APPLICATION OF MACHINE LEARNING MODELS FOR PREDICTIVE MAINTENANCE OF BIOTECHNICAL SYSTEMS
/
UTILIZAREA MODELELOR DE ÎNVĂȚARE AUTOMATĂ PENTRU MENTENANȚA PREDICTIVĂ A SISTEMELOR BIOTEHNICE

Adrian IOSIF[1], Edmond MAICAN[1], Sorin Ştefan BIRIŞ[1], Nicolae-Valentin VLĂDUȚ[2]
[1]Faculty of Biotechnical Systems Engineering, University POLITEHNICA of Bucharest / Romania [2]INMA Bucharest / Romania
E-mail: biris.sorinstefan@gmail.com
DOI: https://doi.org/10.35633/inmateh-75-79

Keywords: *Predictive maintenance, Machine learning, Anomaly detection, Agricultural machinery, Hydraulic system monitoring, Multi-model voting framework, Sensor data analysis.*

ABSTRACT
Ensuring the reliability and efficiency of agricultural machinery is critical for modern farming operations. Traditional maintenance strategies, including corrective and preventive approaches, often lead to unexpected downtime or excessive servicing costs. This study explores the application of machine learning-based predictive maintenance for agricultural equipment, focusing on the hydraulic system of a Massey Ferguson 7700 S tractor. Real-time sensor data was collected, with hydraulic pressure selected as the primary diagnostic metric for detecting early signs of mechanical degradation. A predictive maintenance framework was developed using seven machine learning models: Isolation Forest, One-Class SVM, KMeans, DBSCAN, Autoencoders, Convolutional Neural Networks (CNNs), and XGBoost. These models were individually applied to identify pressure anomalies indicative of potential failures. To enhance detection accuracy, a "Council of the Wise" ensemble approach was introduced, where an anomaly was validated only if at least four of the seven models agreed on its presence. This consensus-based method reduced false positives and improved fault identification reliability. Results demonstrated that integrating multiple models effectively distinguished significant anomalies from noise, capturing both transient mechanical instabilities and gradual wear-related failures. The findings highlight the potential of machine learning-driven predictive maintenance in optimizing maintenance schedules, reducing unplanned downtime, and extending equipment lifespan. This study establishes a scalable, data-driven maintenance approach that enhances the operational resilience of agricultural machinery, ensuring greater efficiency and sustainability in farming operations.

REZUMAT
Asigurarea fiabilității și eficienței utilajelor agricole este esențială pentru agricultura modernă. Strategiile tradiționale de mentenanță, inclusiv abordările corective și preventive, conduc adesea la perioade neprevăzute de indisponibilitate sau la costuri excesive de întreținere. Acest studiu investighează posibilitatea aplicării unui sistem de mentenanță predictivă pentru echipamentele agricole bazat pe învățarea automată, analizînd date achiziționate din sistemul hidraulic al tractorului Massey Ferguson 7700 S. Datele au fost colectate în timp real, presiunea din instalația hidraulică fiind selectată ca parametru principal pentru detectarea timpurie a eventualelor defecțiuni. A fost dezvoltat un sistem de mentenanță predictivă bazat pe șapte modele de învățare automată: Isolation Forest, One-Class SVM, KMeans, DBSCAN, Autoencoders, Convolutional Neural Networks (CNNs) și XGBoost. Aceste modele au fost aplicate individual pentru a identifica anomalii ale presiunii, fapt care ar putea scoate în evidență posibile defecțiuni. Pentru a îmbunătăți acuratețea detecției, a fost introdusă o abordare de tip „sfatul înțelepților”, în care o anomalie a fost validată doar dacă cel puțin patru dintre cele șapte modele au confirmat prezența acesteia. Această metodă bazată pe consens a redus numărul de alarme false și a îmbunătățit fiabilitatea identificării defecțiunilor. Rezultatele au demonstrat că integrarea mai multor modele a permis diferențierea eficientă a anomaliilor semnificative de zgomot, evidențiind atât instabilitățile mecanice tranzitorii, cât și degradările progresive. Concluziile evidențiază potențialul mentenanței predictive bazată pe învățarea automată în optimizarea programelor de întreținere, reducerea timpilor neplanificați de oprire și prelungirea duratei de viață a echipamentelor. Metoda propusă pentru mentenanța utilajelor agricole este scalabilă și poate îmbunătăți semnificativ reziliența operațională, eficiența și sustenabilitatea în sectorul agricol.

INTRODUCTION

Maintenance strategies in industrial and agricultural machinery have evolved to address operational reliability and cost efficiency. Traditional approaches include corrective maintenance, which involves repairing equipment only after a failure has occurred, often leading to unexpected downtime and high repair costs. Preventive maintenance seeks to mitigate these risks by scheduling routine inspections and servicing at fixed intervals, regardless of the actual condition of the equipment. While this approach reduces the likelihood of failures, it may result in unnecessary maintenance and higher operational costs.

Predictive maintenance, powered by machine learning and real-time data, proactively detects early signs of wear or malfunction, replacing fixed schedules and reactive repairs. By identifying subtle anomalies and degradation patterns, it optimizes maintenance scheduling, reducing unnecessary interventions while preventing costly breakdowns. Advanced algorithms enhance precision in fault detection, ensuring efficient resource allocation and extending equipment lifespan. Increasingly adopted in the agriculture and food industries, predictive maintenance improves equipment reliability, minimizes downtime, and lowers maintenance costs. Conventional methods, such as corrective maintenance, addressing failures after they occur, and preventive maintenance, based on predefined schedules, can lead to inefficiencies and unexpected failures. Machine learning models provide a more intelligent, data-driven approach, enhancing fault prediction accuracy and overall system performance.

This paper explores the application of machine learning (ML) models in predictive maintenance for biotechnical systems, evaluating their effectiveness in improving system reliability and efficiency. By examining the capabilities of various machine learning techniques in failure prediction and early fault detection, the study highlights the advantages of a data-driven approach in minimizing downtime and reducing maintenance costs. The potential of predictive maintenance extends beyond cost savings, it enhances operational resilience, ensuring that agricultural and food production systems remain sustainable and efficient in the face of growing industry demands (*Lima et al., 2021*).

Agricultural machinery, such as the Massey Ferguson 7700 S tractor, operates under demanding conditions that require continuous performance monitoring to ensure efficiency and longevity. During various soil operations, including plowing, tillage, and planting, as well as harvesting corn or other crops, the hydromechanical systems of agricultural machinery, including the hydraulic system, engine, and drivetrain, experience dynamic mechanical stresses. Modern real-time condition monitoring systems integrated into such equipment collect high-frequency operational data, capturing parameters such as engine temperature, hydraulic pressure, vibration levels, fuel consumption, and mechanical load fluctuations. These data streams, when processed using machine learning models, allow for the early identification of potential failures, preventing costly breakdowns and minimizing downtime.

Ouadah et al. (2022) evaluates various supervised machine learning algorithms for predictive maintenance, focusing on their classification and regression capabilities. Key algorithms analyzed include Random Forest, Decision Trees, and k-Nearest Neighbors (KNN), tested on both real-world and simulation datasets. The findings indicate that:

- Random Forest and Decision Trees perform similarly in small datasets and excel in regression-based reliability prediction;
- KNN proves to be more effective for classification tasks, particularly in handling large volumes of data;
- the use of vibration analysis and reliability evaluation enhances failure detection accuracy, supporting early fault diagnosis in industrial equipment.

Furthermore, the study highlights vibration analysis as one of the most reliable PdM techniques, particularly for detecting faults in rotating machinery such as misalignment, imbalance, and bearing degradation. Additional monitoring methods include infrared thermography, acoustic analysis, lubricant analysis, and ultrasonic testing, all of which contribute to data collection for ML models. The results emphasize that PdM strategies, when combined with ML techniques, can reduce system downtime, optimize maintenance scheduling, and minimize costs. However, challenges such as data quality, model interpretability, and real-time deployment constraints remain critical factors in PdM implementation. The study concludes that while ML-based predictive maintenance provides a more dynamic and cost-effective approach than traditional methods, continued research is needed to refine models, improve scalability, and enhance real-time decision-making.

Predictive maintenance (PdM) is an essential component of modern industrial operations, aiming to optimize maintenance actions by predicting failures before they occur. Traditional PdM strategies rely on statistical models and machine learning (ML) techniques, but reinforcement learning (RL) has emerged as a promising approach for designing autonomous and adaptive maintenance strategies.

Siraskar et al. (2023) explores how reinforcement learning is applied to predictive maintenance across various industrial contexts, from early fault detection to health index modeling and maintenance scheduling. Unlike conventional ML techniques, which often require large amounts of labeled historical data, RL-based approaches learn optimal maintenance policies dynamically through interaction with the system.

A contribution of this work is the development of a taxonomy of RL-based predictive maintenance strategies, distinguishing between model-based RL, model-free RL, Markov Decision Processes (MDP), Partially Observable MDPs (POMDPs), and Semi-Markov Decision Processes (SMDPs). The paper highlights the advantages of RL over traditional PdM methods, including its ability to handle dynamic, non-stationary environments, optimize maintenance decisions in real time, and reduce overall maintenance costs. The review also outlines practical challenges in applying RL to PdM, including data quality issues, real-time deployment constraints, and the need for robust edge computing solutions. Additionally, the study emphasizes the importance of reward function design, exploration-exploitation trade-offs, and agent training efficiency in ensuring successful RL-based maintenance systems.

Meddaoui et al. (2024) focuses on remaining useful life (RUL) estimation, a key metric in predictive maintenance that determines the expected time before a component or system fails. Various ML models, including support vector machines (SVM), random forests (RF), k-nearest neighbors (KNN), and artificial neural networks (ANN), are evaluated to determine their effectiveness in predicting RUL. The findings indicate that:

- KNN and ANN models achieve the highest classification accuracy, with KNN reaching 95.92% accuracy in failure prediction;
- Random forests (RF) and principal component analysis (PCA) provide the best results for regression-based RUL estimation, demonstrating lower prediction error rates;
- feature engineering and data preprocessing techniques, including normalization, principal component analysis (PCA), and attribute selection methods, significantly impact model performance by reducing data dimensionality and improving predictive accuracy.

Recent advancements in artificial intelligence have revolutionized predictive maintenance strategies, enabling more efficient fault detection, failure prediction, and maintenance optimization (*O'Neil et al., 2022*). A study on deep learning-driven architectures for predictive maintenance highlights how machine learning techniques, particularly convolutional neural networks (CNNs), deep recurrent neural networks (DRNNs), stacked autoencoders (SAEs), and deep belief networks (DBNs), have evolved to address the complex challenges of reliability engineering in industrial systems. *Li et al.* (*2024*) systematically reviews deep learning-driven architectures applied to PdM, examining their advantages, limitations, and adaptability across various industrial scenarios.

