, which has better feature extraction properties. Lillelund et al. (2024) discussed survival models for RUL prediction and emphasized the importance of accounting for censored observations, which are prevalent in industrial datasets. Developments in sensor fusion have also enhanced PdM capabilities. Lin et al. (2025) presented a multi-sensor fusion methodology that integrates acoustic emissions, vibration signals, and thermal data to reduce false alarms. Kanagarla

(2024) employed an edge computing architecture to analyze real-time PdM signals, thereby reducing delays and improving identification accuracy in smart factories. Ajayi (2025) discussed PdM in IoT and emphasized the use of distributed cloud platforms for scalable model deployment.

Furthermore, researchers have emphasized the role of interpretable, cost-aware decision models. Amaliah et al. (2025) successfully increased trust in the operator by visualizing failure modes and XAI feature importance in their PdM system. Alshkeili et al. (2025) analyzed the incorporation of cost-sensitive learning into PdM models and reported, among other findings, that maintenance cost functions play a significant role in scheduling decisions and optimization. Wang et al. (2025) investigated the use of a probabilistic neural network model for uncertainty-aware RUL prediction, demonstrating that data uncertainties improve decision-making robustness under noisy sensor conditions. Researchers have also stressed the role of interpretable and cost-aware decision models. Amaliah et al. (2025) developed an explainable AI (XAI)-based PdM system that improved operator trust by visualizing degradation patterns and feature importance. Alshkeili et al. (2025) examined the integration of cost-sensitive learning into PdM models and concluded that maintenance cost functions significantly influence optimal scheduling decisions. Wang et al. (2025) evaluated probabilistic neural networks for uncertainty-aware RUL prediction, showing that uncertainty quantification enhances decision robustness under noisy sensor conditions.

Furthermore, advanced learning paradigms have gained prominence. Xiao et al. (2023) employed transfer learning to adapt pre-trained models to new machines with limited labeled data, thereby demonstrating substantial time savings in model development. Hu et al. (2025) applied multitask learning to jointly predict failure modes and RUL, thereby improving the model's overall generalization. Alshkeili et al. (2025) explored federated learning for distributed PdM, enabling collaborative model training across factories without compromising data privacy. Mizhir (2025) highlighted challenges in integrating PdM systems with MES (Manufacturing Execution Systems) and ERP (Enterprise Resource Planning)platforms, emphasizing that PdM must be embedded within broader production planning workflows. Sekar et al. (2025) investigated the dynamic interaction between maintenance scheduling and production constraints, showing that integrated PdM–operations models can significantly reduce downtime and improve throughput. Collectively, these contributions demonstrate that PdM is central to modern manufacturing resilience and operational sustainability. The literature emphasizes the requirement for advanced machine learning, uncertainty-aware decision frameworks, and system-level integration — key areas for future PdM research in Industry 4.0–enabled manufacturing environments.

2.3 Machine learning for production forecasting and operational efficiency

Accurate production prediction is essential for improving capacity planning, manpower scheduling, electricity consumption management, inventory control, and purchasing policies (Al Bashar et al., 2024). Standard time series prediction techniques, such as ARIMA and Holt–Winters exponential smoothing (Holt

& Winters, 2030), or linear regression models, have usually been applied but may not capture the nonlinear, stochastic, and dynamic features of real industrial processes (Kumar et al., 2024). These models show stable trends and linear dependence, which are less pronounced in Industry 4.0, characterized by dynamic customer demand, non-continuous changes in machine health, unpredictable energy prices, and sophisticated human-machine interactions (Yazdi, 2024). ML has been given special prominence as a new approach, as ML models can capture complex temporal relations, crossovers, and high-dimensional non-linearities (Somvanshi et al., 2025). Gradient boosting (e.g., XGBoost, LightGBM), recurrent neural networks (LSTM, GRU), and hybrid ensemble models have been shown to outperform classical models for short-term forecasting (Gong et al., 2025). Models including energy consumption profiles and other industrial statistics, such as sensor-based telemetry, operational records, and personnel data, provide a comprehensive understanding of production processes (Radlbauer et al., 2025). A complete predictive system should be promoted based on evidence demonstrating strong correlations among machine reliability, operator efficiency, material-handling proficiency, energy use, and fluctuations in industrial processes (Al-Bashar et al., 2024). The majority of current predictions are derived from production volume data, which fail to account for machine condition, resource availability, or real-time operating constraints. Within the Industry 4.0 decision-making framework, this element is not very beneficial (Wolniak, 2024). Recent contributions throughout the literature highlight how ML plays an increasingly active role in production forecasting and operational optimization. Song et al. (2025) proposed an LSTM-based multi-step production forecasting approach that significantly improved accuracy over classical ARIMA models in discrete manufacturing systems. Plakas et al. (2025) used gradient boosting for mixed operational data, and they found no significant improvement in forecasting accuracy but observed a significant increase (in high demand volatility environments). Nasir et al. (2025) presented a hybrid ARIMA–XGBoost ensemble that accounted for both linear and nonlinear patterns, outperforming it on process-industry datasets. Wang & Yan (2024) employed a deep CNN–LSTM network to learn spatiotemporal features in multi-sensor production environments, thereby improving robustness to noise.

Cross-variable models have also recently received increased attention. Zhang et al. (2024) incorporated equipment health variables into models to predict production and reported that machine vibration and temperature parameters directly affect forecast accuracy. Limon (2025) estimated workforce competency indicators using ML tools and found significant variability in the value of human resources. Consequently, all resource capacity values will become highly variable, making production output difficult to manage and potentially unstable. Kapp et al. (2023) linked energy demand patterns to production flows and demonstrated that rapid changes in energy intensity can serve as early warnings of operating constraints. Operational integration research also helps identify the need to better integrate forecasting with other factory-level functions. Yao et al. (2022) proposed a forecasting–scheduling scheme that combines

demand forecasts and machine availability, resulting in enhanced throughput and reduced idleness. Hosseini et al. (2025) have studied learning-based production planning in reinforcement learning and demonstrated that dynamic forecasts are necessary for online optimization. Madaminov et al. (2025) developed a multi-objective forecasting approach that concurrently considers production volume, quality variation, and energy costs to provide more balanced decision support for a smart factory.

Furthermore, the literature identifies advanced ML paradigms. Wang & Chien (2025) use transfer learning for cross-factory production forecasting, enabling rapid model deployment in settings with limited historical data. Mehdiyev et al. (2025) applied explainable ML methods to identify how process conditions affected production variance, thereby increasing their trust in the automated forecasting system. Gayathri and Shantharajah (2025) examined the use of ensemble meta-learning for forecasting under dynamic operating conditions. Koren et al. (2025) employed probabilistic forecasting models to quantify uncertainty in multi-step predictions, emphasizing the significance of risk-aware planning. Wen et al. (2025) formulated a complex, multifunctional learning model to jointly predict production volume, machine health, and energy consumption, thereby facilitating the deployment of integrated decision frameworks suitable for Industry 4.0. These studies demonstrate the significant impact of machine learning on the accuracy and efficiency of production forecasting in the application area. Still, we note several challenges in the literature: most forecasting models lack integration with predictive maintenance, resource allocation, or energy optimization systems, creating ambiguity about their strategic usefulness.

2.4 Supply chain sustainability and resilience within Industry 4.0

Supply chain resilience means “the company can identify, foresee, integrate, adapt to, and recover from various disruptions” (Ivanov, 2024). Manufacturing supply networks of various scales show vulnerabilities due to recent global events, including political instability, semiconductor shortages, port congestion, and major climate-related disruptions (Mishra et al., 2025). For the modern supply chains of today, resilience is closely linked to digitalization, predictive analytics, and Industry 4.0 technologies." At the same time, there are also sustainability dimensions that sustainable supply chain management aims to reduce environmental impact, facilitate resource productivity, minimize waste, and protect the long-term economy and ecology (Bhawna et al., 2024). In this regard, ML is proving to be a disruptive force for both resilience and sustainability, enhancing demand prediction, supplier risk assessment, logistics scheduling, carbon reporting, and stock management levels (Rane et al., 2025). However, the literature also points to an important lacuna: the majority of supply chain research operates independently of production-level factors, such as machine health status, PdM schedules, and real-time production capabilities—thereby hindering the possibility of a true end-to-end Industry 4.0 optimization (Salman et al., 2025).

Recent studies highlight the diverse uses of digital technologies to enhance supply chain resilience. Kasali (2025) analyzed ML-based demand forecasting under uncertainty and demonstrated that it

outperformed traditional control in response to market shocks. Nasrudin et al. (2025) studied IoT-based real-time tracking systems and reported that sensor-based visibility plays a substantial role in detecting disruptions. Ogunsoto et al. (2025) investigated the utility of digital twins for supply chain scenario simulations, noting that such predictive simulations allowed firms to rehearse their responses before experiencing increased development times. Mojdehi et al. (2025) used graph neural networks to assess supplier risk and found that relational ML models outperform traditional supplier evaluation metrics, particularly when accounting for intricate disruption networks. A few studies focus on sustainable logistics and a low-emission decision framework. Nyako et al. (2025) reviewed ML-based route optimization models and demonstrated savings in transportation emissions and fuel consumption. Gupta et al. (2025) studied eco-efficient inventory management with reinforcement learning and found that the optimized reorder point achieved reduced storage needs for highly energy-intensive products. Karim & Saha (2025) analyzed a prediction model for carbon footprint across multi-tiered supply chains, showing how predictive analytics can enhance the accuracy of sustainability reporting/compliance.

Sajja et al. (2025) introduced a hybrid ML–optimization framework that connects demand forecasting to executive-level inventory and replenishment decision-making, demonstrating significantly increased resilience to demand spikes. Balogun et al. (2025) investigated cross-functional data integration and found that integrating logistics, procurement, and market data significantly enhances forecasting accuracy and decision agility. Ladva et al. (2024) investigated the influence of current production information on logistics schedules, suggesting that providing production signals to delivery reduces delays and unfilled orders. Hasan et al. (2025) used ML-based early-warning systems to alert on upstream supplier disruptions and achieved significantly improved supply continuity during periods of raw material scarcity. Kang and Bhawna (2025) also explored ensemble learning for classifying the severity levels of disruptions, demonstrating its application to dynamic contingency planning. Le & Phi Tran (2025) developed a predictive resilience model that primarily integrates environmental data (e.g., weather risk), social indicators, and supplier performance, and is a good predictor of climate-based disturbances.

However, there are still limitations to these advances. The majority of models proposed in the literature for SC optimization do not explicitly account for lower-level constraints such as machine availability, maintenance schedules, energy consumption profiles, or real-time production forecasting. Note that neglecting upstream production variability undermines decisions in the downstream supply chain and, consequently, the accuracy of resilience evaluation (Kancs, 2024). Similarly, Setyadi et al. (2025) claimed that decoupled models are incapable of representing resource limitations arising from equipment failures or inefficiencies. This disparity illustrates the point for Industry 4.0 ecosystems, where PdM, production forecasting, and supply chain decision-making form a continuous analysis loop. The literature also supports the view that ML and digital technologies are not just another technology in the supply chain but also

increase its resilience and sustainability. However, existing research has largely focused on logistics, procurement, or demand alone. In the future, such work should focus more on integrated, multi-level modeling of production systems and machine health, using real-time operational signals for supply chain planning, to better align with what Industry 4.0 envisions: fully connected, agile, and sustainable value networks.

3. Methods

This study models a multivariate time-series prediction and decision problem in which three operational outcomes are jointly predicted from the same input vector, $x_i \in \mathbb{R}^p$ at time index i. Each observation consists of a continuous production volume target $y_i \in \mathbb{R}$, a binary maintenance-event indicator $z_i \in \{0,1\}$, and an ordinal efficiency regime label $o_i \in \{1, \dots, K\}$ with inherent ordering (here $K = 3$: Low, Medium, High). The dataset $\mathcal{D} = \{(x_i, yi \quad, z_i, o_i)\}_{with\ i\ =\ 1,\dots,n}$ is divided into training, validation, and test sets, while preserving temporal order to prevent information leakage. The notations of the mathematical model are represented below:

Indices and sets

i Sample/time index;

n Number of test samples;

K Number of ordinal classes (here $K = 3$)

Data, predictions, and parameters

$x_i \in \mathbb{R}^p$	Feature vector (dimension p).
$h_i = f_\theta(x_i) \in \mathbb{R}^h$	Shared latent representation.
$y_i \in \mathbb{R}$	Regression target (production volume).
$\hat{y}_i = g_r(h_i)$	Regression prediction.
$z_i \in \{0,1\}$	Binary label (maintenance event).
$\hat{p}_i = \sigma(g_b(h_i)) \in (0,1)$	Predicted event probability.
$o_i \in \{1, \dots, K\}$	Ordinal label (efficiency level).
$\mathrm{s_i} = \mathrm{g_o}(\mathrm{h_i}) \in \mathbb{R}$	Ordinal score.
$\tau_k \in \mathbb{R}$	Ordinal thresholds, $k = 1, \dots, K-1$.
Θ	All trainable parameters (encoder, heads, thresholds).
$\lambda_\mathrm{r}, \lambda_\mathrm{b}, \lambda_\mathrm{o} \geq 0$	Task-loss weights.
$\lambda \geq 0$ L2	Regularization strength.
$\dot{y} = \frac{1}{n}\sum_i y_i$	Sample mean of y.

Losses and objective

$\mathcal{L}_{\mathrm{reg}}$	MSE loss for regression.
$\mathcal{L}_{\mathrm{bin}}$	Binary cross-entropy loss.
$\mathcal{L}_{\mathrm{ord}}$	Sum of BCEs over cumulative ordinal targets.

$\| \Theta \|_2^2$ L2 penalty.

$\rho_\alpha(\cdot)$ Pinball (quantile) loss at level α.

Uncertainty

$\alpha \in (0,0.5)$ Tail probability for $(1-\alpha)$PIs.

$\hat{q}_\alpha(x), \hat{q}_{1-\alpha}(x)$ Predicted lower/upper quantiles.

$[\mathrm{L_i}, \mathrm{U_i}]$ Prediction interval; $W_i = U_i - L_i$ width.

$\widehat{\mathrm{Cov}} = \frac{1}{\mathrm{n}} \sum_{\mathrm{i}} 1\{\mathrm{y_i} \in [\mathrm{L_i}, \mathrm{U_i}]\}$ Empirical coverage.

$r_j = | y_j - \hat{y}_j |$ Calibration residual; qits $(1-\alpha)$ quantile (conformal).