The authors highlight how Industry 4.0 advancements and IIoT (Industrial Internet of Things) have enabled PdM models to process vast industrial datasets in real time. CNNs and autoencoders outperform traditional ML techniques by automatically extracting critical degradation patterns from raw sensor signals. LSTMs and DRNNs are particularly effective in failure prediction and RUL estimation, capturing long-term dependencies in machinery operations. Long Short-Term Memory (LSTM) networks, trained on historical operational data, provided fault predictions with an accuracy exceeding 95 percent. As a result, repair response time was reduced by 40 percent, and maintenance costs were lowered by 20 percent.

Serradilla et al. (2021) provide a comprehensive review of deep learning (DL) applications in predictive maintenance (PdM), addressing challenges in selecting suitable architectures for industrial use cases. Given the increasing volume of industrial data in the Industry 4.0 era, the study highlights how advanced DL techniques can optimize maintenance strategies by predicting failures before they occur, thus minimizing downtime and reducing costs. The authors categorize and analyze various state-of-the-art (SotA) deep learning architectures, including Self-Organizing Maps (SOM), One-Class Neural Networks (OCNN), and Generative Models, examining their adaptability to real-world industrial maintenance scenarios. The paper also systematically evaluates PdM data characteristics, reviews statistical and traditional ML techniques as baselines, and compares DL approaches against benchmark datasets. A key contribution of this work is the comparison of DL models on a turbofan engine degradation dataset, providing insights into performance, adaptability, and reproducibility in PdM applications.

Zhang et al. (*2023*) explores an optimized approach to configuring maintenance service vehicles for agricultural machinery operations. It presents a method for resource allocation that minimizes costs while enhancing the efficiency of maintenance services. Using a service coverage model, the research optimizes vehicle positioning and dispatch to ensure timely repairs while reducing overall operational expenses. An improved genetic algorithm is implemented to address the challenge of optimizing service vehicle deployment, incorporating enhanced selection, crossover, and mutation techniques to refine resource distribution.

Wang et al. (2023) show that autoencoders, a type of neural network designed for unsupervised learning, can effectively extract meaningful fault-related features from large datasets. This approach enables the mapping of failure causes to observed fault phenomena, such as cutter winding, roller clogging, bridge blockage, chain wear, and excessive tool clearance. The results show that autoencoders outperform traditional rule-based and statistical models by dynamically learning complex patterns from sensor data, rather than relying on fixed analytical equations. A key advantage of ML-based predictive maintenance is its ability to identify anomalies before they lead to critical failures. For instance, when analyzing historical harvester failure data, an ML model trained on normalized and preprocessed sensor inputs achieved a prediction accuracy of over 93%, significantly outperforming support vector machines (SVMs), Bayesian classifiers, and sparse self-coding networks. The confusion matrix analysis confirmed that autoencoder-based models exhibited the lowest false-positive and false-negative rates, making them highly reliable for real-time predictive applications. In predictive modeling for refrigeration systems, sensor fusion techniques have been used to combine vibration, temperature, and pressure sensor data. ML algorithms process these multi-source inputs, allowing for early fault detection by recognizing subtle deviations from normal operating conditions. Experimental results indicate that deep learning models reduce false alarm rates by 27% compared to traditional threshold-based monitoring methods, providing a significant improvement in predictive accuracy.

To assess the effectiveness of ML models in predictive maintenance, *Wang et al. (2023)* have conducted comparative analyses of different approaches: self-coding neural networks (autoencoders), sparse self-coding networks, a variant of traditional autoencoders, and RotGBM, a hybrid gradient-boosting model.

Reliability is a fundamental concern in agricultural machinery, as failures in critical components can lead to significant productivity losses, increased maintenance costs, and operational inefficiencies. Traditional failure analysis methods rely on Failure Mode, Effects, and Criticality Analysis (FMECA) to identify high-risk components. However, these approaches often suffer from subjectivity and difficulty in data acquisition. By utilizing data mining techniques, researchers have been able to systematically analyze large datasets of operational faults, extract meaningful failure patterns, and enhance predictive maintenance strategies (*Yang at al., 2022*). The study of failure data collected from grain harvesters has identified the cutter component as the most failure-prone, with blade degradation being a dominant failure mode. Excessive clearance, misalignment, and improper installation contribute to high failure rates, requiring frequent inspections and maintenance interventions. Data-driven failure analysis has enabled the classification of failure causes into four primary categories: failures due to mechanical wear, accidental failures, operational errors, and failures caused by insufficient maintenance.

Risk assessment methodologies have been refined using machine learning models, allowing for a more precise evaluation of failure probabilities and severity. The use of risk matrix analysis, analytic hierarchy process (AHP), and FMECA data visualization has improved the accuracy of risk categorization. Machine learning models trained on failure history have successfully identified patterns in risk distribution, revealing that blockages in the header, speed mismatches between working parts, and cutter malfunctions pose the highest operational risks. The application of these models in predictive maintenance enables real-time fault detection, reducing downtime and optimizing repair schedules (*Kammerer et al., 2021*).

Challenges remain in implementing machine learning-based predictive maintenance at scale. The accuracy of predictions relies heavily on the quality and availability of sensor data, which can be affected by environmental conditions and sensor noise. Additionally, real-time decision-making requires low-latency processing, particularly in remote agricultural locations with limited connectivity. Advances in edge computing and federated learning are expected to address these limitations by enabling decentralized data processing, reducing dependence on cloud-based infrastructure.

Table 1 summarizes the key machine learning methods discussed in the recent studies, highlighting their specific applications in predictive maintenance and the advantages they offer in improving reliability, reducing downtime, and optimizing maintenance schedules.

Table 1

AI Methods for Predictive Maintenance

Method	Application in Predictive Maintenance	Advantages
Convolutional Neural Networks (CNNs)	Analyzes vibration and temperature data to detect faults like bearing degradation, misalignment, and lubrication failures.	Extracts both time-domain and frequency-domain features, improving predictive accuracy.

Method	Application in Predictive Maintenance	Advantages
Deep Recurrent Neural Networks (DRNNs)	Tracks component degradation over time, improving failure prediction in conveyor belts, gearboxes, and hydraulic systems.	Captures complex degradation trends that traditional methods miss.
Long Short-Term Memory (LSTM) & Gated Recurrent Units (GRUs)	Handles long-term dependency modeling, predicting remaining useful life (RUL) in time-series maintenance data.	Enables robust prediction of maintenance needs for systems requiring long-term analysis.
Stacked Autoencoders (SAEs)	Enhances feature extraction from high-dimensional sensor data, reducing reliance on manual feature selection.	Automates feature learning, enhancing predictive model efficiency.
Deep Belief Networks (DBNs)	Offers a probabilistic framework for modeling uncertainty in sensor-based failure prediction models.	Handles noisy sensor data effectively, increasing reliability of predictions.
Transfer Learning	Adapts models trained in one environment to new industrial conditions with minimal retraining.	Reduces the need for extensive labeled datasets, improving model adaptability.
Reinforcement Learning	Optimizes maintenance scheduling by dynamically learning from equipment performance data.	Continuously refines maintenance strategies based on real-time feedback.
Digital Twin Technology	Creates virtual representations of physical assets to simulate operational conditions and improve failure prediction accuracy.	Enhances predictive modeling by reducing reliance on physical testing, optimizing resource use.

Liu et al. (*2024*) provide a comprehensive overview of predictive maintenance, integrating Industry 4.0 concepts such as IoT, big data, and AI. The book highlights predictive maintenance as a key advancement over traditional methods, emphasizing real-time condition monitoring, fault diagnosis, and predictive analytics to enhance equipment performance and reliability. A major focus is the shift from conventional maintenance to AI-driven models using sensor data, digital twins, and machine learning for early fault detection. Digital twin technology enables real-time simulations and failure forecasting, while predictive algorithms optimize maintenance schedules, reducing unplanned downtime and improving spare parts management.

A case study on Vestas illustrates how SCADA systems and predictive maintenance enhance wind turbine performance through real-time data analysis. The book also explores applications in agriculture and food processing, where IoT-enabled sensors anticipate failures in farm machinery, irrigation systems, and food processing equipment. AI-driven monitoring of refrigeration units and conveyor belts ensures food safety and reduces spoilage risks.

Additionally, the book examines AI-based automation techniques, including genetic algorithms, particle swarm optimization, and neural networks, for optimizing repair schedules. Challenges such as multi-source data fusion and sensor integration are addressed, emphasizing the need for robust predictive models. The work underscores the role of Industry 4.0 in advancing predictive maintenance, making industrial systems smarter, more reliable, and cost-efficient. The Table 2 provides a summarized overview of key insights from *Liu et al. (2024)*. This summary captures the essential aspects of predictive maintenance methodologies, including its integration with Industry 4.0, applications of AI and machine learning, the role of digital twins, sensor technologies, and industry-specific implementations in agriculture and food processing. Additionally, the table highlights the primary challenges and future directions of AI-driven predictive maintenance.

Table 2

Extended Summary of Intelligent Predictive Maintenance

Topic	Key Insights
Predictive Maintenance Concept	Transition from reactive and preventive maintenance to AI-driven predictive strategies.
Integration with Industry 4.0	Use of Cyber-Physical Systems (CPS), Industrial IoT (IIoT), and big data for real-time monitoring.
AI and Machine Learning Models	Implementation of deep learning, Bayesian networks, fault tree analysis, and neural networks for failure prediction.
Digital Twin Applications	Creation of virtual replicas for failure forecasting, operational simulations, and process optimization.
Sensor Technologies	Integration of vibration sensors, acoustic emission monitors, and temperature sensors for real-time anomaly detection.

Topic	Key Insights
Applications in Agriculture (Extended Insight)	Additional insights on predictive maintenance applications in farm equipment, including IoT-enabled sensors on combine harvesters, irrigation systems, and agricultural processing equipment, reducing failures and optimizing performance.
Applications in Food Industry (Extended Insight)	Beyond the book’s scope: AI-driven predictive maintenance for monitoring refrigeration units, conveyor belts, and packaging automation, ensuring food safety, reducing spoilage, and minimizing production disruptions.
Challenges and Future Directions	Challenges in data integration, sensor reliability, and computational resource requirements; advancements in edge computing and federated learning proposed as solutions.

Bala et al. (2024) explore the integration of artificial intelligence (AI) with edge computing for predictive maintenance in industrial machines, addressing key challenges in data processing and model deployment. The study categorizes existing approaches into three main architectures: cloud-based training with edge deployment, edge-based training for enhanced data privacy, and hybrid cloud-edge training. The review highlights the advantages of edge computing, such as reduced latency and improved availability, while also acknowledging challenges like limited computational power and the need for lightweight AI models. Additionally, the paper discusses emerging research directions, including synthetic data generation, transfer learning applications, and secure communication protocols for edge AI. This work contributes to the evolving landscape of AI-driven predictive maintenance by assessing the trade-offs between centralized and distributed learning architectures.

Glock et al. (2024) introduce a novel unsupervised CPD framework, Predict and Compare (P&C), which integrates predictive machine learning models with statistical validation methods to improve anomaly detection reliability. Their approach aligns closely with ensemble-based strategies, such as the “Council of the Wise” framework employed in this study, which aggregates multiple model outputs to enhance detection accuracy. Using both deep learning models (LSTM) and statistical models (ARIMA), the Predict and Compare method successfully identifies structural shifts while mitigating false positives caused by long-term trend variations. The insights from *Glock et al.* (*2024*) reinforce the importance of combining predictive modeling with statistical validation in predictive maintenance applications. Their findings serve as a valuable reference point for enhancing multi-model anomaly detection frameworks, further supporting the scalability and adaptability of predictive maintenance strategies across diverse industrial settings.

MATERIALS AND METHODS

The predictive maintenance analysis was conducted using data collected from a Massey Ferguson 7700 S tractor, a widely used model in modern precision agriculture. Data acquisition took place in May 2023, focusing on the tractor’s hydraulic system, which plays a crucial role in various agricultural operations, including plowing, lifting, and operating auxiliary implements. Given that hydraulic systems in agricultural machinery are subject to high operational loads and varying environmental conditions, they are critical components for predictive maintenance assessment.

The Massey Ferguson 7700 S series is known for its high-performance engine, advanced transmission options, and robust hydraulic capabilities, making it an ideal candidate for predictive maintenance analysis. This model is typically equipped with a 6.6L AGCO Power engine, delivering 140 to 280 horsepower, depending on the variant. Its Dyna-VT continuously variable transmission (CVT) allows for smooth and efficient power delivery, enhancing fuel efficiency and reducing mechanical stress. However, the hydraulic system is one of the most stressed components in the tractor, as it directly influences the performance of attached implements and overall machine efficiency (https://www.masseyferguson.com/en/product/tractors/mf-7700-s.html).

The hydraulic system of the 7700 S series operates within a pressure range of 140 to 200 bar, providing the necessary force for lifting and operating various implements. This system consists of variable displacement pumps, pressure relief valves, hydraulic cylinders, and control units that regulate fluid flow and pressure. Given that hydraulic failures are often attributed to fluid contamination, wear in seals and valves, pressure fluctuations, and overheating, predictive maintenance techniques are crucial in detecting early signs of degradation before failures occur.

This study evaluates predictive maintenance strategies in agriculture using AI-driven failure prediction and sensor-based diagnostics. By integrating these methodologies, the objective was to reduce unplanned downtime, optimize maintenance schedules, and enhance the operational lifespan of the tractor’s hydraulic system, ultimately contributing to increased agricultural productivity.

The predictive maintenance framework for agricultural and food industry machinery, specifically for the Massey Ferguson 7700 S tractor's hydraulic system, operates through an integrated process of data acquisition, processing, analysis, and decision-making.

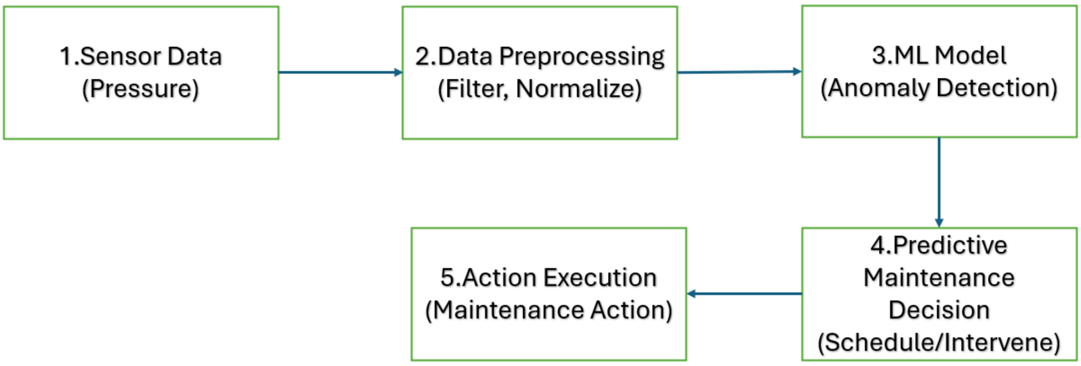

Fig. 1 - Data flow from sensor data collection through data preprocessing, ML model analysis (e.g., anomaly detection, failure prediction), maintenance decision-making, and action execution

The diagram in figure 1 follows a structured workflow that ensures early failure detection, optimizes maintenance scheduling, and improves equipment longevity.

1. Sensor Data Collection. The system begins with real-time data acquisition from multiple IoT-enabled sensors embedded in the hydraulic system of the tractor. These sensors continuously monitor key performance indicators (KPIs) such as:
    - hydraulic pressure fluctuations detect inefficiencies or valve blockages;
    - fluid temperature variations identify overheating risks or cooling inefficiencies;
    - vibration levels in hydraulic actuators and pumps indicates mechanical wear or imbalance;
    - operational workload and duty cycle variations determines stress levels on the system.
2. Data Transmission and Preprocessing. The collected sensor data is transmitted via wireless communication to an edge computing unit or cloud-based storage system. Here, the raw data undergoes preprocessing, including:
    - noise filtering – removes irrelevant fluctuations caused by environmental interference;
    - feature extraction – selects the most relevant sensor parameters for analysis;
    - data normalization – ensures consistency in input values across different operating conditions.
3. Machine Learning Models for Anomaly Detection. The processed data is then analyzed using machine learning (ML) models trained to recognize normal vs. abnormal operating patterns. Additionally, moving average and moving median filters were used as comparative statistical methods to smooth noisy sensor readings and identify abrupt fluctuations. While these are not machine learning techniques, they serve as baseline references to validate the anomaly detection models. Key AI techniques used include:
    - Time-Series Forecasting (LSTM, ARIMA – not applied in this study) to detect deviations from expected trends;
    - Isolation Forest and One-Class SVM for detecting anomalies by isolating rare and abnormal sensor readings;
    - K-Means and DBSCAN clustering to categorize different failure modes and identify outliers indicative of early-stage mechanical degradation;
    - Autoencoders for feature extraction and unsupervised anomaly detection, reconstructing normal operational states and identifying anomalies when deviations exceed a predefined threshold;
    - Convolutional Neural Networks (CNNs) to analyze spatial-temporal correlations in sensor data, particularly useful for identifying nonlinear relationships in hydraulic system degradation;
    - XGBoost for failure prediction, leveraging gradient-boosted decision trees to analyze feature importance and classify high-risk operational states.
4. Predictive Maintenance Decision Making. Based on the analysis, the system provides predictive alerts regarding potential failures before they occur. The maintenance strategy includes:
    - immediate alerts for critical failures when a component is on the verge of breakdown;
    - scheduled maintenance recommendations to address wear before it escalates;
    - spare parts inventory management ensuring the availability of necessary components based on predicted failure rates.

5. Action Execution and Maintenance Action Optimization. The system integrates with farm management software (FMS) or maintenance dashboards, allowing technicians to access real-time condition reports and historical performance logs. This enables:
   - proactive intervention rather than reactive maintenance;
   - minimized downtime by repairing components before catastrophic failure;
   - cost reduction by avoiding unnecessary replacements and optimizing resource allocation.

To evaluate the performance of the hydraulic system, real-time sensor data was collected, encompassing pressure measurements, flow rates, temperature variations, and vibration levels. These parameters were continuously tracked to detect potential anomalies associated with seal wear, pump inefficiencies, valve obstructions, or hose degradation.

For this study, output pressure was selected as the primary diagnostic metric due to its critical role in assessing the health and efficiency of hydraulic pumps. Variations in pressure readings can reveal early signs of degradation, helping maintenance teams anticipate failures before they lead to system breakdowns. Several aspects of hydraulic pump performance can be inferred by closely monitoring pressure fluctuations, making it a valuable tool for predictive maintenance. One of the most critical insights that output pressure provides is an indication of internal component wear. A gradual drop in pressure over time may signal wear in essential components such as pistons, vanes, or gears, reducing overall efficiency and leading to performance losses. Additionally, monitoring pressure allows for the early detection of cavitation, a damaging condition where vapor bubbles form due to pressure inconsistencies and collapse violently, causing internal damage to the pump. While output pressure is an effective parameter for diagnosing hydraulic pump health, relying solely on it may not provide a complete picture of the system's condition. It is best used in combination with other parameters, such as flow rate, temperature, and vibration analysis, to create a more comprehensive predictive maintenance strategy. A multi-sensor approach helps improve diagnostic accuracy, allowing for a deeper understanding of the system's overall health and facilitating early anomaly detection *(Herrera-Granados et al., 2024*).

In this stage of the research, output pressure was selected as the primary health indicator of the hydraulic pump, with data recorded under constant torque conditions. The decision to maintain a constant torque setting provided a controlled operational environment, ensuring that any fluctuations in pressure readings could be attributed to internal factors within the pump or hydraulic system, rather than variations in external load or torque. By isolating pressure as the key variable, this study aimed to enhance the reliability of pressure-based diagnostics. Any deviations detected in the output pressure were more easily linked to specific mechanical issues such as component wear, cavitation, sealing degradation, or hydraulic resistance variations. These pressure anomalies act as early warning signals, allowing maintenance teams to take preventive action before critical failures occur.

The moving average is a widely used technique that smooths data by calculating the mean value over a predefined window of time. By averaging out short-term fluctuations, it provides a clearer representation of gradual changes in system behavior. For instance, if hydraulic pressure readings are recorded every second, applying a moving average over one minute interval can reveal long-term trends while minimizing the impact of transient noise. Large deviations from the moving average, such as sudden spikes or drops, can serve as early indicators of abnormal conditions. A sustained upward drift in temperature, for example, may suggest the gradual failure of a cooling system, while a sharp pressure drop could point to a hydraulic leak or valve malfunction. The moving median, in contrast, calculates the median value within a chosen time window, making it particularly effective at handling data sets with sudden, irregular spikes or short-lived anomalies. Unlike the moving average, which smooths all fluctuations evenly, the moving median focuses on the central tendency of the data, filtering out outliers that could distort the analysis. This property makes it highly valuable in predictive maintenance scenarios where sensor noise, electrical interference, or momentary load changes can create misleading spikes in the data. By emphasizing typical values rather than being influenced by extreme variations, the moving median helps in identifying genuine trends without overreacting to temporary fluctuations.