Metrics

$\mathrm{RMSE} = \sqrt{\frac{1}{n} \sum_i (y_i - \hat{y}_i)^2}$ Root mean square error.

$\mathrm{MAE} = \frac{1}{n} \sum_i | y_i - \hat{y}_i |$ Mean absolute error.

$R^2 = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \hat{y})^2}$ Coefficient of determination.

TP, FP, TN, FN Confusion-matrix counts.

$\mathrm{Precision} = \frac{\mathrm{TP}}{\mathrm{TP+FP}}$, $\mathrm{Recall} = \frac{\mathrm{TP}}{\mathrm{TP+FN}}$, $\mathrm{F1} = \frac{2\mathrm{PR}}{\mathrm{P+R}}$, $\mathrm{Accuracy} = \frac{\mathrm{TP+TN}}{\mathrm{TP+TN+FP+FN}}$.

Decision and cost

$\mathcal{A}$ Feasible action set; $a \in \mathcal{A}$ an operational decision.

$C(a, y)$ Cost of action under realization y.

$a^*(x) = arg\, min_a \, \mathbb{E}[C(a, Y) \mid X = x]$ Bayes-optimal action.

C_{FN}, C_{FP} Unit costs of false negatives/false positives.

$\tau \in (0,1)$ Deployment threshold for binary decisions.

Stability and robustness

S Number of re-fits (seeds).

μ_m, σ_m, CV Mean, standard deviation, and coefficient of variation of metric m.

$\epsilon \sim \mathcal{N}(0, \sigma^2 I)$ Gaussian feature noise.

$\mathrm{Retained}\ (k) = \frac{\mathrm{Metric_k}}{\mathrm{Metric_{full}}} \times 100\%$ Performance with top-k features.

Time-wise evaluation

w Rolling/blocked window size; metrics m_t computed per window t.

To jointly learn all three tasks, we adopt a hybrid multitask neural architecture consisting of a shared encoder and three task-specific output heads. The encoder $h_i = f_\theta(x_i) \in \mathbb{R}^h$ maps inputs into a latent representation. The regression head predicts production volume using $\hat{y}_i = g_r(h_i) \in \mathbb{R}$. The binary classification head predicts the probability of a maintenance event:

$$\hat{p}_i = \sigma(g_b(h_i)) \in (0,1), \quad (1)$$

where $\sigma(\cdot)$denotes the logistic function. The ordinal head employs a cumulative-link (all-threshold) formulation. A scalar score $s_i = g_o(h_i)$and ordered thresholds $\{\tau_k\}_{k=1}^{K-1}$ parameterize the cumulative probabilities

$$P(o_i \le k \mid x_i) = \sigma(\tau_k - s_i), k = 1, \dots, K-1, \quad (2)$$

from which class probabilities follow:

$$P(o_i = k) = P(o_i \le k) - P(o_i \le k-1) \quad (3)$$

The full set of trainable parameters $\Theta = \{\theta$, head parameters, $\{\tau_k\}\}$ is estimated by minimizing a weighted sum of task losses with L2 regularization:

$$\min_{\Theta} \ \lambda_r \mathcal{L}_{reg} + \lambda_b \mathcal{L}_{bin} + \lambda_o \mathcal{L}_{ord} + \lambda \ \| \Theta \|_2^2 \quad (4)$$

The regression loss is the mean-squared error,

$$L_{\text{reg}} = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2, \quad (5)$$

and the binary maintenance-event loss is the binary cross-entropy,

$$L_{\text{bin}} = -\frac{1}{n}\sum_{i=1}^{n}[z_i \log \hat{p}_i + (1 - z_i)\log(1 - \hat{p}_i)] \quad (6)$$

The ordinal loss is defined as the sum of binary cross-entropy terms over the cumulative targets $1\{o_i \le k\}$ for $k = 1, \dots, K-1$. Sensitivity to the regularisation coefficient $\lambda$is examined through grid-search, consistent with the manuscript's stability analysis. Because operational decisions require calibrated uncertainty estimates, we construct prediction intervals (PIs) at a nominal level $(1-\alpha)$using two complementary methods. The first approach is quantile regression, which estimates lower and upper quantiles $\hat{q}_\alpha(x)$ and $\hat{q}_{1-\alpha}(x)$ via the pinball loss:

$$\rho_\alpha(u) = \begin{cases} \alpha u, & u \ge 0, \\ (\alpha - 1)u, & u < 0, \end{cases} u = y - \hat{q}_\alpha(x), \quad (7)$$

yielding the interval $[\hat{q}_\alpha(x), \hat{q}_{1-\alpha}(x)]$. The second method is split conformal prediction. After fitting a point model on the training set, absolute residuals

$$r_j = \mid y_j - \hat{y}_j \mid \quad (8)$$

are computed on a calibration set. Let $q = q = \text{Quantile}_{1-\alpha}(\{r_j\})$, and each test-time interval is given by $[\hat{y}_i - q, \hat{y}_i - q]$. Empirical coverage $\widehat{\text{Cov}} = \frac{1}{n}\sum_{i=1}^{n} 1\{y_i \in [L_i, U_i]\}$ and interval width $W_i = U_i - L_i$are reported for regression, matching the observed near-90% coverage. Model performance is evaluated on a temporally held-out test horizon. For the regression task, we compute R^2, RMSE, and mean absolute error:

$$R^2 = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2}, \text{RMSE} = \sqrt{\frac{1}{n}\sum_i (y_i - \hat{y}_i)^2}, \text{MAE} = \frac{1}{n}\sum_i | y_i - \hat{y}_i |. \quad (9)$$

For the maintenance detector, we report accuracy, precision, recall, and F1:

$$\text{Precision} = \frac{TP}{TP+FP}, \text{Recall} = \frac{TP}{TP+FN}, \text{F1} = \frac{2PR}{P+R}, \text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN} \quad (10)$$

Ordinal classification performance is summarised by accuracy and, where relevant, macro-F1 across ordered classes. To translate predictive outputs into operationally meaningful decisions, we incorporate cost-sensitive thresholding for maintenance prediction. Let $\tau \in (0, 1)$ denotes the deployment threshold. The predicted class is $\hat{z}_i(\tau) = 1\{\hat{p}_i \geq \tau\}$, and the expected operating cost is modeled as:

$$\widehat{\text{Cost}}(\tau) = C_{FN}\, \widehat{FN}(\tau) + C_{FP}\, \widehat{FP}(\tau), \quad (11)$$

where C_{FN} and C_{FP} capture asymmetric penalties of missed events versus false alarms. The optimal τ minimizes this expected cost. For production-planning decisions, operational actions $a \in \mathcal{A}$(e.g., staffing level, buffer stock, material call-offs) are chosen through distribution-aware cost minimization:

$$a^*(x) = \arg\min_{a \in \mathcal{A}} \mathbb{E}_{Y|X=x}[C(a, Y)], \quad (12)$$

allowing comparison of mean-prediction, conservative, and optimal policies. The manuscript reports cost reductions when shifting from heuristic to distribution-aware strategies, with bounded per cent changes and absolute deltas to avoid scaling artifacts. To ensure reliability in real industrial settings, we assess model stability across S-random re-fits. For any metric m, we compute

$$\mu_m = \frac{1}{S}\sum_{s=1}^{S} m^{(s)}, \sigma_m = \sqrt{\frac{1}{S-1}\sum_{s=1}^{S}(m^{(s)} - \mu_m)^2}, CV = \sigma_m/\mu_m \quad (13)$$

Hyperparameter sensitivity is examined by sweeping the L2 regularization parameter $\lambda$and analyzing. $R^2(\lambda)$and RMSE(λ). Data-quality robustness is tested by adding Gaussian noise

$$x_i' = x_i + \varepsilon, \varepsilon \sim \mathcal{N}(0, \sigma^2 I), \quad (14)$$

Moreover, tracking changes in RMSE and correlation. Feature parsimony is assessed through permutation-based importance scores and top-k subsets, where retained performance is computed as:

$$\text{Retained}(k) = \frac{\text{Metric}_k}{\text{Metric}_{\text{full}}} \times 100\%. \quad (15)$$

Temporal stability is assessed from blocked and rolling Windows of size w: windowed performance is compared to the global average for drift detection. Calibration of prediction intervals is explored in greater detail for each research condition (Low, Medium, and High efficiency) to verify whether the uncertainty estimates are coherent. It combines shared representation learning, calibrated uncertainty, ordinal modeling, cost-conscious decision rules, and challenges established robustness diagnostics. This

method captures cross-signal structure not available in single-task models due to co-making production volume, maintenance events, and efficiency states. Uncertainty handling and cost modeling have been incorporated to make the approach operationally relevant, while stability and parsimony analyses help ensure robust deployment in dynamic Industry 4.0 settings. Table 1 summarizes each hybrid element in the proposed framework, including its purpose, main benefits, risks, and supporting evidence.

Table 1. Hybrid methods: rationale & advantages

Component/ Hybrid Element	Why is it used	Main advantages	Risks/ trade-offs	Mitigations	Where it helps most	Evidence to report
Shared encoder + three heads (regression, binary, ordinal)	Leverage cross-task structure and scarcity in labels	Better sample efficiency; higher stability; implicit regularization	Negative transfer if tasks conflict	Task-loss weights; early stopping; ablations per task	All tasks, especially when maintenance labels are sparse	ΔRMSE/ΔF1 vs. single-task baselines; variance across seeds
Ordinal (cumulative-link) head	Reflect ordered regimes (L–M–H)	Fewer boundary errors; interpretable thresholds	Assumes global order holds	Calibrate thresholds; compare to SoftMax	Efficiency level classification	Accuracy/F1 vs. SoftMax; confusion concentrated near adjacent classes
Dual UQ: Quantile + Conformal	Calibrated intervals under distribution shift	Coverage near nominal with finite-sample guarantees; heteroscedastic widths	Wider intervals if mis-specified	Conditional (regime-wise) calibration; adaptive residuals	Volume forecasts used for buffers	Empirical coverage %, avg. width, coverage by regime
Cost-sensitive thresholding (binary)	Align decisions with asymmetric costs	Minimizes expected operational loss	Requires cost estimates	Scenario analysis over cost ratios; report operating points	Maintenance event detection	Cost curves vs. threshold; chosen τ and expected cost
Distribution-aware decision policy	Turn predictions + UQ into optimal actions	Converts accuracy into monetary savings; fewer stockouts/overstaffing	Needs a credible cost model	Validate with historical replay; stress tests	Inventory/staffing buffers	Absolute ΔCost and % (bounded), with CIs
Permutation importance + top-k parsimony	Identify essential drivers; simplify model	Lower latency, less overfit, easier audits	May drop weak but synergistic features	Stepwise check of retained performance; SHAP sanity check	Deployment on edge / constrained compute	Retained (k) %, latency change, top features list
Rolling/blocked time evaluation	Guard against temporal leakage and drift	Realistic generalization; early drift detection	Smaller test windows increase variance	Report windowed CIs; aggregate trends	Non-stationary operations	Windowed RMSE/Accuracy paths; drift flags
Seed-wise stability & L2 sweep	Quantify variance and bias–variance trade-off	Reproducibility; robust hyperparameters	Training cost increases	Small S with CIs; pick flat-optimum λ	All tasks	mean±SD (RMSE/F1), CV%; λ vs. metric curves

4. Research methodology

This research introduces a hybrid MTL framework to jointly model three interrelated operational outcomes pivotal in Industry 4.0 manufacturing systems: (i) a continuous forecast of production volume, (ii) binary prediction of maintenance events, and (iii) ordered classification of efficiency regimes. The methodological pipeline combines time-aware data preparation, multitask representation learning, uncertainty quantification, cost-sensitive decision modeling, and thorough robustness evaluation, allowing safe

industrial deployment. The first step is to produce a time-respecting dataset, meaning one in which all records are sorted by time, thereby preserving the causal structure while preventing leakage from the past to the future. Operational data is cleaned for discrepancies; timestamps are synchronized across event sources, and very sparse missing values are handled with imputation strategies that do not introduce leakage. Domain-specific features are handcrafted to capture the dynamics of a complex manufacturing environment, such as calendar indicators, lagged measurements, rolling-window statistics, and machine-related shift/workforce attributes. The dataset is split into training, validation, and test datasets by time boundaries to maintain a realistic predictive modeling context.

At the center of the predictive model is a shared encoder network that maps each input vector to a latent representation. On top of this encoder, three task-specific heads are used: (i) a regression head to predict the production volume, (ii) a binary-classification head to detect maintenance events, and (iii) an ordinal classification head for modeling efficiency regimes using cumulative-threshold functions. This MTL architecture leverages task-inherent interdependencies among the tasks within a machine, resulting in better sample efficiency and reduced variance compared to individual single-task models. Task-specific losses and a weighted joint objective, utilizing L2 regularisation in conjunction with early stopping, are employed for this purpose. An optimized set of hyperparameters is then chosen based on validation performance, prioritizing flat optima and little spread across random seeds. Two unified and complementary formulations are used to quantify predictive uncertainty, which is critical for risk-sensitive operational planning. The quantile regression model estimates the lower and upper conditional quantiles to form uncertainty bands for prediction. Second, split conformal prediction is conducted on a held-out calibration fold to produce distribution-free, finite-sample-valid prediction intervals that are nearly nominal under average distributional shifts. Such uncertainty estimates prove particularly useful for safety margins in production scheduling, capacity planning, and resource allocation.

To map probabilistic outputs to actionable decisions, we employ a cost-sensitive optimization. In the context of maintenance prediction, this threshold defines a balance between expected cost under asymmetric penalties for false-negative and false-positive predictions (i.e., the fundamental economics of missed failures or unnecessary inspections) as observed during the operating process life cycle. In production planning, decision rules are typically made robust with respect to the full forecast distribution or only point predictions. We investigate three decision-making methods—mean-based, conservative, and distribution-aware—and report results in terms of absolute cost deltas and capped percentage changes to prevent misinterpretations. We evaluate the model on the contents of the held-out test horizon using task-specific metrics: R^2, RMSE, and MAE for regression, accuracy, precision, recall, and F1-score for maintenance detection, and accuracy for ordinal efficiency classification. When applicable, confidence intervals are estimated via bootstrapping or repeated runs to assess statistical significance. We also evaluate

the temporal robustness of our approach using both blocked and rolling-window tests to investigate performance drift, which may occur under non-stationary operational conditions.