Combining the moving average and moving median offers a more comprehensive approach to anomaly detection. While the moving average is useful for identifying long-term trends and gradual deviations, the moving median ensures that brief, isolated spikes do not lead to false alarms. When applied together, these techniques provide a balanced perspective, helping maintenance teams distinguish between short-term disturbances and sustained changes that require intervention. This dual approach enhances predictive maintenance by allowing for early detection of subtle shifts in equipment behavior, ensuring that emerging faults can be addressed proactively before they lead to failures.

By incorporating both techniques into predictive maintenance strategies, organizations can optimize their ability to monitor system performance, improve failure prediction accuracy, and reduce the risk of unexpected breakdowns. These statistical smoothing methods serve as valuable tools in data preprocessing, supporting more advanced machine learning models by ensuring that input data remains stable, interpretable, and free from misleading noise. When integrated with real-time monitoring systems, moving averages and moving medians contribute to a more robust and reliable maintenance framework, ultimately improving operational efficiency and reducing downtime across industrial and agricultural applications.

One of the most valuable findings was that the most significant anomaly patterns were not necessarily extreme values, but rather subtle deviations that did not conform to typical operational trends. These nuanced anomalies, often overlooked by traditional threshold-based methods, reinforce the importance of leveraging multiple machine learning techniques to gain a more comprehensive perspective on system health. Given that no single model can be expected to detect every relevant anomaly with absolute accuracy, a more effective approach involves a form of ensemble learning inspired by the "wisdom of the crowd" principle.

This study proposes a method for estimating anomaly detection points, referred to as the "Council of the Wise." This approach relies on an ensemble decision-making system, where an odd number of models independently analyze incoming sensor data and "vote" on whether a given point should be classified as an anomaly requiring maintenance intervention. This method ensures a more balanced and reliable assessment, reducing the likelihood of false positives and false negatives. By aggregating the decisions of multiple diverse models, this ensemble approach enhances predictive accuracy and increases confidence in maintenance recommendations. For instance, if five out of the seven models flag a particular data point as anomalous, the system classifies it as an anomaly and triggers an alert for maintenance action. If only a minority of models detect an issue, the system may either discard the anomaly or mark it as a lower-priority concern for further observation. This voting mechanism mitigates the risk associated with the limitations of any single model, creating a more resilient and adaptive predictive maintenance framework.

By combining different models, each with distinct strengths in identifying various types of anomalies – whether based on density, clustering, reconstruction errors, or decision trees – the "Council of the Wise" approach maximizes the effectiveness of predictive maintenance. It reduces dependency on any single algorithm, ensuring a more robust and interpretable decision-making process. Ultimately, this proposed method enhances anomaly detection accuracy, optimizes maintenance scheduling, reduces downtime, and extends equipment lifespan, making it a valuable contribution to predictive maintenance strategies.

The “Council of the Wise” framework, where seven machine learning models independently analyze sensor data and “vote” to determine anomalies, shares conceptual similarities with the Digital Twin paradigm while serving a distinct function within predictive maintenance (*Pulcini and Modoni, 2024*). Both approaches leverage real-time data-driven decision-making and advanced computational models to enhance predictive accuracy and optimize system performance. A Digital Twin is a virtual replica of a physical system, continuously updated with sensor data to enable real-time monitoring, fault prediction, and operational optimization. In contrast, the “Council of the Wise” operates as an ensemble anomaly detection mechanism, where multiple machine learning models process the data independently and reach a consensus on potential failures. This ensemble approach enhances reliability by reducing false positives and ensuring that only significant deviations from normal operational patterns are flagged for maintenance interventions. Integrating the “Council of the Wise” within a Digital Twin framework could enhance predictive maintenance strategies by providing an additional decision-making layer. While the Digital Twin maintains a real-time simulation of system behavior, the ensemble-based “voting” mechanism would refine anomaly detection, improving failure prediction accuracy and ensuring more confident maintenance recommendations. This hybrid methodology would combine the benefits of continuous situational awareness and high-confidence anomaly detection, leading to proactive and precise intervention strategies, minimizing downtime, and optimizing resource allocation in industrial and agricultural applications.

In the field of predictive maintenance, accurately identifying early signs of mechanical degradation is essential to minimizing unplanned downtime and optimizing maintenance schedules. One of the key challenges in this domain is distinguishing genuine system anomalies from natural operational variations, particularly in complex industrial environments where data trends are influenced by multiple factors. Change point detection (CPD) techniques play a crucial role in this process, enabling systems to identify transitions between different operational states without being misled by normal fluctuations.

To develop a robust predictive maintenance framework, this study employed eight machine learning models: Isolation Forest, One-Class SVM, KMeans, DBSCAN, Autoencoders, Convolutional Neural Networks

(CNNs), XGBoost, and Gaussian Naïve Bayes. Each model was selected based on its ability to detect anomalies in hydraulic system pressure data, leveraging different approaches such as unsupervised learning, clustering, deep learning, and probabilistic classification. The following sections provide a detailed examination of how each model was applied and its effectiveness in identifying potential failures.

Isolation Forest is a powerful and efficient tool in predictive maintenance, particularly suited for detecting anomalies in systems where deviations in sensor data can indicate early signs of equipment wear, malfunction, or inefficiencies. It is an unsupervised machine learning algorithm designed to differentiate normal operating conditions from rare, abnormal events, making it highly effective in identifying potential failures before they escalate. Unlike One-Class SVM, which constructs a boundary around normal data points, Isolation Forest isolates anomalies by assessing how quickly a given data point can be separated from the majority of the dataset. This method works on the principle that anomalies are typically sparse and differ significantly from normal data, making them easier to isolate (*Murphy, 2002; Xiao, 2022*).

Isolation Forest is based on Isolation Trees (iTrees), which recursively divide data points by selecting random features and split values. For a given dataset X with n samples, an Isolation Tree is built using the following recursive function:

$$Partition(X) = \begin{cases} X, if\ |X| = 1\ (single\ point\ left) \\ Partition(X_L) \cup Partition(X_R), otherwise \end{cases} \tag{1}$$

where X_L and X_R are the left and right subsets after a random split on a randomly chosen feature. The height $h(x)$ of a data point x in an Isolation Tree represents the number of splits required to isolate x.

The key idea behind Isolation Forest is that anomalies are easier to isolate than normal points, so they tend to have shorter average path lengths in the Isolation Trees.

$$E\big(h(x)\big) = 2H(n-1) - \frac{2(n-1)}{n} \tag{2}$$

where $H(n)$ is the harmonic number, approximated as:

$$H(n) \approx ln(n) + 0.5772156649\ (Euler - Mascheroni\ constant) \tag{3}$$

Since anomalies tend to be isolated faster, their expected path length $E(h(x))$ is smaller than that of normal points.

To quantify how anomalous a point is, Isolation Forest computes an anomaly score $s(x)$ based on the path length:

$$s(x) = 2^{-\frac{E(h(x))}{c(n)}} \tag{4}$$

where $c(n)$ is an adjustment factor for normalizing path lengths:

$$c(n) = 2H(n-1) - \frac{2(n-1)}{n} \tag{5}$$

The anomaly score $s(x)$ falls within the range [0,1]. If we have a high $s(x)$ ($s(x)$~1) is likely an anomaly (isolated quickly). If we have a low $s(x)$ ($s(x)$~0) is likely normal data (requires many splits to isolate). A typical threshold for anomaly classification is $s(x)>0.6$ or $s(x)>0.7$ based on empirical tuning.

One-Class Support Vector Machine (One-Class SVM) is a kernel-based anomaly detection model that learns a boundary around normal data and identifies points outside this boundary as anomalies. Unlike traditional SVMs used for classification, One-Class SVM is an unsupervised algorithm that operates on a single class of data (*Murphy, 2002; Xiao, 2022*).

One-Class SVM finds a hyperplane or hypersphere that best encloses normal data, separating it from potential outliers. The optimization problem is formulated as:

$$min\frac{1}{2}||w^2|| + \frac{1}{vn}\sum_{i=1}^{n}\xi_i - \rho \tag{6}$$

where w is the normal vector defining the hyperplane, $\xi_i \geq 0$ are slack variables allowing some flexibility for soft margins, ρ is the decision boundary threshold, v is the hyperparameter controlling the fraction of outliers allowed in the dataset (typically between 0 and 1), n is the number of training samples. The goal is to maximize the margin while allowing a small fraction of the training points to be considered anomalies.

Once trained, One-Class SVM uses the following decision function to classify new data points:

$$f(x) = sign(w \cdot \phi(x) - \rho) \tag{7}$$

where $\phi(x)$ is a feature transformation function (applied via a kernel), ρ is the threshold learned from the training data. A new point x is classified as normal if $f(x) \geq 0$ or anomalous if $f(x)<0$.

One-Class SVM is a widely used tool in predictive maintenance, particularly effective for detecting anomalies in systems where normal operating behavior dominates, but occasional deviations may indicate potential faults, wear, or system degradation. Its ability to model normal conditions and identify unexpected variations makes it a valuable method for early fault detection, allowing maintenance teams to intervene before minor issues develop into costly failures. In predictive maintenance applications, data is continuously collected from sensors monitoring critical parameters and other performance indicators. One-Class SVM is trained exclusively on this normal operating data, constructing a mathematical boundary that encapsulates expected values. Once trained, the model evaluates new sensor readings by determining whether they fall within or outside this learned boundary. Any data point that significantly deviates from the norm is flagged as an anomaly, signaling potential irregularities in the system.

KMeans is a widely used clustering technique that partitions data into a predefined number of clusters based on similarity. In predictive maintenance, this method is particularly effective for analyzing historical sensor data to identify typical operating states. By defining clusters that represent normal working conditions, KMeans can evaluate new data points in real-time. If a new observation falls significantly outside the boundaries of these predefined clusters, it is flagged as an outlier or anomaly. This approach is well-suited for relatively stable systems where operational patterns remain consistent (*Murphy, 2002; Xiao, 2022*).

The goal of K-Means is to minimize the total variance within clusters, defined as:

$$J = \sum_{i=1}^{k} \sum_{x \in C_i} \|x - \mu_i\|^2 \tag{8}$$

where J is the total sum of squared distances (cost function), k is the number of clusters, C_i is the set of points assigned to cluster i, x is a data point in cluster C_i, μ_i is the centroid of cluster i, $\|x-\mu_i\|^2$ represents the squared Euclidean distance between a point and its cluster center. The algorithm seeks to find the optimal centroids μ_i that minimize this cost function.

Each data point is assigned to the closest centroid using the Euclidean distance:

$$c(x) = arg\ min_i \|x - \mu_i\| \tag{9}$$

where $c(x)$ is the cluster index assigned to x. The function finds the centroid μ_i that minimizes the distance to x.