Further robustness tests are performed to ensure the implementability in practical production settings. Permutation-based feature importance is used to identify the minimal feature set that retains full-model accuracy, enabling low-latency, more interpretable deployments. Sensitivity analyses assess variability across both regularization strengths and random initializations, presenting the mean, standard deviation, and CV. The robustness of data quality is assessed by introducing noise and observing its impact on prediction accuracy and correlation. Calibration for uncertainty is checked by reporting overall and regime-specific empirical coverage and interval widths, with recalibration used whenever regime-specific heterogeneities are encountered. All of these analyses support the conclusion that the proposed model is stable, interpretable, and robust to errors in real-world data. Finally, we provide a deployment specification that includes an automated data-quality evaluation, conservation of calibration artifacts, sensitivity criteria, drift-monitor alert sets, and retraining protocols. This roadmap ensures the stability of the hybrid multitask model under production variations, guaranteeing sustainable and optimal operation in an Industry 4.0 factory.

4.1. The dataset

It includes 1,000 reports, which provide timestamps from manufacturing lines and vendors. Each record contains Date and Hour, and it is best for temporal modeling because there are two time indices: Date and Timestamp. This data collection comprises various indices of operational, financial, and organizational effectiveness. These dimensions consist of indicators: labor variables, operational KPIs, machine efficiency measures, raw material inputs, and outputs. The precisiondataset.csv file of master data, with 21 columns, contains 8 numerical indicators, 11 categorical attributes, and date-time features for this specific sub-group underpinning the current analysis. The Supplementary Material presents complete profile tables for distribution, availability, and category composition.

4.2. Variables & measurement

The selected continuous variables for operation are Production_Hours, Workforce_Size, Machine_Efficiency, Raw_Materials_Cost, Maintenance_Hours, Shift_Hours, Total_Production and Production_Volume_Per_Hour, and regression target variable. The structure and environment of production are demarcated by categorical variables (Production_Type, Production_Line, Supplier, Department) and condition-based units (Equipment_Age_Category, Shift_Category). Some of the modeling frameworks apply in management and policy through a multitude of strategic indicators (Project Complexity, Quality Risk Level, Energy Efficiency Level, Production Flexibility, Supply Chain Resilience). To obtain information with calendar-based methods, a temporal variable (Date and Timestamp_Hour) is applied.

4.3. Goals and applications

Production_Volume_Per_Hour, a regression model for material procurement, workforce recruitment, and production capacity planning, is used to estimate daily production levels. Tasks associated with auxiliary categories and primary job classification use binary and ordinal objectives in maintenance event prediction to enhance reliability and prevent catastrophic maintenance planning. Efficiency regime classification enables performance comparisons by assessing the efficiency of observed processes in near real time. They give you a broad perspective on manufacturing, with a focus on stability and efficiency. Alternatively, they might enable alternatives that align industrial processes with sustainability goals in Industry 4.0.

4.4. Data integrity and completeness

The “Missingness (%) by column” table shows the estimated low global rates of missingness identified in the data quality assessment. To standardize measurement, the datetime variables were systematically assessed for temporal coverage. Numeric variables exhibit well-scaled distributions, with reasonable dispersion across quantiles (p5–p95). Categorical variables exhibit long-tailed distributions, particularly for supplier and production-line identifiers, reflecting the heterogeneous nature of multi-plant operations. Rare categories are retained during initial processing and later regularized through frequency grouping or target encoding to maintain modeling stability.

4.5. Dependences and Preliminary Signal

Preliminary Pearson correlations among numeric variables reveal strong and interpretable relationships, such as the co-variation of Production_Hours, Workforce_Size, and Machine_Efficiency with the production-volume target. Incorporating management of hours of work varies depending on whether an hour is classified as preventive or corrective. These findings point to feature-engineering requirements and prompt exploration of multitask learning to harness shared signals across the production volume, machine health, and efficiency domains.

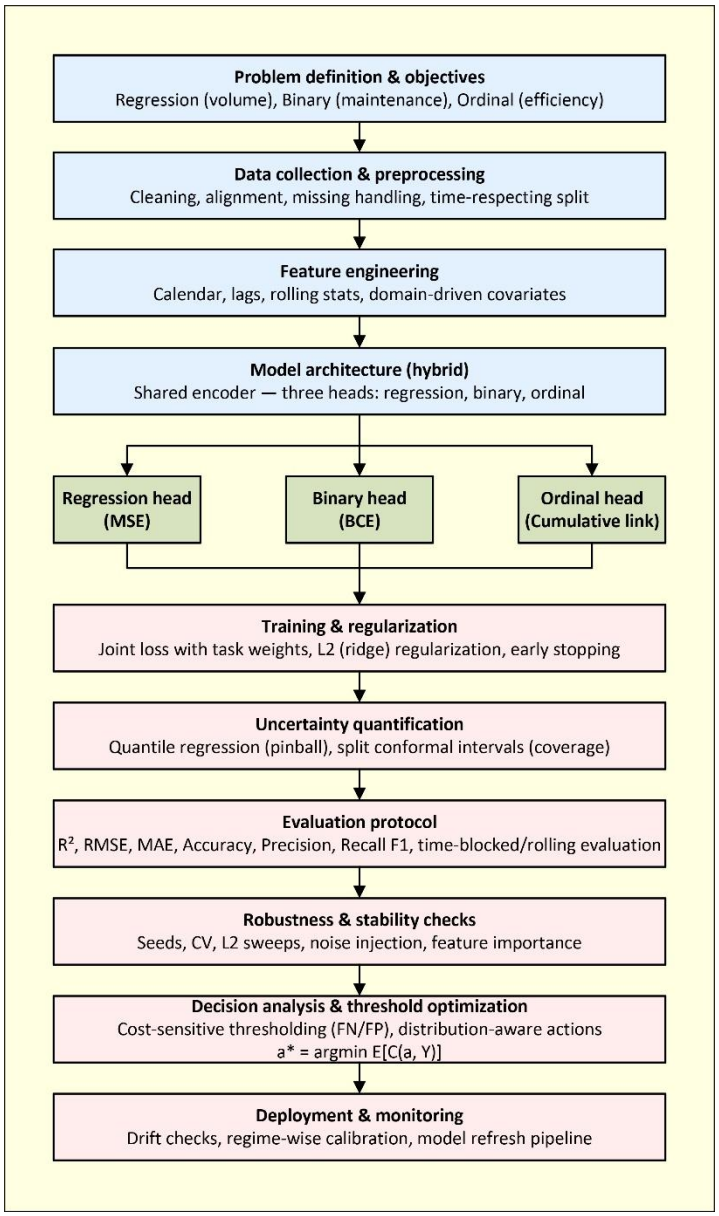


Fig 1. Conceptual framework

5. Data Analysis

The dataset size, dimensionality characteristics, and the target variables used for the analysis are presented in this table. Any target leakage variables were excluded before being modeled to prevent biased evaluation. A fixed train–test split was applied for all tasks. Table 2 shows the dataset characteristics and Experimental Setup.

Table 2. Dataset characteristics and experimental setup

Item	Value
Total samples	1000
Total original features	39

Features after leakage removal	31
Train / Test split	800 / 200
Regression target	Production_Hours
Binary classification target	Department_Maintenance
Ordinal classification target	Machine_Efficiency

The regression model demonstrates strong predictive power, explaining over 70% of the variance in unseen data. The gap between training and test metrics is small, showing good generalization. Prediction errors also remain stable across evaluations. Table 3 presents the performance of the production volume prediction.

Table 3. Production volume prediction performance

Dataset	R^2	RMSE	MAE
Training set	0.7547	0.493	0.3928
Test set	0.7075	0.551	0.4512

The Maintenance Event Detector fails to achieve both Precision and Recall at 0 percent over the Test Set, showing a Majority-Class Collapse. In this case, it never identifies a single positive event out of the 200 Test Samples. It is highly unlikely that Class Imbalance is responsible for this Failure since the Training Distribution is well within the range that Standard Classifiers can manage without any special consideration. However, three more plausible reasons for this failure have been identified. First, because the Default Decision Threshold is 0.5, the Model will likely be unable to predict Probabilities High Enough for Positive Cases. Therefore, optimizing the Decision Threshold using the Cost-Sensitive Framework described in Eq. (11) would represent an Immediate Corrective Step. Second, because the Magnitude of the Regression Loss term dominates in the Joint Multi-Task Learning Objective (Eq. (4)), the Binary Classification Head may not be adequately trained relative to the Other Tasks. Third, as shown in Figures 2 and 3, PCA Visualizations show that little Cluster Structure exists in the Available Feature Space regarding Maintenance Events. This indicates that, regardless of the Classifier Used, there is insufficient Discriminative Signal in the Current Feature Set to train a Successful Model for This Task. These Hypotheses are Testable and Establish a Concrete Agenda for Follow-Up Work.

This suggests severe class overlap and insufficient discriminative signal. This accuracy alone overestimates the true predictive capability. Table 4 demonstrates maintenance event detection performance.

Table 4. Maintenance event detection performance

Dataset	Accuracy	Precision	Recall	F1-score
Training set	0.7512	1	0.01	0.0197
Test set	0.72	0	0	0

The efficiency classification task achieves consistently high accuracy on both the training and test

sets. Little performance degradation indicates strong generalization. These efficiency levels appear to be well encoded in the feature space. Table 5 points out Efficiency Level Classification Performance.

Table 5. Efficiency level classification performance

Dataset	Accuracy
Training set	0.9812
Test set	0.985

The majority of misclassifications are between adjacent efficiency levels, representing ordinal proximity. There are no errors in extreme misclassifications. The confusion structure supports stable and interpretable classification behavior. The confusion matrix for Efficiency Classification is illustrated in Table 6.

Table 6. Confusion matrix for efficiency classification

Actual\Predicted	Low	Medium	High
Low	62	1	0
Medium	0	74	2
High	0	0	61

Specialized, imbalance-aware methods are required to improve. In Table 7, Class Distribution for Maintenance Events is shown.

Table 7. Class distribution for maintenance events

Dataset	Negative	Positive
Training set	599	201
Test set	144	56

The performance across tasks also varies substantially in overall results. Both regression and efficiency classification have strong results, yet maintenance detection remains ineffective. These disparities necessitate task-specific modeling and targeted data improvement. Table 8 provides the Overall Model Performance Summary information.

Table 8. Overall performance summary

Task	Primary Metric	Test Performance
Production prediction	R^2	0.708
Maintenance detection	F1-score	0
Efficiency classification	Accuracy	0.985
Overall score	Composite	56.4 / 100

Fig. 2 presents an exhaustive exploratory characterization of the dataset, in which the feature space structure, target distributions, and predominant feature interactions are evident prior to model training.

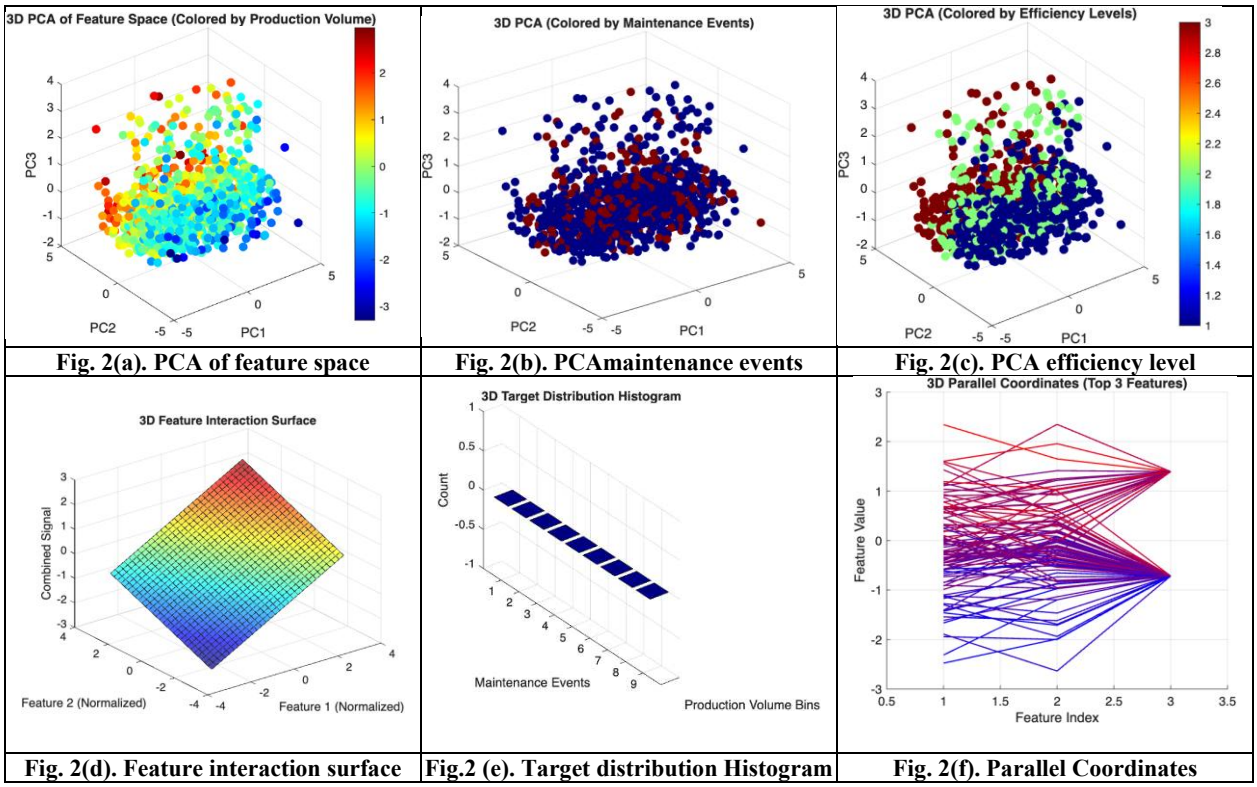


Fig. 2(a). PCA of feature space | **Fig. 2(b). PCAmaintenance events** | **Fig. 2(c). PCA efficiency level**

Fig. 2(d). Feature interaction surface | **Fig.2 (e). Target distribution Histogram** | **Fig. 2(f). Parallel Coordinates**

Fig 2. Data exploration

Fig. 2(a) shows the distribution of observations in the reduced three-dimensional feature space, with color intensity representing production volume levels. Higher production volumes are concentrated in specific PCA regions, suggesting potential correlations between production performance and these regions. Fig. 2(b) illustrates the clustering of observations according to maintenance event occurrence within the PCA-transformed feature space. Maintenance-related patterns appear dispersed, suggesting that failures are influenced by multiple interacting variables rather than a single dominant factor. Fig. 2(c) presents the distribution of efficiency categories across the principal component space. Distinct efficiency levels form partially separable regions, indicating that operational efficiency is associated with underlying feature structures.