The algorithm iterates until centroids no longer change significantly or the decrease in cost function is below a threshold:

$$\mu_i^{(t+1)} - \mu_i^{(t)} < \epsilon \tag{10}$$

where t is the iteration index, ϵ is a small predefined tolerance level.

DBSCAN (Density-Based Spatial Clustering of Applications with Noise) offers a more flexible approach to anomaly detection by clustering data points based on density rather than predefining a fixed number of clusters. Unlike KMeans, which assumes that all data points should belong to a cluster, DBSCAN labels low-density points as noise or anomalies. This makes it particularly effective in handling complex sensor data where operating conditions may fluctuate, and anomalies do not necessarily conform to a fixed pattern (*Xiao, 2022*).

DBSCAN uses two important hyperparameter: ε (neighborhood radius) that defines the maximum distance within which points are considered neighbors and *min_pts* – minimum points – that defines the minimum number of points required within a neighborhood of radius ε for a point to be considered a core point.

To measure similarity between points, DBSCAN typically uses the Euclidean distance:

$$d(x_i, x_j) = \sqrt{\sum_{k=1}^{d} (x_{ik} - x_{jk})^2} \tag{11}$$

where $d(x_i, x_j)$ is the Euclidean distance between two points x_i and x_j, d is the number of dimensions, x_{ik} and x_{jk} are the k-th features of points i and j. Other distance metrics like Manhattan distance or Cosine similarity can be used depending on the application.

A point x in the dataset is classified into one of three categories:

- core point, if it has at least min_pts neighbors within distance

$$|\{x_j \in D | d(x_i, x_j) \le \varepsilon\}| \ge min\_pts \tag{12}$$

- border point, if point lies within ε-distance of a core point but has fewer than *min_pts* neighbors;
- noise (outlier), if it is not a core point and does not lie within the ε-radius of any core point.

In predictive maintenance, DBSCAN is useful for identifying irregular or infrequent operating states that might not fit within predefined categories. For example, in rotating machinery, vibration readings often exhibit natural fluctuations due to varying loads and speeds (*Laaradj et al., 2023*; *Popescu et al., 2022*). DBSCAN can distinguish between expected variations and truly anomalous behavior, such as imbalance, misalignment, or excessive friction, by detecting low-density, outlier points in the data. Since DBSCAN does not require prior knowledge of the number of clusters, it is particularly advantageous for monitoring systems with variable operating conditions, where unexpected failures may arise in non-uniform patterns.

Both KMeans and DBSCAN contribute to a proactive maintenance strategy by enabling early detection of anomalies and potential faults. KMeans is best suited for systems with predictable and well-structured operational states, where deviations from established clusters can reliably indicate anomalies. In contrast, DBSCAN excels in complex and noisy environments where system behavior is more variable, allowing for greater adaptability in detecting rare but significant failures. By combining these clustering techniques with other anomaly detection methods, maintenance teams can improve predictive maintenance accuracy, optimize intervention timing, and extend the lifespan of critical machinery.

Convolutional Neural Networks (CNNs) are becoming increasingly valuable in predictive maintenance, particularly for detecting anomalies in sensor data that exhibit spatial or temporal patterns. While CNNs are traditionally known for their success in image processing, they have also proven to be highly effective in analyzing sequential data, making them a powerful tool for identifying irregularities in time-series sensor readings. Their ability to capture local dependencies within complex datasets allows them to recognize subtle variations that may indicate early signs of equipment degradation (*Abo-Habaga et al., 2024*; *Subburaj et al., 2025*; *Yun et al., 2024*; *Zhang et al., 2024*; *Taner et al., 2024*).

The fundamental operation in CNNs is convolution, which is applied between an input feature map and a set of learnable filters (kernels). Mathematically, the convolution operation for a 2D input (e.g., sensor data) is:

$$Z(i,j) = \sum_m \sum_n X(i-m, j-n) K(m,n) \quad (13)$$

Where:

X(i,j) is input feature map (e.g., sensor data), *K(m,n)* convolution kernel (filter), *Z(i,j)* output feature map after convolution, *(m,n)* is kernel size (height and width). This convolution operation slides the kernel over the input and computes weighted sums, capturing local patterns like edges, trends, or changes in sensor data (*Murphy, 2002; Xiao, 2022)*.

In predictive maintenance applications, CNNs are trained on historical sequences of normal operating data collected from sensors monitoring key parameters. By applying convolutional layers, CNNs extract patterns from these sequences, learning what constitutes normal fluctuations under different operating conditions. Unlike traditional anomaly detection methods that rely on manually defined thresholds, CNNs autonomously learn meaningful features from raw sensor data, enhancing their ability to detect deviations that signal emerging faults. Once trained, a CNN continuously processes new sensor data, comparing it to previously learned patterns. When it identifies abnormal trends – such as unexpected pressure spikes, sudden drops, or irregular variations in vibration frequency – it flags them as potential anomalies. These deviations may indicate developing issues such as wear in mechanical components, internal leaks in hydraulic systems, or imbalances in rotating machinery. By detecting these early-stage anomalies, CNNs enable maintenance teams to investigate and address potential failures before they escalate into severe damage or system downtime.

One of the primary advantages of CNNs in predictive maintenance is their ability to handle multidimensional data, making them particularly useful in systems where multiple sensor inputs need to be analyzed simultaneously. For example, in industrial machinery equipped with an array of sensors measuring different parameters, CNNs can correlate fluctuations across these variables, identifying complex patterns that may not be apparent through single-variable analysis. This capability allows for a more comprehensive assessment of equipment health, improving failure prediction accuracy.

Autoencoders, a specialized type of neural network, are highly effective in predictive maintenance for detecting anomalies in complex and nonlinear data patterns. Their primary function is to learn a compressed representation of normal operating data and then attempt to reconstruct it as accurately as possible. Any significant differences between the reconstructed output and the original data indicate deviations from expected patterns, which may signal the onset of equipment faults or mechanical degradation (*Xiao, 2022*).

An autoencoder aims to encode input data into a lower-dimensional representation and then reconstruct it as accurately as possible. It consists of two main parts: encoder that compresses input data into a latent space representation and decoder that reconstructs the original input from the latent space (*Hajgató et al., 2022*). Given an input x, an autoencoder learns two functions:

1. encoding function $z = f_\theta(x)$, where z is the latent representation (compressed encoding) and f_θ is the encoder function (usually a neural network with parameters θ);
2. decoding function $\hat{x} = g_\phi(z)$ where $\hat{x}$ is the reconstructed input, g_ϕ is the decoder function (another neural network with parameters ϕ).

The goal of an autoencoder is to minimize the reconstruction error between x and $\hat{x}$.

Using XGBoost in combination with Isolation Forest creates a highly effective approach to anomaly detection in predictive maintenance, leveraging the strengths of both algorithms to enhance the accuracy and reliability of fault identification. This hybrid strategy takes advantage of Isolation Forest's ability to quickly isolate outliers while utilizing XGBoost's advanced pattern recognition to refine anomaly detection, reducing false positives and ensuring maintenance teams focus on critical system issues. Isolation Forest serves as the initial anomaly detection mechanism, analyzing historical sensor data – in this study, but could be also pressure, temperature, and vibration readings – to establish a model of normal operating behavior. By isolating outliers, it effectively flags deviations that may indicate potential failures. However, because Isolation Forest does not differentiate between minor fluctuations and critical anomalies, it can sometimes generate false positives, triggering unnecessary alerts. To refine these results, XGBoost is introduced as a second layer of analysis.

XGBoost optimizes a regularized loss function:

$$L = \sum_{i=1}^{N} l(y_i, \hat{y}_i) + \sum_{k=1}^{K} \Omega(f_k) \tag{14}$$

where $l(y_i, \hat{y}_i)$ is loss function (e.g., Log Loss for classification, MSE for regression), $\Omega(f_k)$ is regularization term to prevent overfitting, *K* number of trees, f_k prediction function for tree *k*.

Each XGBoost tree predicts the anomaly probability by using gradient boosting, where each new tree corrects the previous errors (*Murphy, 2002*):

$$g_i = \frac{\partial L}{\partial \hat{y}_i}, h_i = \frac{\partial^2 y}{\partial \hat{y}_i{}^2} \tag{15}$$

where g_i is gradient (first derivative of loss function), h_i is hessian (second derivative of loss function). These terms adjust tree weights for better classification of anomalies.

For each tree leaf, the new score is updated as:

$$w_j = -\frac{\sum g_i}{\sum h_i + \lambda} \tag{16}$$

where w_j is weight assigned to leaf *j*, λ is a regularization parameter.

The final anomaly score is a combination of Isolation Forest's anomaly score and XGBoost's refined classification probability:

$$S_{final}(x) = \alpha \cdot s_{IF}(x) + (1-\alpha) \cdot P_{XGB}(x) \tag{17}$$

where $s_{IF}(x)$ is Isolation Forest anomaly score, P_{XGB} is XGBoost probability of being an anomaly and α a weight parameter (controls importance of each method). A higher $S_{final}(x)$ means a higher likelihood of an anomaly.

Gaussian Naïve Bayes is a probabilistic classification model widely applied in predictive maintenance, particularly for anomaly detection and failure prediction in industrial systems. It operates based on Bayes' theorem, modeling each feature as following a Gaussian distribution (*Murphy, 2002; Xiao, 2022*).

The fundamental formula is:

$$P(C_k|X) = \frac{P(X|C_k)P(C_k)}{P(X)} \tag{18}$$

where $P(X|C_k)$ is the posterior probability of class C_k given the feature vector *X*, $P(X|C_k)$ is the likelihood, the probability of observing *X* given class C_k, $P(C_k)$ is the prior probability of class C_k, *P*(*X*) is the marginal probability of the feature vector *X*, serving as a normalizing constant.

The assumption of Gaussian distribution allows the model to estimate the likelihood of different operating conditions based on historical sensor data, making it well-suited for systems where key parameters such as temperature, pressure, and vibration fluctuate within predictable ranges. Naïve Bayes assumes conditional independence of features, meaning:

$$P(X|C_k) = \prod_{i=1}^{n} P(X_i|C_k) \tag{19}$$

The predicted class is the one that maximizes the posterior probability:

$$C^* = arg\ max_{C_k} P(C_k|X) \tag{20}$$

Despite its strengths, Gaussian Naïve Bayes has limitations that must be considered when applying it to predictive maintenance. The model assumes independence between features, which may not always hold true in complex industrial systems where multiple variables interact. This simplification can lead to misclassifications, particularly when the underlying data distribution deviates significantly from a normal Gaussian pattern. Additionally, the model does not capture intricate dependencies between parameters as effectively as more advanced machine learning methods such as tree-based models or deep learning approaches.