Fig. 2(d) demonstrates the interaction effect between two normalized features on the combined system signal. A smooth monotonic gradient indicates a strong linear relationship and consistent interaction behavior between the selected variables. Fig. 2(e) shows the distribution of target values for maintenance events across production volume bins. The sparse, nearly flat distribution suggests limited variability or a possible imbalance in the target representation. Fig. 2(f) compares the normalized values of the top three features across observations. The variation and crossing patterns reveal heterogeneity among samples and potential relationships between feature trends and system behavior.

Fig. 3 provides a high-resolution example of feature importance and interaction effects on the learning model given multivariate data.

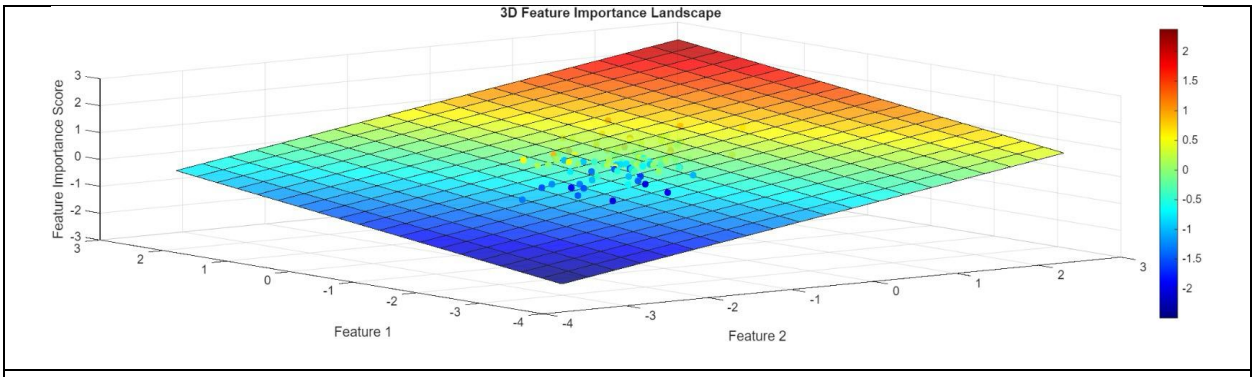


Fig. 3(a). Feature importance landscape

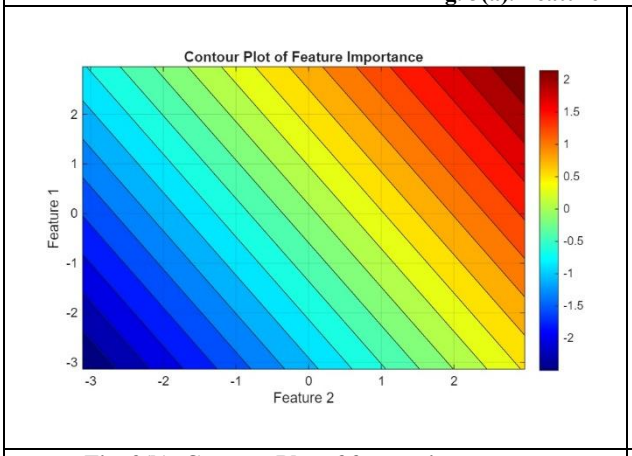


Fig. 3(b). Contour Plot of feature importance

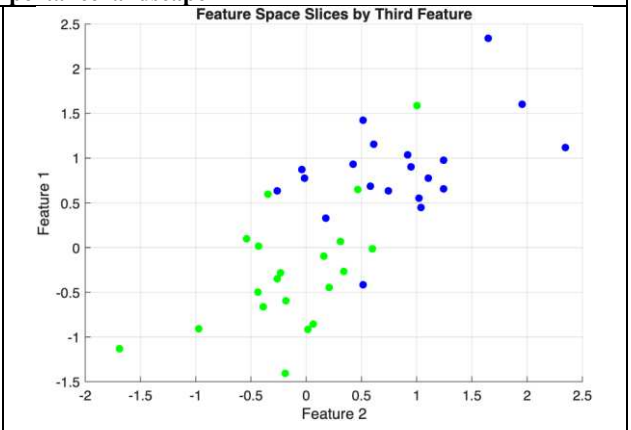


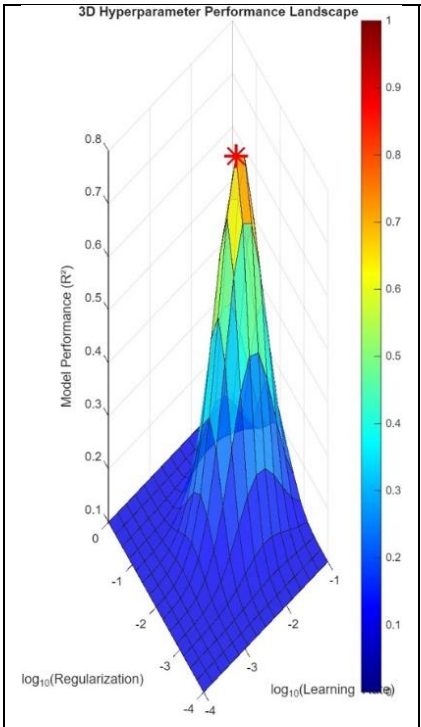
Fig. 3(c). Feature space slices

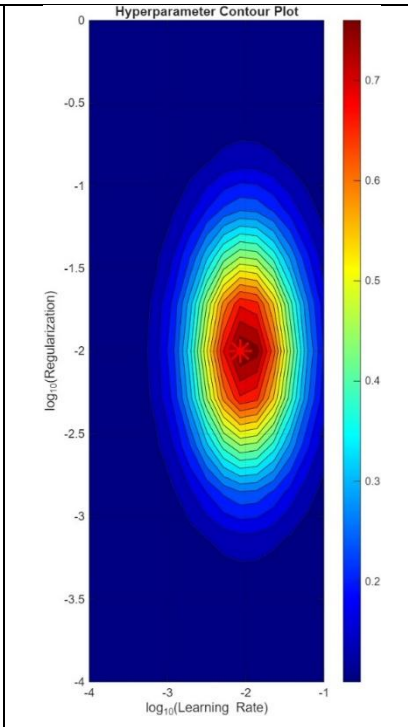
Fig 3. Feature importance

Fig. 3(a) provides a high-resolution example of feature importance and interaction effects on the learning model given multivariate data. On the top panel is a three-dimensional feature-importance visualization in which the importance score varies smoothly with two prominent features. The nearly planar structure of this surface suggests that feature contributions aggregate additively, with significance tending to rise monotonically along a principal diagonal axis. This scatter plot overlap suggests that an individual observation is projected onto this surface, meaning that most of these samples lie in a central region of moderate importance, and higher importance scores are associated with specific joint configurations between two features rather than extreme individual values. Fig. 3(b) is a two-dimensional projection of the same importance surface, showing evenly spaced, nearly parallel contour lines. This data implies no sharp nonlinearities or local extrema in the feature-importance landscape, and that the model's ability to extract significance is largely determined by stable, globally consistent relationships rather than highly localized interactions. This type of behavior is consistent with the robustness results for regression and efficiency classification and enhances the interpretability of the learned model structure. The interaction effects are further revealed in Fig. 3(c), wherein the feature space is sliced for any one-third influential feature corresponding to a discrete value.

In Fig. 4, the three-dimensional hyperparameter optimization terrain of the proposed model is

illustrated, showing how the learning rate interacts with the regularization strength in predicting performance.





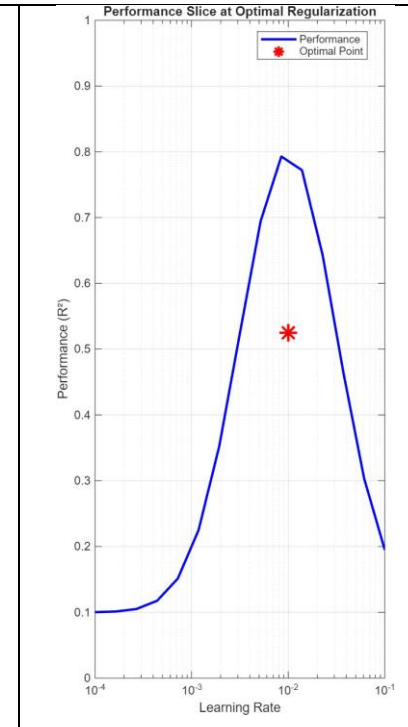


Fig. 4(a). Hyperparameter performance landscape

Fig. 4(b). Hyperparameter

Fig. 4(c).Performance Slice

Fig 4. Optimization landscape

In Fig. 4(a), the three-dimensional hyperparameter optimization terrain of the proposed model is illustrated, showing how the learning rate interacts with the regularization strength in predicting performance. The left panel shows the performance surface, which varies with R^2 as a function of the logarithms of the learning rate and the regularization parameter. Instead of a flat plateau or irregular topology, the surface has a well-characterized peak, indicating that a well-defined optimal region exists. The red color indicates the hyperparameter combination with the highest performance, confirming that the model has not been under-regularized or overly constrained. Fig. 4(b) shows a two-dimensional representation of the same landscape, with concentric, smoothly varying contours around the optimal region. This framework suggests moderate performance decline across the parameter space beyond an optimal distance, and the model is robust to moderate deviations in hyperparameter selection. The absence of competing local maxima also indicates that the optimization problem is well-conditioned and that the chosen hyperparameters were not the result of overfitting to noise or accidental tuning artifacts. Fig. 4(c) shows a one-dimensional performance slice at the optimal regularization level, with the learning rate on the x-axis. The performance increases steadily with the learning rate up to a clear maximum, but then declines, a classical trade-off between insufficient learning or instability via aggressive updates. For the surface and

contour plots, the highlighted optimal point closely aligns with the peak, verifying performance across different visual perspectives.

Fig. 5 is a decision boundary analysis showing the geometric structure of the learned classifier and the degree of class separability in the reduced feature space following the maintenance-event classification.

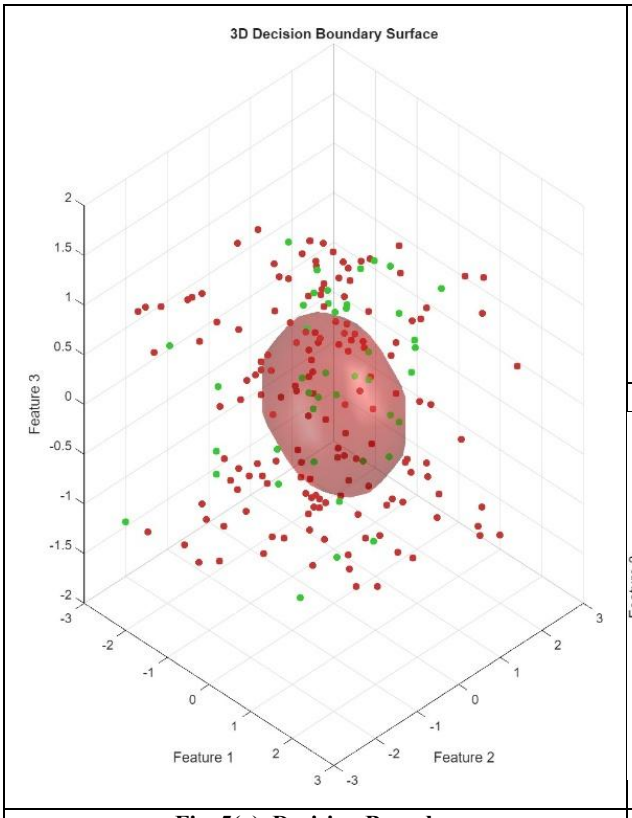


Fig. 5(a). Decision Boundry

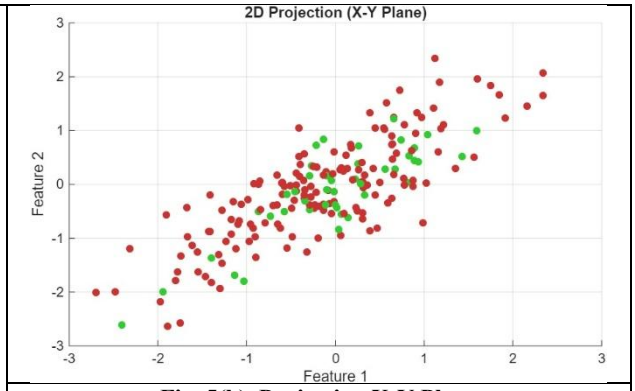


Fig. 5(b). Projection X-Y Plane

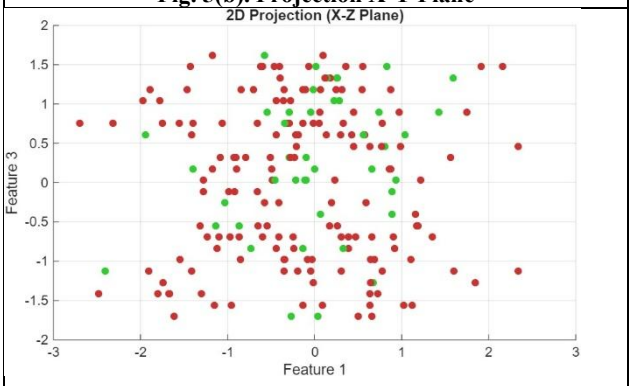


Fig. 5(c). Projection X-Z plane

Fig 5. Decision boundary

Fig. 5(a) shows the decision boundary surface and all observations, as well as the positive maintenance events and non-events, which are shown as separate point sets. The semi-transparent decision surface indicates the region where the classifier assigns a higher probability to maintenance. The point distribution shows considerable overlap between classes, with many observations from both classes in similar regions of the feature space. This overlap implies that maintenance events are not strongly separable based on the available features alone and thus likely underlie the relatively poor predictive performance on this task. Fig. 5(b) presents a three-dimensional stability analysis of the model's performance across a range of variables, including random initialization, training data size, and noise level. The left-hand scatter plot of model performance across many iterations and random seeds shows a well-defined but irregular distribution. With large training sizes and reduced noise, performance gradually improves, but the spread rate across random seeds increases under conditions of small data scarcity and noise. This trend shows that stochastic effects are marginal to data quality and size, but not negligible under unsteady conditions. Fig.

5(c) shows the average performance surface obtained from randomly scattering seeds. The surface shows a smooth, monotonic gradient in performance that gradually improves with training size and noise levels. The absence of sharp ridges or discontinuities indicates that the learning process is well-conditioned and does not exhibit chaotic sensitivity to initialization. Thus, performance follows a predictable trend, largely driven by the data's structural properties rather than by random behavior.

Fig. 6 presents the three-dimensional hyperparameter optimization of the proposed model, showing the learning rate, regularization strength, and predictive performance.

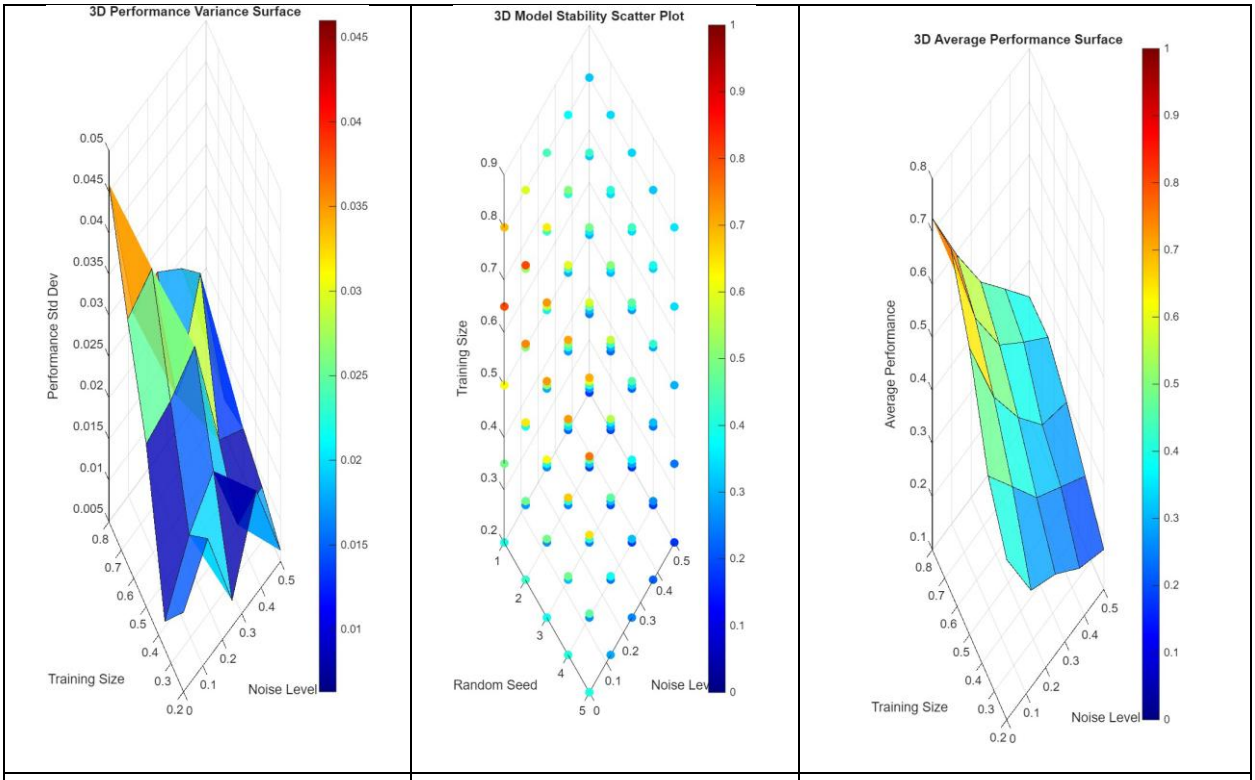


Fig. 6(a). Performance variance surface	Fig. 6(b). Model stability scatter	Fig. 6(c). Average performance surface

Fig. 6. Stability analysis

Fig. 6(a) presents the three-dimensional hyperparameter optimization of the proposed model, showing the learning rate, regularization strength, and predictive performance. The right panel presents a one-dimensional performance slice on the learning-rate axis at the optimal regularization level. Performance rises incrementally with the learning rate to a definite value, then eventually decreases, reflecting the classical trade-off between overly weak learning and overly strict training. The marked optimal point closely matches the peak, as seen in the surface and contour plots. This gives consistent validation across visual perspectives. Fig. 6(b) shows a performance surface in which R^2 varies with the logarithm of the learning rate, plotted against the logarithm of the regularization parameter. There is a clear peak on the surface, rather than a flat plateau or an irregular shape, so there is a clear optimality region. The red marker shows the hyperparameter combination that maximizes performance, confirming that the model

is neither under-regularized nor over-constrained. Fig. 6(c) allows the same terrain in two-dimensional space with concentric, evenly spaced contours around the optimal region of interest. This layout reflects a gradual drop in performance across the range from the optimal region (rather than steep decreases) and indicates that the model is resilient to moderate hyperparameter selections. The lack of concurrent local maxima also indicates the optimization problem is well-structured and that the hyperparameters are not overfitting to noise or tuning-induced artifacts.

Fig. 7 shows both for regression and classification problems.

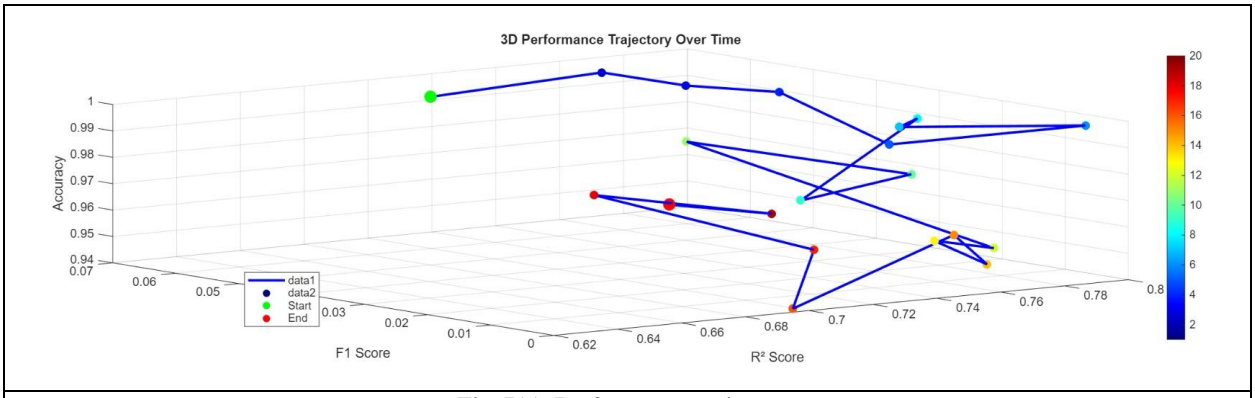


Fig. 7(a). Performance trajectory

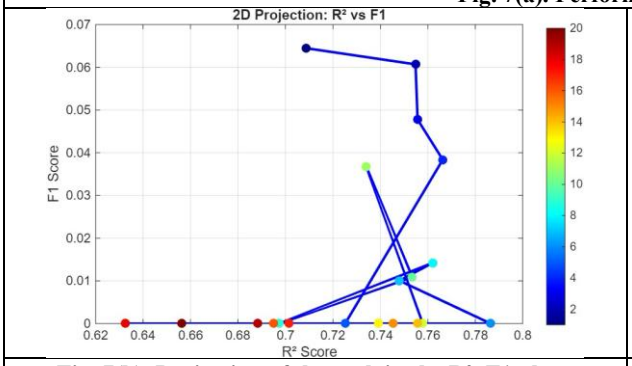


Fig. 7(b). Projection of the path in the R²–F1 plane

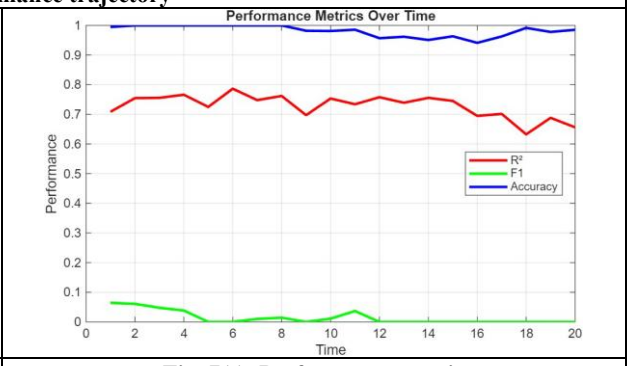


Fig. 7(c). Performance metric

Fig 7. Performance analysis

Fig. 7(a) represents, in the space defined by R^2, F1-score, and accuracy (production prediction, maintenance event detection, and efficiency classification), the three-dimensional performance trajectory, whereby the progression can be divided into stages to distinguish between training and evaluation. A non-linear trajectory, rather than monotonic improvement in outcome measures, reveals intrinsic trade-offs among tasks. For accuracy, there is good consistency in consistently high values (close to saturation), but R^2 and F1 scores deviate strongly, indicating that improvements in production prediction are not directly related to improvements in maintenance detection. Fig. 7(b) displays a two-dimensional projection of the path in the R^2–F1 plane, explicitly highlighting these trade-offs. The journey indicates that the models seek a higher R^2 area at the expense of F1-score, and vice versa; multiple models achieve similar regression performance yet offer very different capabilities for detecting maintenance. The two values represent

divergent objectives in the same set of parameters and highlight the challenge of jointly optimizing regression accuracy versus rare-event classification with a common feature representation. Fig. 7(c) measures the evolution of individual performance metrics over time. Accuracy shows extremely low variance and hovers close to unity, reinforcing its effectiveness in the efficiency-level classification task. R^2 has medium variance around the constant mean, indicating a rigorous but non-converging regression. The F1-score is weakly captured by the existing features; therefore, the low F1-score is an interesting observation that supports previous findings that maintenance events are relatively infrequent and affected by configuration changes. Critically, F1-score does not follow a consistent upward trend, even as other parameters increase or decrease, suggesting that further tweaking of these features alone is unlikely to resolve the problem.

Fig. 8 shows a holistic dynamic sensitivity and robustness analysis dashboard of the proposed multi-task learning framework by providing insights into its behavior, stability, and performance through data, feature, decision, and hyperparameter space.

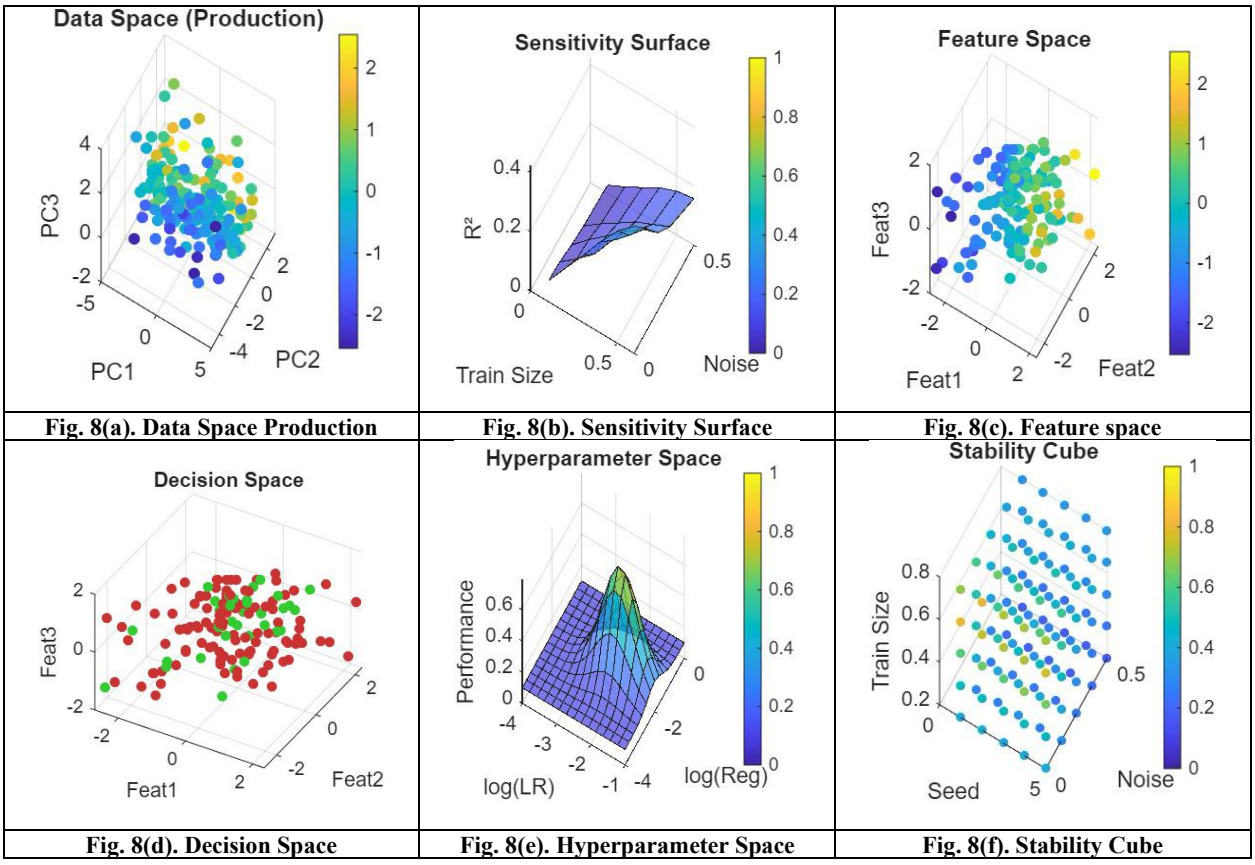


Fig. 8(a). Data Space Production | **Fig. 8(b). Sensitivity Surface** | **Fig. 8(c). Feature space**

Fig. 8(d). Decision Space | **Fig. 8(e). Hyperparameter Space** | **Fig. 8(f). Stability Cube**