Each of these eight machine learning models contributed with insights into anomaly detection, leveraging different mathematical approaches to identify deviations from normal operating conditions. Interestingly, some of the models identified similar patterns of abnormal behavior, reinforcing their reliability in predictive maintenance and suggesting consistency in detecting potential failures.

An important observation in this study was that Gaussian Naïve Bayes failed to detect any anomalies, making it unsuitable for inclusion in the final anomaly “voting” process. This is most likely due to the fact that Naïve Bayes relies on well-defined probability distributions and requires a larger dataset or extended training time to accurately learn the underlying patterns in the data. Given the complexity and variability of hydraulic pressure fluctuations, its assumptions about feature independence and Gaussian distribution may not have aligned well with the dataset, leading to an inability to effectively distinguish anomalies from normal operating conditions. This outcome underscores a broader challenge in artificial intelligence: the difficulty in predicting why certain models perform effectively in specific contexts while others fail to recognize meaningful patterns. Given its inability to identify anomalies in this dataset, Gaussian Naïve Bayes was excluded from the final decision-making process, ensuring that only models capable of reliably detecting deviations contributed to the "Council of the Wise voting” framework.

The implementation of all machine learning models was carried out using Python 3.9. The experiments were conducted using widely adopted open-source libraries, including Scikit-learn, TensorFlow, Keras, XGBoost, NumPy, and Matplotlib. Model training and evaluation were performed on a workstation equipped with an Intel Core i7 processor, 32 GB of RAM, and the Ubuntu 20.04 LTS operating system. This configuration provided sufficient computational resources to train the models efficiently and ensured the reproducibility of results in a controlled and stable software environment.

The dataset used in this study was recorded during real-life plowing operations in May 2023, using a Massey Ferguson 7700 S tractor equipped with a hydraulic pressure sensing system. In order to minimize external variability and isolate equipment behavior, an approximately constant torque level was maintained throughout the entire recording period, regulated by the onboard electronic control system. The dataset includes measurements sampled at a frequency of 20 Hz and later aggregated using harmonic means calculated over one-minute intervals, providing a stable temporal structure suitable for anomaly detection models. The models based on unsupervised learning, such as Isolation Forest, One-Class SVM, K-means, and Autoencoder, were trained using data presumed to represent normal system behavior, as no labeled anomalies were available at the time of analysis. For the deep learning models (Autoencoder and 1D CNN), the training was performed exclusively on segments of data identified as non-anomalous, with model evaluation conducted by observing reconstruction errors and deviation from learned representations. No separate validation dataset was required in the traditional supervised sense, as the focus of the analysis was on outlier and anomaly detection in time-series sensor data using models that learn from normal conditions. The detection results were cross-referenced with expert knowledge and system behavior to verify the correctness of anomaly localization.

The decision to average the pressure values was based on the operational characteristics of hydraulic systems, where most mechanical degradation processes and failure precursors evolve over extended periods rather than instantaneously. High-frequency fluctuations, while relevant in some cases, can introduce unnecessary noise in anomaly detection models, potentially leading to false positives. By computing one-minute mean values, we retained essential variations indicative of progressive component wear, valve inefficiencies, or pressure instabilities while minimizing the impact of momentary perturbations caused by sensor noise or minor operational transients. Additionally, reducing the data dimensionality improved computational efficiency. Despite this averaging process, the dataset still captured meaningful pressure variations associated with the hydraulic system’s health, ensuring that predictive maintenance insights were not compromised.

RESULTS

To comprehensively assess anomalies in the hydraulic system of the Massey Ferguson 7700 S tractor, seven machine learning models were applied: Isolation Forest, One-Class SVM, KMeans, DBSCAN, Autoencoders, CNNs, and XGBoost. Each of these models independently analyzed the sensor data, identifying deviations from normal operating conditions.

Figure 2 presents a comparative analysis between Moving Average, Moving Median, and the anomalies detected by Isolation Forest in the hydraulic pressure data. The Moving Average smooths out short-term fluctuations by averaging data points within a defined window, providing a clearer trend of pressure variations. Similarly, the Moving Median filters out transient spikes and noise, making it particularly effective in handling abrupt outliers that might distort the overall trend. A significant portion of the anomalies flagged by Isolation Forest aligns with points where the Moving Average and Moving Median exhibit noticeable deviations or sudden shifts, suggesting that the system experienced abnormal fluctuations in hydraulic pressure. However, while the Moving Average and Moving Median indicate general pressure trends, they do not inherently classify anomalies, unlike Isolation Forest, which explicitly isolates unusual patterns.

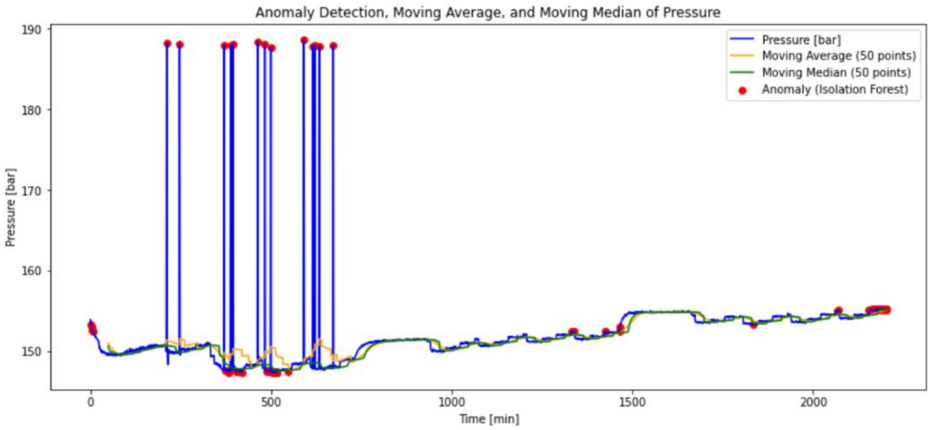

Fig. 2 – Moving Average and Moving Median compared with the anomalies detected by Isolation Forest in the hydraulic pressure data

This comparison highlights that statistical smoothing techniques like Moving Average and Moving Median can aid in visually interpreting pressure fluctuations, but they lack the ability to proactively detect anomalies with the same precision as machine learning-based models. The observed correlation between some of the Isolation Forest anomalies and deviations in the statistical trends reinforces the effectiveness of machine learning models in predictive maintenance, as they can effectively differentiate between normal system variations and potential faults.

The following figures (Figure 3 to 9) present the results of each machine learning model, showcasing their performance in detecting anomalies over time within the hydraulic pressure data. Each plot highlights pressure fluctuations and flagged anomalies, allowing for a comparative analysis of detection performance.