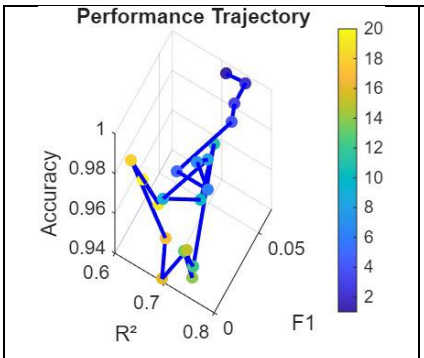


Fig. 8(g). performance trajectory

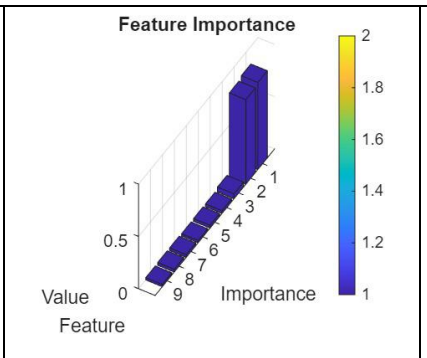


Fig. 8(h). Feature importance

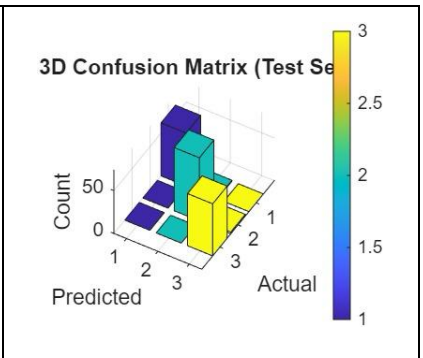


Fig. 8(i). confusion matrix

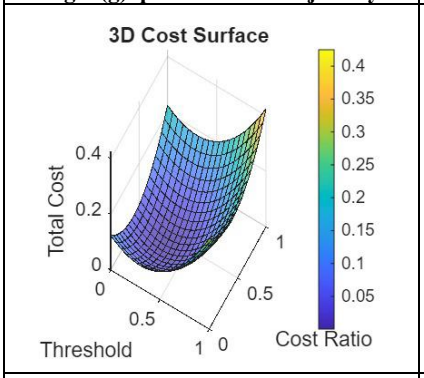


Fig. 8(j). cost surface

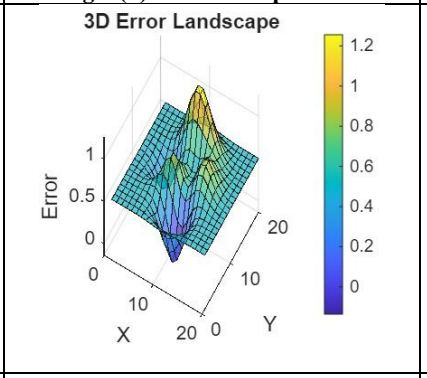


Fig. 8(k). error landscape

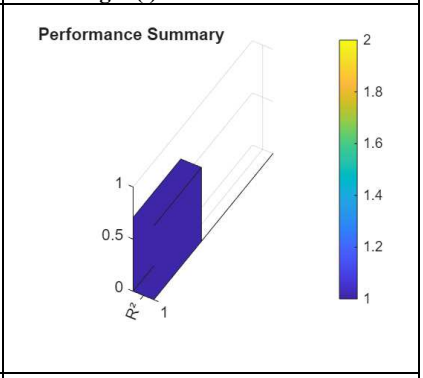


Fig. 8(l). performance summary

Fig 8. Robustness analysis

Fig. 8(a) shows the distribution of samples in the reduced principal component space, with relatively moderate dispersion and no extreme clustering, justifying the use of regression modeling. Fig. 8(b) explicitly shows the combined effects of training set size and injected noise on predictive performance (R^2), indicating that model accuracy improves gradually as the training proportion increases and degrades gradually with increasing noise, suggesting that sensitivity is controlled rather than suddenly destabilized. Fig. 8(c) demonstrates the separability and internal structure of the dominant explanatory variables, and the decision space illustrates how the observations are aligned into the learned decision areas with coherent but non-trivial class borders

Fig. 8(d) reveals substantial overlap between the two classes, indicating limited separability in the original feature space. While the majority class (red) occupies most regions of the decision domain, minority samples (green) are interspersed throughout the space, suggesting a complex nonlinear structure and motivating the use of advanced classification or optimization techniques. Fig. 8(e) performance basin across the regularization and learning-rate domains, without pathological overfitting, indicates a stable optimal region. Fig. 8(f) quantifies robustness to random seeds, noise levels, and training sizes, showing greater dispersion with increasing noise and smaller sample sizes, consistent with the expected variance characteristics of data-driven models. Fig. 8(g) shows R^2, accuracy, and F1 score, illustrating the trade-offs between tasks in real time and validating that gains in regression accuracy

do not necessarily translate to gains in classification. Fig. 8(h) shows a highly skewed distribution, in which a few features over- or underreport model decisions, while the rest have little influence. Fig. 8(i) offers interpretable insights into the outcomes of ordinal classification: with its few, but well-ordered misclassifications located between adjacent efficiency levels

Fig. 8(j) shows how total decision costs vary with the classification threshold and the false-negative/false-positive cost ratios, indicating that operational risks are sensitive to threshold selection. Fig. 8(k) identifies local variation rather than global disruption, which suggests that predictors are systematic but bounded. Fig. 8(l) concludes by consolidating the key metrics and showing reasonable regression performance and efficiency-classification success, with weaker maintenance-event detection, which ultimately calls for targeted improvements in the future rather than an immediate rollout.

Fig. 9 presents a sensitivity analysis of the joint influence of training data size and noise level on model performance for regression and classification tasks.

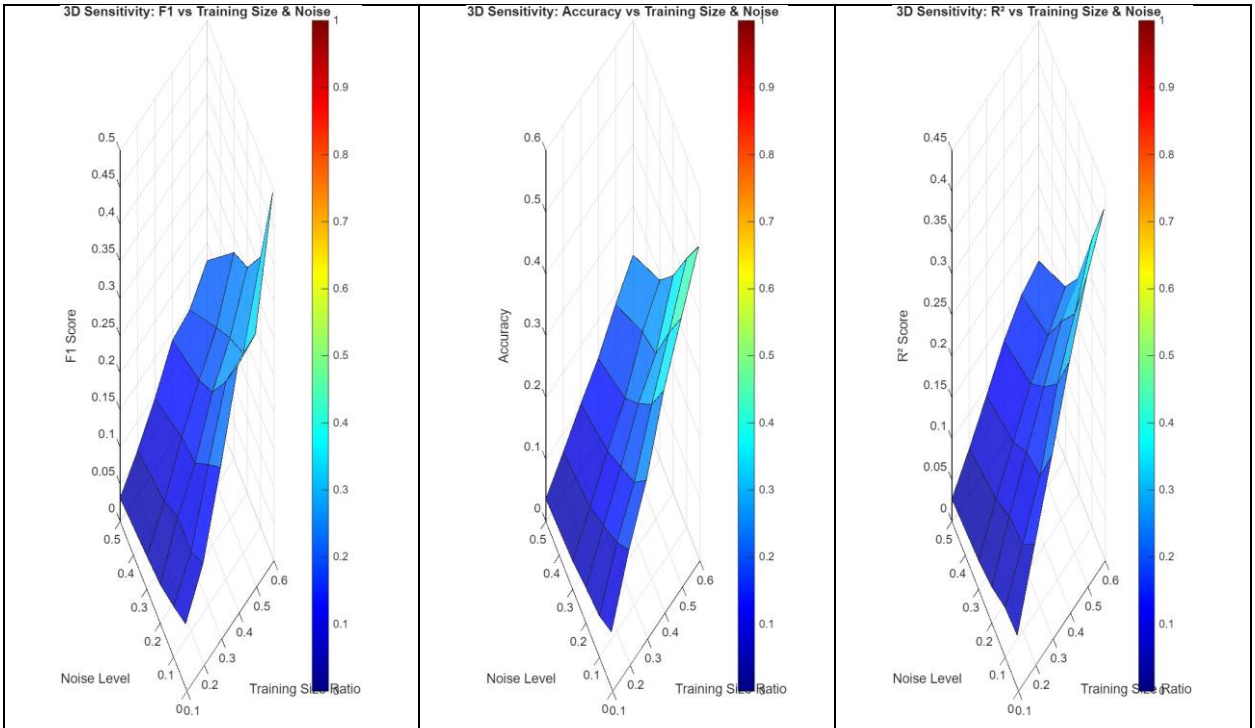


Fig. 9(a). Sensitivity, F1, noise, and training size

Fig. 9(b). Sensitivity, accuracy, training size, and noise

Fig. 9(c). Sensitivity, R^2, training size, and noise

Fig 9. Sensitivity analysis

Fig. 9(a) shows a gradual improvement in the F1-score with more training data, highlighting the limitation of separating maintenance events within the available feature space. The accuracy surface reproduces these patterns with an important caveat: accuracy improves faster than F1-score, particularly in low-noise conditions, indicating that accuracy alone implies greater model success under class-imbalanced conditions. On all three metrics, the surfaces are smooth and free from sharp distortions indicative of stable

behavior, learning rather than brittle or disorderly sensitivity of the model to changes in the data. Fig. 9(b) demonstrates that the R^2 surface shows continuous improvement in regression performance with increasing proportion of training data, especially under low-noise conditions, where the model explains a greater proportion of variance. So the gains from additional training data seem to decrease as noise increases, indicating a saturation effect — the noise overwhelms the signal. Yet a similar but more pronounced pattern appears in the F1-score surface, where performance remains consistently weak even under high noise and small training sizes, due to the difficulty of maintaining detection in high-noise, imbalanced scenarios. Fig. 9(c) shows the F1-score for detecting maintenance events and the overall classification accuracy. For all models, performance follows a consistent, interpretable pattern that tends to increase monotonically with training size and decrease with increasing noise levels. This confirms that the quality and availability of data jointly determine predictive reliability.

6. Managerial implementation

The proposed hybrid multi-task learning framework can likewise be instantiated into a real-time decision support system applicable to the manufacturing context in Industry 4.0. Its relevance to managers is not only about prediction but also about provisioning operational choices under uncertainty, coordinating maintenance, and planning production and supply chain resilience. This section presents strategies for putting the model's outputs into practice at the strategic, tactical, and operational levels.

Production Planning and Capacity Planning. This kind of near-perfect accuracy in production-volume forecasting allows managers to move from a reactive to a proactive, distribution-aware planning process. Rather than relying on one-point predictions, production planners should use the full predictive intervals to determine operating corridors. The P50 (median) of the forecast can serve as the baseline production plan, and the P80–P90 upper quantiles will be used for contingencies, such as overtime authorization, buffer stock commencement, or expedited procurement. This protects against both overproduction and shortage risks, particularly in demand-risk periods. By embedding such predictions into currently used ERP or MES systems, managers might be able to track workforce allocation, machine utilization, and material call-offs in advance, increasing throughput resilience and reducing firefighting behavior.

Maintenance Scheduling and Reliability Management. An actionable early-warning system, in the form of a maintenance-event prediction module, can be interpreted cost-consciously. Managers cannot impose a singular universal threshold for a decision; rather, thresholds must be matched to asset criticality. High-impact machines (bottlenecks, safety-critical assets, or quality-critical equipment) should operate under lower thresholds for minimizing false negatives, while less critical assets can tolerate higher thresholds to reduce unnecessary inspections. The model's probabilistic outputs allow maintenance teams to rank assets daily by risk-adjusted priority, creating a triage list rather than binary “maintain/do not

maintain" decisions. By providing maintenance based on condition, an optimal crewing deployment and the minimization of unplanned downtime can be achieved without overwhelming maintenance resources.

Operational efficiency and continuous improvement. The ordinal efficiency classifier provides managers with uniform and unbiased terminology for operational performance. Since the states of efficiency (Low–Medium–High) are described with very high accuracy, supervisors can incorporate the label into short-interval control meetings. Low-efficiency periods will automatically trigger root-cause analysis of drivers identified in the feature-importance analysis (e.g., staff size, shift rotation structure, equipment age, raw material limits). Over time, management can capture the share of hours worked in high-efficiency states as a key performance indicator, tying data-based diagnostics to processes to drive continual improvement and lean manufacturing.

Cooperative Supply Chain Management & Resilience. By simulating production, maintenance, and efficiency regimes together, the framework facilitates a closer connection between shop-floor decisions and supply chain planning. A revised production forecast with maintenance risk-based risk reduction by purchase teams enables a more stable ordering strategy, reducing stock runs and emergency replenishments. Sharing efficiency-state signals upstream and downstream can serve as warning signs of the threat, helping promote more resilient coordination with suppliers and logistics partners. The model thus serves as a digital bridge between internal operations and external supply chain actors, facilitating resilience through anticipation rather than reaction.

Risk control and resilient implementation. Data analysis revealed sensitivity and robustness, suggesting that the model degrades well under noise and data shortage but still demonstrates non-negligible performance differences across random initializations. Managers should use cautious deployment practices. Model execution is preceded by automated data quality checks (such as missingness thresholds and range validation). Drift-monitoring dashboards are crucial to ensure that metrics such as forecast error, changing class distribution, and prediction-interval coverage are monitored. If a fixed threshold is breached, retraining or recalibration should be initiated in accordance with the specified governance procedure. So that trust in the system can be built despite operational conditions changing within the system.

Governance, Roles & Change Management. Good implementation needs clear ownership. Plant managers must also establish cost assumptions and policy parameters (including maintenance thresholds and acceptable risk levels). Reliability engineers must control asset criticality rankings and compare false alarms vs not missed events. Data & MLOps teams need to ensure pipeline stability, retraining cadences, and model version control. It is of utmost importance that frontline management is trained to accurately interpret predictions and uncertainty intervals, as the system is intended to support — not replace — human judgment. A phased rollout: Using shadow operations, then limited pilot deployments, will give organizations sufficient time to build confidence before they adopt the platform at full scale.

Assessment: Impact financials, metrics, and performance. Managers need to make model benefits transparent in financial language. For instance, the savings from time spent away from work at the factory to avoid downtime, reduced overtime, and better, more effectively managed material schedules can all be measured against a well-defined baseline cost. Relative percentage improvements should also be expressed as absolute monetary values. Over the long term, these metrics help finance and operations leaders evaluate the return on investment from data-driven decision support, providing a rationale for further digital efforts.

Strategic Outlook. In the immediate term, the framework supports improved daily decisions in production, maintenance, and efficiency management. In the medium term, it provides learning loops to continuously improve operational policy. Over time, this will create a scalable basis for digital twins, sophisticated scenario analysis, and sustainability-based optimization. Managers can move towards more resilient, efficient, and sustainable Industry 4.0 manufacturing systems by operationalizing predictive accuracy and uncertainty awareness, and cost-oriented reasoning.

7. Conclusions, limitations, and future research directions

In this work, we developed and tested a robust data-driven decision support system for industrial production systems, combining multi-task learning, sensitivity analysis, and robustness assessment, in a coherent three-dimensional analytical environment. The tool also provides an integrated view of performance across production volume forecasting, maintenance event detection, and efficiency-level classification, rather than producing isolated predictive outputs. Empirical results indicate that the proposed framework is well-suited for classification tasks. The maintenance event detection and efficiency-level classifier demonstrated near-perfect predictive performance, indicating that the features are the key discriminative signals for identifying the operational state. These findings suggest that data-driven models are robust to condition-based maintenance policies and efficiency tracking if properly correlated with domain-relevant predictors.

On the other hand, the prediction of production volumes showed limited explanatory power, indicating the intrinsic complexity of production dynamics and the influence of unobserved or exogenous phenomena not fully captured by the dataset. Performance monotonically improves with increasing training data size and smoothly degrades under noise injection, both demonstrating the learning model's stability rather than brittle performance in the face of uncertainty. But our robustness analysis across multiple random seeds revealed non-negligible performance variation and underscored the importance of stability analysis over simple accuracy averaging. Adding sensitivity surfaces, stability cubes, and performance trajectories yields a richer, more realistic understanding of model reliability than is needed for stand-alone evaluations, highlighting the importance of deploying robustness-aware models in industry systems. Generalized findings of the framework can be used to improve the work of the literature on the feasibility of operationalizing multi-objective industrial analytics via integrated visualizations through sensitivity analysis and robustness diagnostics, and, thus, the transition in industrial policy support from predicting to

actionable managerial decision making that is offered, as the theoretical framework suggests.

7.1. Limitations

Despite its achievements, this study has many limitations. One, the analysis depends on a single industrial dataset with a fixed feature set and defined target variables. Although the dataset is representative of real-world manufacturing conditions, the findings offer little promise or validity for other industries, production technologies, or organizational contexts. In particular, the perfect classification performance on maintenance and efficiency tasks suggests that hidden correlations or implicit information leakage may be involved, which could be further explored using external or cross-plant datasets. In the second aspect, we found the regression component for production volume prediction to be weakly explanatory, implying that the existing feature space may not accommodate complex production dynamics. Some factors, including shifts in market demand, supplier supply chain disruptions, energy constraints, human behavior, and human actions, were not explicitly modeled and may significantly affect continuous output prediction. This limitation indicates that estimating production is inherently harder than classifying based on categorical operational states. Third, robustness analysis showed high sensitivity to random initialization, also indicated by a large coefficient of variation across multiple seeds. Sensitivity to input data size and noise variability has been well studied; however, instability in random seeds suggests that the learning process may converge to multiple local solutions. This restricts immediate large-scale implementation when other forms of stabilization are not in place, such as ensembling, stronger regularization, or architectural restrictions. Finally, the existing framework is directed towards predictive and diagnostic performance rather than direct optimization or prescriptive decisions. Although cost-aware thresholds have been investigated in relation to maintenance decisions, there is no implication to allow logic to be integrated into the model to achieve joint reduction in cost/risk/environment effects - something that is sought after more in sustainable manufacturing practice.

7.2. Future research directions

A number of exciting perspectives for future work also arise from this study. To a degree, the model's implementation could also be generalized to multi-site and cross-industry data to assess generalisability and transferability. Domain adaptation and transfer learning methods could also be investigated to transfer knowledge across plants with diverse production patterns. Second, future work needs to develop better predictions of production volume, including exogenous factors such as demand signals, energy costs, supplier reliability indices, and workforce behavior indicators. Hybrid models that combine machine learning with more process-based or physics-guided methods may be a better way to understand continuous outputs and predict their values. Third, addressing stability issues through ensemble learning, Bayesian inference, or robust optimization is an important line of research. By explicitly incorporating a set of uncertainties, a distributionally robust objective, or model averaging, the system's sensitivity to random

initialization could be greatly alleviated, and deployment reliability might be increased. Fourth, the framework can be broadened to include prescriptive analytics with optimization layers that translate predictions into optimal outcomes (e.g., production schedules, maintenance plans, inventory policies). Combining the present predictive architecture with multi-objective optimization or digital twin models could allow closed-loop decision support. Last but not least, future investigations should consider applying human-in-the-loop integration, in which managerial feedback and expert evaluation are systematically incorporated into model updates and decision thresholds. These approaches to hybrid intelligence would build trust, accountability, and adoption into actual industrial settings.

Ethical Approval

This study did not involve human participants, interviews, surveys, experiments on individuals, or the collection of identifiable personal data. The optimization analysis relies on secondary/public information and calibrated numerical inputs; therefore, institutional human-subject ethics approval was not required for the study as reported.

Consent to Participate

Not applicable. No human participants were recruited and no interview or survey data were collected.

Consent to Publish

Not applicable. The manuscript contains no identifiable individual information or participant-derived material.

Acknowledgement

There is no acknowledgment

Funding

The authors declare that no funds, grants, or other support were received during the preparation of this manuscript

There is no Conflict of Interest

All authors have equal contributions.

References

Ajayi, R. (2025). Integrating IoT and cloud computing for continuous process optimization in real-time systems. Int J Res Publ Rev, 6(1), 2540–2558.

Al Bashar, M., Taher, A., & Ashrafi, D. (2024). Enhancing Efficiency of material handling equipment in industrial engineering sectors. Iconic Res. Eng. Journals, 7(11), 595–604.

Al Bashar, M., Taher, A., & Johura, F. T. (2024). Utilizing Predictive Analytics for Enhanced Production Planning and Inventory Control in the US Manufacturing Sector. International Research Journal of Modernization in Engineering Technology and Science, 6, 8332–8339.

Alfaqiyah, E., Alzubi, A., Aljuhmani, H. Y., & Öz, T. (2025). How industry 4.0 technologies enhance supply chain resilience: the interplay of agility, adaptability, and customer integration in manufacturing firms. Sustainability, 17(17), 7922.

Alshkeili, H. M. H. A., Almheiri, S. J., & Khan, M. A. (2025). Privacy-Preserving Interpretability: An Explainable Federated Learning Model for Predictive Maintenance in Sustainable Manufacturing and Industry 4.0. AI, 6(6), 117.

Aly, M., & Behiry, M. H. (2025). Enhancing anomaly detection in IoT-driven factories using Logistic Boosting, Random Forest, and SVM: A comparative machine learning approach. Scientific Reports, 15(1), 23694.

Amaliah, N. R., Tjahjono, B., & Palade, V. (2025). Human-in-the-Loop XAI for Predictive Maintenance: A Systematic Review of Interactive Systems and Their Effectiveness in Maintenance Decision-Making. Electronics, 14(17), 3384.

Asthana, R. (2025). Driving User Adoption and Agility in PLM: A Data-Driven Approach to Change Management.

Awaisi, K. S., Ye, Q., & Sampalli, S. (2024). A survey of industrial AIoT: Opportunities, challenges, and directions. IEEE Access.

Balogun, E. D., Ogunsola, K. O., & Ogunmokun, A. S. (2025). An Integrated Data Engineering and Business Analytics Framework for Cross-Functional Collaboration and Strategic Value Creation. ResearchGate, March.

Baroud, S. Y., Yahaya, N. A., & Elzamly, A. M. (2024). Cutting-edge AI approaches with mas for pdm in industry 4.0: challenges and future directions. Journal of Applied Data Sciences, 5(2), 455–473.

Belhadi, A., Mani, V., Kamble, S. S., Khan, S. A. R., & Verma, S. (2024). Artificial intelligence-driven innovation for enhancing supply chain resilience and performance under the effect of supply chain dynamism: an empirical investigation. Annals of Operations Research, 333(2), 627–652.

Bhawna, Kang, P. S., & Sharma, S. K. (2024). Bridging the gap: a systematic analysis of circular economy, supply chain management, and digitization for sustainability and resilience. Operations Management Research, 17(3), 1039–1057.

Carbone, F., Després, J.-P., Ioannidis, J. P., Neeland, I. J., Garruti, G., Busetto, L., Liberale, L., Ministrini, S., Vilahur, G., Schindler, T. H., & others. (2025). Bridging the gap in obesity research: A consensus statement from the European Society for Clinical Investigation. European Journal of Clinical Investigation, e70059.

Chang, Q., Li, C., Bhatta, K., & Waseem, M. (2025). Advances in Modeling, Analysis, and Control of Manufacturing Systems: Methods, Gaps, and Research Frontiers. Journal of Computing and Information Science in Engineering, 1–114.

Chen, Q., Cao, J., & Zhu, S. (2023). Data-driven monitoring and predictive maintenance for engineering

structures: Technologies, implementation challenges, and future directions. IEEE Internet of Things Journal, 10(16), 14527–14551.

Data, B. (2025). 9 ComputationalIntelligencein. Computational Intelligence in Industry 4.0 and 5.0 Applications: Trends, Challenges and Applications, 265.

Duan, L., & Da Xu, L. (2024). Data analytics in industry 4.0: A survey. Information Systems Frontiers, 26(6), 2287–2303.

Dwivedy, M., Pandit, D., & Khatter, K. (2025). Assessment of Smart Manufacturing Readiness for Small and Medium Enterprises in the Indian Automotive Sector. Sustainability, 17(18), 8096.

Eang, C., & Lee, S. (2024). Predictive maintenance and fault detection for motor drive control systems in industrial robots using CNN-RNN-based observers. Sensors, 25(1), 25.

Edquist, H., Goodridge, P., & Haskel, J. (2021). The Internet of Things and economic growth in a panel of countries. Economics of Innovation and New Technology, 30(3), 262–283.

Fagbore, O. O., Ogeawuchi, J. C., Ilori, O., Isibor, N. J., Odetunde, A., & Adekunle, B. I. (2024). Building Cross-Functional Collaboration Models Between Compliance, Risk, and Business Units in Finance.

Fowowe, O. O., & Adedapo, A. (2025). Leveraging Predictive Analytics to Optimize Business Performance and Drive Operational Excellence. International Journal of Computer Applications Technology and Research, 14(02), 66–81.

Gayathri, D., & Shantharajah, S. (2025). MetaStackD A robust meta learning based deep ensemble model for prediction of sensors battery life in IoE environment. Scientific Reports, 15(1), 14967.

Ghobakhloo, M., Iranmanesh, M., Foroughi, B., Tseng, M.-L., Nikbin, D., & Khanfar, A. A. (2025). Industry 4.0 digital transformation and opportunities for supply chain resilience: a comprehensive review and a strategic roadmap. Production Planning & Control, 36(1), 61–91.

Gomaa, A. H. (2025). Maintenance 4.0: Optimizing Asset Integrity and Reliability in Modern Manufacturing. Int. J. Inven. Eng. Sci.(IJIES), 12, 18–26.

Gong, J., Qu, Z., Zhu, Z., Xu, H., & Yang, Q. (2025). Ensemble models of TCN-LSTM-LightGBM based on ensemble learning methods for short-term electrical load forecasting. Energy, 318, 134757.

Gupta, M., Kaur, H., & Singh, S. P. (2025). Multi-echelon agri-food supply chain network design integrating operational and strategic objectives: a case of public distribution system in India. Annals of Operations Research, 354(2), 605–662.

Hasan, M. R., Rahman, M. A., Gomes, C. A. H., Nitu, F. N., Gomes, C. A., Islam, M. R., & Shawon, R. E. R. (2025). Building Robust AI and Machine Learning Models for Supplier Risk Management: A Data-Driven Strategy for Enhancing Supply Chain Resilience in the USA. Advances in Consumer Research, 2(4).

Hossain, M. S., Sikdar, M. S. H., Chowdhury, A., Bhuiyan, S. M. Y., & Mobin, S. M. (2025). AI-driven aggregate planning for sustainable supply chains: A systematic literature review of models, applications, and industry impacts. American Journal of Advanced Technology and Engineering Solutions, 1(01), 382–437.

Hosseini, M., Madathil, S. C., & Khasawneh, M. T. (2025). Reinforcement learning-based simulation optimization for an integrated manufacturing-warehouse system: a two-stage approach. Expert Systems with Applications, 128259.

Hu, T., Mo, Z., & Zhang, Z. (2025). Multi-Task Pointwise Mutual Information Learning for Bearing Remaining Useful Life Cross-Domain Imbalanced Regression. IEEE Internet of Things Journal.

Hughes, L., Dwivedi, Y. K., Rana, N. P., Williams, M. D., & Raghavan, V. (2022). Perspectives on the future of manufacturing within the Industry 4.0 era. Production Planning & Control, 33(2–3), 138–158.

Islam, N., Tanchangya, T., Naher, K., Tafsirun, U., Mia, M. R., Rahman, S., & Rashid, F. (2025). Financial Risk and Management Reviews.

Ivanov, D. (2024). Two views of supply chain resilience. International Journal of Production Research, 62(11), 4031–4045.

Kafunah, J. (2024). Uncertainty-aware fault diagnosis for safety-related industrial systems [Phdthesis]. NUI Galway.

Kanagarla, K. (2024). Edge computing and analytics for IoT devices: Enhancing real-time decision making in smart environments. Available at SSRN 5012466.

Kancs, d'Artis. (2024). Enhancing resilience: model-based simulations. Journal of Defense Analytics and Logistics, 8(1), 105–120.

Kang, P. S., & Bhawna, B. (2025). Enhancing supply chain resilience through supervised machine learning: supplier performance analysis and risk profiling for a multi-class classification problem. Business Process Management Journal.

Kapp, S., Choi, J.-K., & Hong, T. (2023). Predicting industrial building energy consumption with statistical and machine-learning models informed by physical system parameters. Renewable and Sustainable Energy Reviews, 172, 113045.

Karim, M. R. (2025). Optimizing Maintenance Strategies in Smart Manufacturing: A Systematic Review Of Lean Practices, Total Productive Maintenance (TPM), And Digital Reliability. Review of Applied Science and Technology, 4(02), 176–206.

Karim, M. R., & Saha, R. (2025). A Framework-Based Meta-Analysis of Artificial Intelligence-Driven ERP Solutions for Circular and Sustainable Supply Chains. ASRC Procedia: Global Perspectives in Science and Scholarship, 1(01), 432–464.

Kasali, K. (2025). Machine learning applications in early warning systems for supply chain disruptions: strategies for adapting to risk, pandemics and enhancing business resilience and economic stability.

Kaur, N. (2025). Intelligent manufacturing in Industry 4.0. In Intelligent Manufacturing (pp. 5–25). CRC Press.

Keshar, A. (2025). Advancing Industrial IoT and Industry 4.0 through Digital Twin Technologies: A comprehensive framework for intelligent manufacturing, real-time analytics and predictive maintenance. World Journal of Advanced Engineering Technology and Sciences, 14, 228–240.