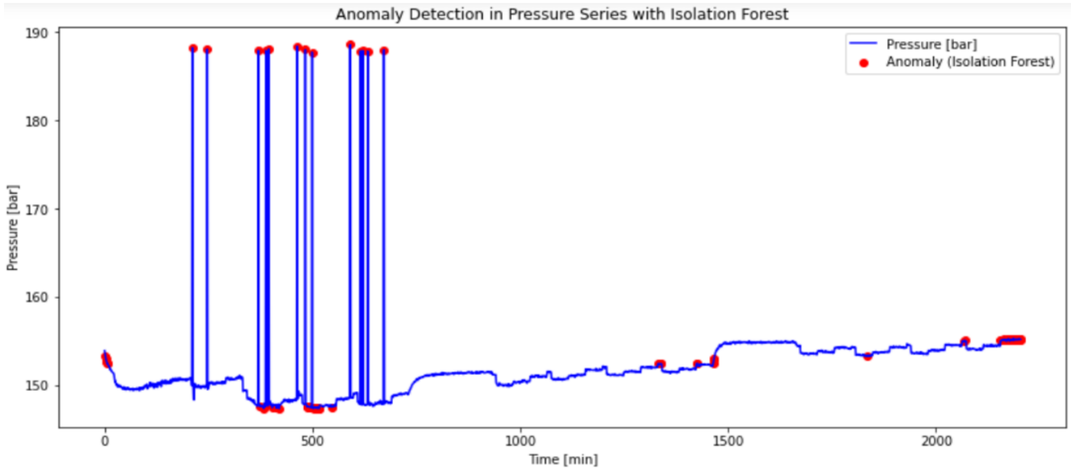

Fig. 3 – Isolation Forest results

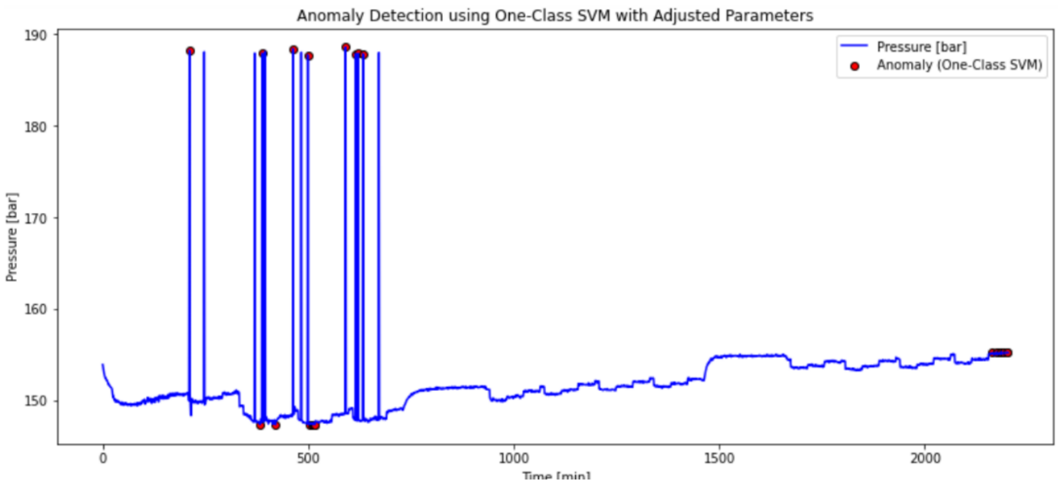

Fig. 4 – One-Class SVM results

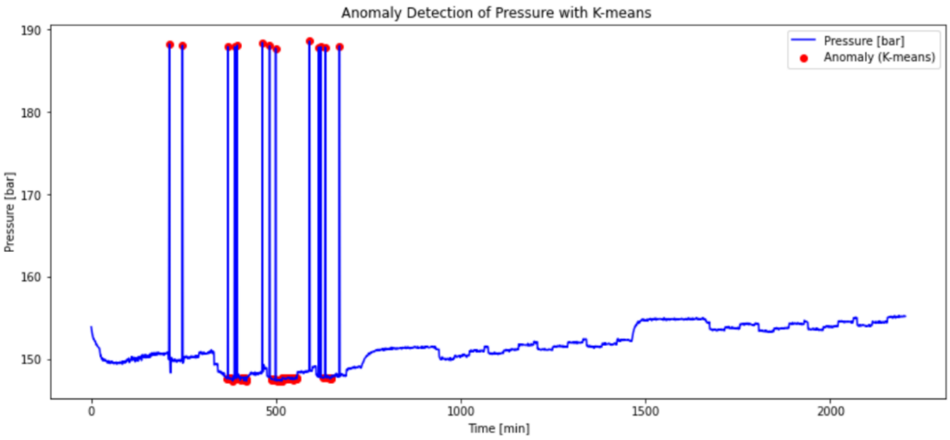

Fig. 5 – KMeans clustering results

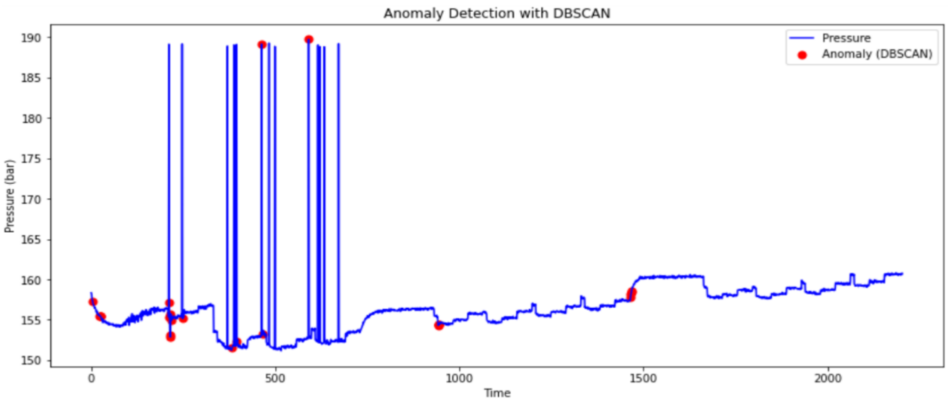

Fig. 6 – DBSCAN results

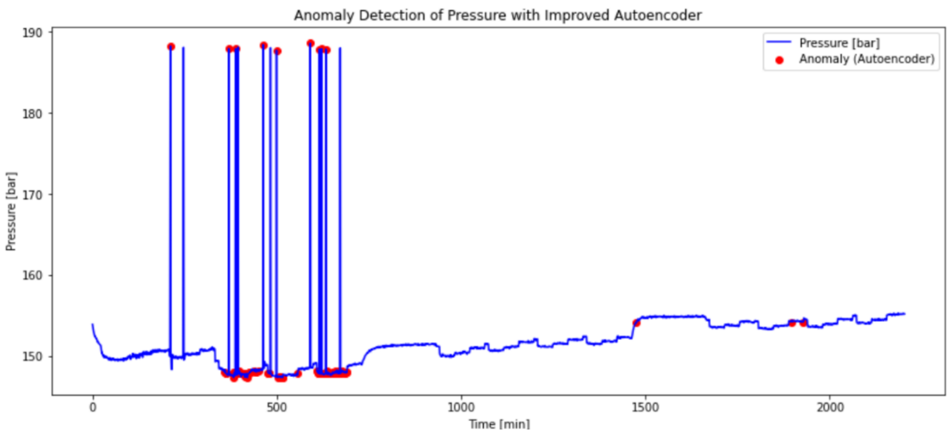

Fig. 7 – Autoencoder results

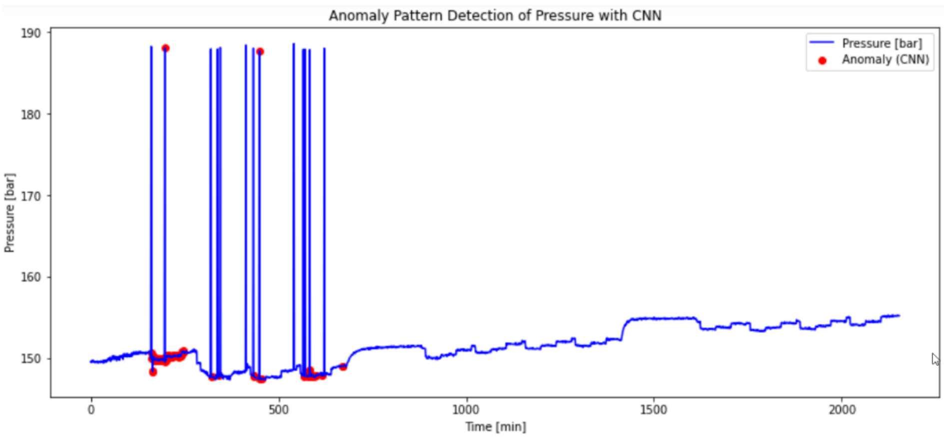

Fig. 8 – CNN results

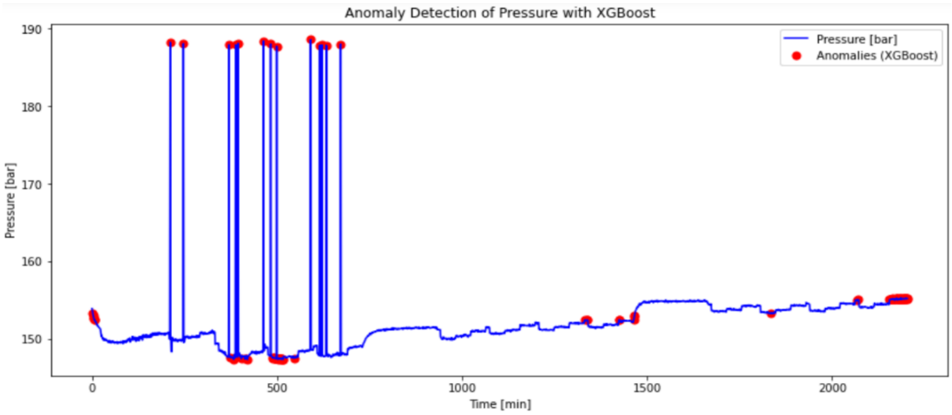

Fig. 9 – XGBoost results

Each figure provides a visual representation of how each algorithm performed in identifying anomalies, offering insights into their sensitivity, reliability, and potential limitations. The final anomaly decision (Table 3 and Figure 10), however, was determined using the "Council of the Wise" method, where an anomaly was validated only if at least four out of the seven models concurred on its presence. This consensus-based approach enhances the robustness of anomaly detection by reducing false positives and ensuring that flagged anomalies are significant deviations from normal operating conditions. By integrating these machine learning models, a robust anomaly detection framework was developed, ensuring greater accuracy in predictive maintenance. The results demonstrate that while some models are more sensitive, others provide a higher precision in detecting real faults.

Table 3 presents critical pressure anomalies identified through a consensus-based approach, where multiple machine learning models independently flagged specific points as outliers. These anomalies, identified under the specified experimental conditions, represent deviations from normal operational behavior in the hydraulic system of the Massey Ferguson 7700 S tractor. “The Council of the Wise” method, proposed in this study, combines the outputs of seven anomaly detection models (Isolation Forest, One-Class SVM, KMeans, DBSCAN, Autoencoders, CNNs, and XGBoost). Instead of relying on a single detection method, the approach ensures robustness by requiring that at least four models agree on an anomaly before it is classified as an outlier. This majority “voting” mechanism helps mitigate false positives and enhances detection accuracy.

Each entry in the table corresponds to a timestamp (Time, in minutes) and a recorded pressure value (Pressure, in bar), which has been identified as anomalous. These anomalies may indicate incipient mechanical degradation, hydraulic system inefficiencies, or potential faults, such as: valve malfunctions, leading to sudden pressure drops or spikes; cavitation or fluid instability, which can create short-term fluctuations in pressure; internal component wear, affecting system performance.

Since pressure variations in hydraulic systems are inherently dynamic, the use of this “voting”-based ensemble method significantly could improve reliability in predictive maintenance. These flagged anomalies warrant further inspection and maintenance intervention before they evolve into critical failures, thereby reducing downtime and improving system longevity.

Table 3

Final anomalies

Time (s)	Pressure (bar)
2164	151.71
2175	151.81
2186	151.81
2203	151.83
211	188.71
247	188.62
370	188.50
378	145.43
381	145.40
483	188.52
499	188.34
500	145.03
502	144.98
504	145.13
505	145.08
507	145.05
509	144.95
515	144.95
516	144.95
591	188.80
621	188.36
633	188.34
672	188.62

While this study focuses primarily on hydraulic pressure as a key indicator for predictive maintenance, additional sensor data – such as vibration, temperature, and load variation – were also collected during the experimental phase. However, to ensure clarity, reproducibility, and focus in this first phase of model evaluation, only pressure-based analysis was presented. These additional sensor signals are currently being processed and will be integrated into a multi-sensor anomaly detection framework in future work. The inclusion of multiple sensor streams is expected to enhance the robustness of the models and allow more nuanced fault detection and diagnostics.

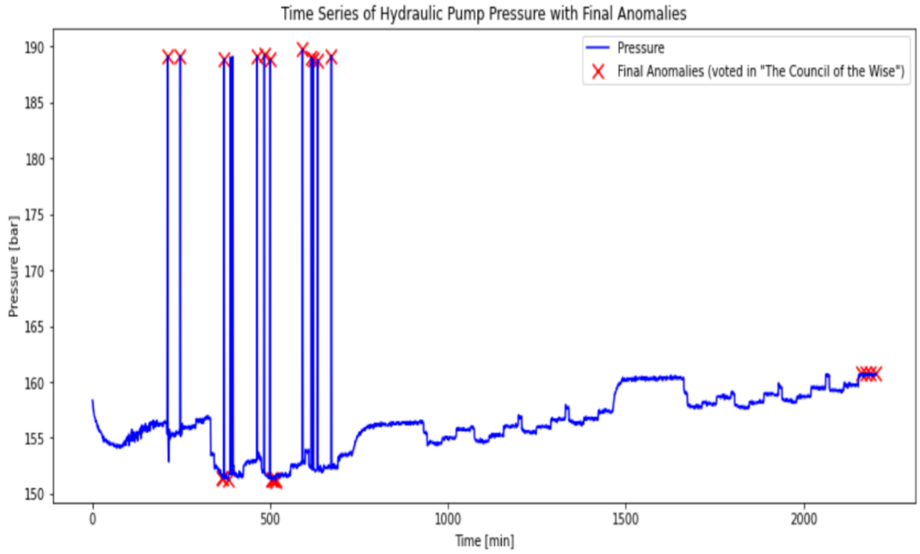

Fig. 10 – Final ”voted” anomalies

CONCLUSIONS

The final results, summarized in Table 3, confirm the effectiveness of the "Council of the Wise" approach in identifying critical anomalies in the hydraulic system of the Massey Ferguson 7700 S tractor. The detected anomalies indicate significant deviations from normal operating conditions, with pressure values fluctuating beyond expected ranges. This validates the robustness of the multi-model “voting” framework, which successfully filtered out noise, reduced false positives, and preserved the most relevant anomalies.