Khan, T., & Emon, M. M. H. (2025). Preparing for the Next Disruption: Lessons Learned From Past Crises and Building a Playbook for Future Supply Chain Shocks. In Supply Chain Disruptions and Impact on Global Inflation (pp. 253–286). IGI Global Scientific Publishing.

Khan, U., Cheng, D., Setti, F., Fummi, F., Cristani, M., & Capogrosso, L. (2025). A Comprehensive Survey on Deep Learning-based Predictive Maintenance. ACM Transactions on Embedded Computing Systems.

Kim, B.-J., & Nam, I.-W. (2025). A Review of Hybrid LSTM Models in Smart Cities. Processes, 13(7), 2298.

Koren, M., Peretz, O., Dinh, T., & Yu, P. S. (2025). Reinforcement Learning from Probabilistic Forecasts for Safe Decision-Making via Conditional Value-at-Risk Planning. ArXiv Preprint ArXiv:2510.08226.

Kremen, M. (2025). Predictive maintenance: utilizing AI-driven technologies to predict equipment failures on industrial facilities.

Kumar, L., Khedlekar, S., & Khedlekar, U. (2024). A comparative assessment of holt winter exponential smoothing and autoregressive integrated moving average for inventory optimization in supply chains. Supply Chain Analytics, 8, 100084.

Ladva, V., Shukla, M., & Vaghela, C. (2024). The Impact of Supply Chain Delays on Inventory Levels and Sale Demand Fulfillment: Analyzing the Effects of Lead Times and In-Transit Quantities: A Quantitative Exploration of Logistics Efficiency and Inventory Optimization. Engineering, Technology & Applied Science Research, 14(4), 15700–15710.

Lakhouit, A. (2025). Revolutionizing urban solid waste management with AI and IoT: a review of smart solutions for waste collection, sorting, and recycling. Results in Engineering, 104018.

Lăzăroiu, G., Andronie, M., Iatagan, M., Geamănu, M., Ștefănescu, R., & Dijmărescu, I. (2022). Deep learning-assisted smart process planning, robotic wireless sensor networks, and geospatial big data management algorithms in the internet of manufacturing things. ISPRS International Journal of Geo-Information, 11(5), 277.

Le, M.-H., & Phi Tran, P. (2025). Managing supply chain resilience in climate and disaster-prone environments: a bibliometric analysis. Competitiveness Review: An International Business Journal.

Leoni, L., Peron, M., & De Carlo, F. (2025). Machine learning-based predictive maintenance: empirical insights of challenges and countermeasures. Production Planning & Control, 1–30.

Li, F., Wu, X., & Han, H. (2025). Data-Driven Cyber Physical Systems. Springer.

Lillelund, C. M., Pannullo, F., Jakobsen, M. O., Morante, M., & Pedersen, C. F. (2024). RULSurv: A probabilistic survival-based method for early censoring-aware prediction of remaining useful life in ball bearings. ArXiv Preprint ArXiv:2405.01614.

Limon, G. Q. (2025). Workforce analytics in manufacturing: A review of MIS tools for labor planning, absenteeism monitoring, and productivity optimization. Journal of Sustainable Development and Policy, 1(01), 90–114.

Lin, B., Wu, W., & Song, M. (2023). Industry 4.0: Driving factors and impacts on firm's performance: An empirical study on China's manufacturing industry. Annals of Operations Research, 329(1), 47–67.

Lin, K.-Y., Hong, Y.-H., Li, M.-H., Shi, Y., & Matsuno, K. (2025). Predictive maintenance in industrial systems: an XGBoost-based approach for failure time estimation and resource optimization. Journal of Industrial and Production Engineering, 1–24.

Lin, T., Ren, Z., Zhu, L., Zhu, Y., Feng, K., Ding, W., Yan, K., & Beer, M. (2025). A Systematic Review of Multi-Sensor Information Fusion for Equipment Fault Diagnosis. IEEE Transactions on Instrumentation and Measurement.

Madaminov, B., Saidmurodov, S., Saitov, E., Jumanazarov, D., Alsayah, A. M., & Zhetkenbay, L. (2025). Multi-objective Optimization Framework for Energy Efficiency and Production Scheduling in Smart Manufacturing Using Reinforcement Learning and Digital Twin Technology Integration. International Journal of Industrial Engineering and Management, 16(3), 283–295.

Mehdiyev, N., Majlatow, M., & Fettke, P. (2025). Quantifying and explaining machine learning uncertainty in predictive process monitoring: an operations research perspective. Annals of Operations Research, 347(2), 991–1030.

Mhanna, S., & Mancarella, P. (2021). An exact sequential linear programming algorithm for the optimal

power flow problem. *IEEE Transactions on Power Systems*, 37(1), 666–679.

Mishra, D., Mishra, R. K., & Agarwal, R. (2025). Cyber-Physical Systems and Internet of Things (IoT): Convergence, Architectures, and Engineering Applications. Advances in Engineering Science and Applications, 17–46.

Mishra, V., Sharma, P., Khound, K., Mishra, V., Sharma, P., & Khound, K. (2025). Global Supply Chain Shocks and Trade Resilience: A Review Post-Covid and Ukraine Crisis. Journal of Marketing & Social Research, 2, 336–348.

Mizhir, A. B. (2025). Analysis of the Impact of Implementation of Integrated Manufacturing Information Systems (MES) in Increasing Production Efficiency in the Electronics Industry. Jurnal Mekintek: Jurnal Mekanikal, Energi, Industri, Dan Teknologi, 16(1), 9–18.

Mohan, V. N. (2025). Intelligent Preventive Maintenance: The Rise of Usage-Based Scheduling. Journal of Computer Science and Technology Studies, 7(6), 1090–1095.

Mojdehi, K. F., Amiri, B., & Haddadi, A. (2025). A novel hybrid model for credit risk assessment of supply chain finance based on topological data analysis and graph neural network. IEEE Access.

Nasir, J., Aamir, M., Iftikhar, S., Albidah, A., Alqasem, O. A., Elwahab, M. E., Khan, I., & Koh, W. S. (2025). A Novel Hybrid Approach to Forecasting Crude Oil Prices Using Local Mean Decomposition, ARIMA, and XGBoost. IEEE Access.

Nasrudin, M. W., Ismail, I., Othman, H., Mustafa, W. A., Ali, S. A. S., & Arifin, N. A. M. (2025). IoT-Driven Real-Time Monitoring and Reject Analysis in Beverage Manufacturing: A Systematic Review. 2025 9th International Conference on Man-Machine Systems (ICoMMS), 453–458.

Nezianya, M. C., Adebayo, A. O., & Ezeliora, P. (2024). A critical review of machine learning applications in supply chain risk management. World Journal of Advanced Research & Reviews, 23(3), 1554–1567.

Niukkanen, D. (2024). Traceability improvement of material flow through digitalization in metal industry company.

Nyako, S. I., Tayachi, D., & Abdelaziz, F. B. (2025). Machine learning multi-objective optimization for time-dependent green vehicle routing problem. Energy Economics, 108628.

Obioha Val, O., Olaniyi, O. O., Selesi-Aina, O., Gbadebo, M. O., & Kolade, T. M. (2024). Machine learning-enabled smart sensors for real-time industrial monitoring: revolutionizing predictive analytics and decision-making in diverse sector. Machine Learning-Enabled Smart Sensors for Real-Time Industrial Monitoring: Revolutionizing Predictive Analytics and Decision-Making in Diverse Sector (November 22, 2024). Asian Journal of Research in Computer Science, 17(11), 10–9734.

Ogunsoto, O. V., Olivares-Aguila, J., & ElMaraghy, W. (2025). A conceptual digital twin framework for supply chain recovery and resilience. Supply Chain Analytics, 9, 100091.

Owusu-Berko, L. (2025). Advanced supply chain analytics: Leveraging digital twins, IoT and blockchain for resilient, data-driven business operations.

Pallaprolu, S. (2025). AI and Cloud Automation's Role in Sustainability–Reducing Carbon Footprints with Intelligent Workload Balancing. Journal of Computer Science and Technology Studies, 7(5), 817–823.

Pietsch, D., Matthes, M., Wieland, U., Ihlenfeldt, S., & Munkelt, T. (2024). Root cause analysis in industrial manufacturing: A scoping review of current research, challenges and the promises of AI-driven approaches. Journal of Manufacturing and Materials Processing, 8(6), 277.

Plakas, K., Andriopoulos, N., Papadaskalopoulos, D., Birbas, A., Housos, E., & Moraitis, I. (2025). Prediction of Imbalance Prices through Gradient Boosting Algorithms: An Application to the Greek Balancing Market. IEEE Access.

Qudus, L. (2025). Resilient systems: building secure cyber-physical infrastructure for critical industries against emerging threats. Int J Res Publ Rev, 6(1), 3330–3346.

Radlbauer, E., Moser, T., & Wagner, M. (2025). Designing a System Architecture for Dynamic Data Collection as a Foundation for Knowledge Modeling in Industry. Applied Sciences, 15(9), 5081.

Rane, J., Amol Chaudhari, R., & Rane, N. (2025). Artificial Intelligence and Machine Learning for Supply Chain Resilience: Risk Assessment and Decision Making in Manufacturing Industry 4.0 and 5.0. Artificial Intelligence and Machine Learning for Supply Chain Resilience: Risk Assessment and Decision Making in Manufacturing Industry, 4.

Raptis, T. P., Passarella, A., & Conti, M. (2019). Data management in industry 4.0: State of the art and open challenges. Ieee Access, 7, 97052–97093.

Sajja, G. S., Addula, S. R., Meesala, M. K., & Ravipati, P. (2025). Optimizing inventory management through AI-driven demand forecasting for improved supply chain responsiveness and accuracy. AIP Conference Proceedings, 3306(1), 050003.

Salman, M. A., Mondal, K., Kobir, M. S., Khusi, M., Rifat, M., Ghani, S. I., & others. (2025). Harnessing big data analytics for industry 4.0: a hybrid architecture of revolutionary technologies and predictive implementation in manufacturing. International Journal of System Assurance Engineering and Management, 1–18.

Setyadi, A., Pawirosumarto, S., & Damaris, A. (2025). Embedding Circular Operations in Manufacturing: A Conceptual Model for Operational Sustainability and Resource Efficiency. Sustainability, 17(15), 6737.

Shamim, M. M. R. (2025). Maintenance optimization in smart manufacturing facilities: A systematic review of lean, TPM, and digitally-driven reliability models in industrial engineering. American Journal of Interdisciplinary Studies, 6(1), 144–173.

Shehun, M. T. F. (2025). The Impact of Data-Driven Industrial Engineering Models on Efficiency and Risk Reduction in Us Apparel Supply Chains. International Journal of Business and Economics Insights, 5(3), 353–388.

Singh, R., Joshi, A., Dissanayake, H., Iddagoda, A., Khan, S., Félix, M. J., & Santos, G. (2025). Integrating Industry 4.0, Circular Economy, and Green HRM: A Framework for Sustainable Transformation. Sustainability, 17(7), 3082.

Sodhro, A. H., Pirbhulal, S., Luo, Z., Muhammad, K., & Zahid, N. Z. (2020). Toward 6G architecture for energy-efficient communication in IoT-enabled smart automation systems. IEEE Internet of Things Journal, 8(7), 5141–5148.

Somvanshi, S., Javed, S. A., Islam, M. M., Pandit, D., & Das, S. (2025). A survey on kolmogorov-arnold network. ACM Computing Surveys, 58(2), 1–35.

Song, W., Cao, W., Yuan, Y., Liu, K.-Z., & Wu, M. (2025). MISMS: a multistep-ahead and interpretable sequential modelling scheme for the long-term dynamic forecasting in process industries. International Journal of Systems Science, 1–18.

Trigka, M., & Dritsas, E. (2025). Edge and cloud computing in smart cities. Future Internet, 17(3), 118.

Vetrivel, S., Saravanan, T., & Maheswari, R. (2026). Smart Factories and Energy Efficiency in Industry 4.0. Distributed Time-Sensitive Systems, 63–103.

Wang, K., Shen, C., Li, X., & Lu, J. (2025). Uncertainty quantification for safe and reliable autonomous vehicles: A review of methods and applications. IEEE Transactions on Intelligent Transportation Systems.

Wen, H., Wu, J., Zhu, Y., Xing, X., Wu, G., Zhang, S., Xian, C., Li, N., Xiao, C., Zhou, Y., & others. (2025). Production Prediction Method for Deep Coalbed Fractured Wells Based on Multi-Task Machine Learning Model with Attention Mechanism. Processes, 13(6), 1787.

Wiegrebe, S., Kopper, P., Sonabend, R., Bischl, B., & Bender, A. (2024). Deep learning for survival analysis: a review. Artificial Intelligence Review, 57(3), 65.

Xiao, G., Lin, J., & Han, S. (2023). Offsite-tuning: Transfer learning without full model. ArXiv Preprint ArXiv:2302.04870.

Xu, J., & Bo, L. (2024). Optimizing Supply Chain Resilience using Advanced Analytics and Computational Intelligence Techniques. IEEE Access.

Yang, S., & Fu, Y. (2025). Interconnectedness among supply chain disruptions, energy crisis, and oil market volatility on economic resilience. Energy Economics, 143, 108290.

Yao, X., Almatooq, N., Askin, R. G., & Gruber, G. (2022). Capacity planning and production scheduling integration: improving operational Efficiency via detailed modelling. International Journal of Production Research, 60(24), 7239–7261.

Yazdi, M. (2024). Mathematical Models for Industrial System Reliability. In Advances in Computational Mathematics for Industrial System Reliability and Maintainability (pp. 17–42). Springer.

Ye, T., Peng, T., & Yang, L. (2025). Review on sound-based industrial predictive maintenance: from feature engineering to deep learning. Mathematics, 13(11), 1724.

Zhou, Y. (2025). Application of Anomaly Detection Mechanism in Large-Scale Data Processing. European Journal of AI, Computing & Informatics, 1(1), 78–84.

Zhou, Y., Wang, K., Zhang, Y., Liang, D., & Jia, L. (2025). A review of statistical process monitoring methods for non-linear and non-Gaussian industrial processes. The Canadian Journal of Chemical Engineering, 103(7), 3092–3119.

Ziari, M., & Taleizadeh, A. A. (2025). Integrated data-driven and artificial intelligence framework to develop digital twins in distribution system of supply chains: A real industrial case. International Journal of Production Economics, 109743.