The temporal distribution of anomalies reveals two key insights. Clusters of anomalies occurring within short intervals suggest transient mechanical instabilities, possibly due to sudden pressure surges, valve malfunctions, or localized hydraulic inefficiencies. In contrast, isolated anomalies at different time points indicate gradual wear or intermittent system irregularities that may not be detected by traditional threshold-based monitoring. By using seven machine learning models, this ensemble-based approach improves predictive maintenance reliability by incorporating diverse detection capabilities. The ability to reduce false alarms while capturing early indicators of mechanical degradation enhances its practical applicability for agricultural machinery maintenance. While this study highlights the potential of the approach, future research could further refine adaptive sampling techniques to enhance data utilization. Retaining raw high-frequency sensor data where beneficial, alongside dynamic aggregation methods, may improve both detection accuracy and computational efficiency. Additionally, optimizing the data preprocessing pipeline could help achieve a better balance between real-time responsiveness and long-term trend analysis. This study focused primarily on hydraulic pressure as an indicator, but incorporating additional sensor data – temperature fluctuations, vibration analysis, mechanical load variations, and fuel consumption metrics – could provide a more comprehensive failure prediction framework. Expanding monitored parameters would enhance fault detection accuracy, root cause analysis, and predictive capabilities.

It is acknowledged that under typical agricultural operations, torque is not constant and may vary due to factors such as soil resistance, implement load, or maneuvering dynamics. In the current experimental setup, torque was intentionally held approximately constant in order to isolate the effects of internal system behavior on hydraulic pressure, and to validate the ability of machine learning models to detect anomalies under controlled conditions. Future stages of this research will involve data collection under variable torque conditions, representing more realistic operating environments. This will allow the models to learn and adapt to fluctuations caused by load dynamics, and assess the robustness of anomaly detection under more complex field conditions.

Overall, these findings demonstrate that the multi-model "voting" framework is a promising predictive maintenance methodology, offering a reliable, automated, and proactive solution for identifying failures before they escalate into costly breakdowns. Future research should explore its adaptability to different agricultural machinery types and operational conditions, refining the approach through advanced AI techniques and real-time deployment.

REFERENCES

[1] Abo-Habaga, M., Ismail, Z., Moustafa, N., Okasha, M., (2024), Developing an Automatic Precision Seeding Unit (APSU) for Pot Seed Planting, December 2024, *INMATEH Agricultural Engineering* 74(3). pp. 260-272, doi: https://doi.org/10.35633/inmateh-74-22.

[2] Bala, A., Rashid, R.Z.J., Ismail, I., Oliva, D., Muhammad, N., Sait, S., Al-Utaibi, K., Amosa, T.I., Memon, K.A., (2024), Artificial intelligence and edge computing for machine maintenance-review, *Artificial Intelligence Review* 57:119, doi: https://doi.org/10.1007/s10462-024-10748-9.

[3] Glock, A.-C., Sobieczky, F., Fürnkranz, J., Filzmoser, P., Jech, M., (2024), Predictive change point detection for heterogeneous data, *Neural Computing and Applications* 36, pp. 16071–16096, doi:https://doi.org/10.1007/s00521-024-09846-0.

[4] Hajgató, G., Wéber, R., Szilágyi, B., Tóthpál, B., Gyires-Tóth, B., Hős, C., (2022), PredMaX: Predictive maintenance with explainable deep convolutional autoencoders, *Advanced Engineering Informatics* 54 101778, doi: https://doi.org/10.1016/j.aei.2022.101778.

[5] Herrera-Granados, G., Misaka, T., Herwan, J., Komoto, H., Furukawa, Y., (2024), An experimental study of multi-sensor tool wear monitoring and its application to predictive maintenance, *The International Journal of Advanced Manufacturing Technology* 133, pp. 3415–3433, doi: https://doi.org/10.1007/s00170-024-13959-0.

[6] Kammerer, C., Gaust, M., Küstner, M., Starke, P., Radtke, R., Jesser, A. (2021), Motor Classification with Machine Learning Methods for Predictive Maintenance, *IFAC PapersOnLine* 54-1, pp. 1059–1064, available online at www.sciencedirect.com.

[7] Laaradj, S., Abdelkader, L., Mohamed, B., Mourad, N., (2023), Vibration-based fault diagnosis of dynamic rotating systems for real-time maintenance monitoring, *The International Journal of Advanced Manufacturing Technology* 126, pp. 3283–3296, doi: https://doi.org/10.1007/s00170-023-11320-5.

[8] Li, Z., He, Q., Li, J., (2024), A survey of deep learning-driven architecture for predictive maintenance, *Engineering Applications of Artificial Intelligence* 133, 108285, doi: https://doi.org/10.1016/j.engappai.2024.108285.

[9] Lima, A.L, da C. D., Aranha, V.M., Carvalho, C.J de L., Nascimento, E.G.S., (2021), Smart predictive maintenance for high-performance computing systems: a literature review, T*he Journal of Supercomputing* 77, pp. 13494-13513, doi: https://doi.org/10.1007/s11227-021-03811-7.

[10] Liu, M., Li, L., Yan, F., (2024), *Intelligent Predictive Maintenance*, Springer, ISBN 978-981-97-2676-9, doi: https://doi.org/10.1007/978-981-97-2677-6.

[11] Meddaoui, A., Hachmoud, A., Hain, M., (2024), Advanced ML for predictive maintenance: a case study on remaining useful life prediction and reliability enhancement, *The International Journal of Advanced Manufacturing Technology* 132, pp. 323–335, doi: https://doi.org/10.1007/s00170-024-13351-y.

[12] Murphy, K.P., (2022), *Probabilistic Machine Learning*, MIT Press, ISBN 978-0-262-04682-4, LC record https://lccn.loc.gov/2021027430.

[13] O'Neil, R, Diallo, C., Khatab, A. (2022), A Novel Predictive Selective Maintenance, Strategy Using Deep Learning and Mathematical Programming, *IFAC PapersOnLine* 55-10, pp. 1207-1212, available online at www.sciencedirect.com.

[14] Ouadah, A., Leila Zemmouchi-Ghomari, L., Salhi, N. (2022), Selecting an appropriate supervised machine learning algorithm for predictive maintenance, *The International Journal of Advanced Manufacturing Technology* 119, pp. 4277-4301, doi: https://doi.org/10.1007/s00170-021-08551-9.

[15] Popescu, T., Aiordachioaie, D., Culea-Florescu, A. (2022), Basic tools for vibration analysis with applications to predictive maintenance of rotating machines: an overview, T*he International Journal of Advanced Manufacturing Technology* 118, pp. 2883–2899, doi: https://doi.org/10.1007/s00170-021-07703-1.

[16] Pulcini, V., Modoni, G., (2024), Machine learning-based digital twin of a conveyor belt for predictive maintenance, *The International Journal of Advanced Manufacturing Technology* 133, pp. 6095-6110, doi: https://doi.org/10.1007/s00170-024-14097-3.

[17] Serradilla, O., Zugasti, E., Rodriguez, J., Zurutuza, U., (2021), Deep learning models for predictive maintenance: a survey, comparison, challenges and prospects, *Applied Intelligence* 52, pp. 10934-10964, doi: https://doi.org/10.1007/s10489-021-03004-y.

[18] Siraskar, R., Kumar, S., Patil, S., Bongale, A., Kotecha, K., (2023), Reinforcement learning for predictive maintenance: a systematic technical review, *Artificial Intelligence Review* 56, pp. 12885–12947, doi: https://doi.org/10.1007/s10462-023-10468-6.

[19] Subburaj, S.D.R., Eswaramoorthy, C., Latha, V.G., Chinnasamy, R.K.P., (2025), Efficient Pest Detection Through Advanced Machine Learning Technique, *Current Agriculture Research Journal* 12(3), pp. 1127-1134, DOI: http://dx.doi.org/10.12944/CARJ.12.3.08.

[20] Taner, A., Mengstu, M.T., Selvi, K. Ç., Duran, H., Gür, I., Ungureanu, N., (2024), *Apple Varieties Classification Using Deep Features and Machine Learning*, Agriculture, 14, 252. doi: https://doi.org/10.3390/agriculture14020252.

[21] Wang, X., Zhang, G., Yao, J., Lian, J., Yang, X., (2023), Fault Prediction Model of Corn Grain Harvester Based on Self-Coding Neural Network, *INMATEH - Agricultural Engineering* 70(2), pp. 221-231, doi: https://doi.org/10.35633/inmateh-70-22.

[22] Xiao, P., (2022), *Artificial Intelligence Programming with Python*, Wiley, ISBN 978-1-119-82086-4, LC record https://lccn.loc.gov/2022931189.

[23] Yang, X., Zhang, G., Yao, J., Lian, J., Xin Wang, X., Lv, D., Yujie Dend, Y., Zhang, A., (2022), Reliability Analysis of Grain Combine Harvesters Based on Data Mining Technology, *INMATEH - Agricultural Engineering* 67(2), pp. 211-220, doi: https://doi.org/10.35633/inmateh-67-21.

[24] Yun, Y., Ma, Z., An, X., Li, D. (2024), *On-Site Intelligent Grading System for Fruits and Vegetables: Case Studies on Tomato and Apple*, pp. 61- 85, in Yun, Y., Sheng, W., Zhang, Z., (Editors), Advanced Sensing and Robotics Technologies in Smart Agriculture, Springer, ISBN 978-981-97-6440-2, doi: https://doi.org/10.1007/978-981-97-6441-9.

[25] Zhang, C., Wang, N., Li, C., Sun, J., Jiang, Q., Han, X., Wang, J. (2024), Research on Defect Identification of Yu-Lu-Xiang Pears Based on Improved Lightweight Residual Neural Network Model, *INMATEH - Agricultural Engineering* 74(3), pp. 117-126, doi: https://doi.org/10.35633/inmateh-74-10.

[26] Zhang, W., Zhao, B., Li, G., Zhou, L., Niu, K., Gao, S., (2023), Configuration Optimization Method of Agricultural Machinery Cluster Operation Maintenance Service Vehicle, *INMATEH – Agricultural Engineering* 69(1), pp. 537-548, doi: https://doi.org/10.35633/inmateh-69-51.