

A MILP Solution of Joint Production and Maintenance Planning with a Predicted Remaining Useful Life Through Boost Method

Selin Özdemir

Industrial Engineering Department
Galatasaray University, Istanbul, Türkiye
selin.aygun@ogr.gsu.edu.tr

M. Ebru Angün

Assoc. Prof. Dr., Industrial Engineering Department
Galatasaray University, Istanbul, Türkiye
Istanbul, 34349, Turkey
eangun@gsu.edu.tr

Abstract

This study proposes an integrated production and maintenance planning framework that combines machine learning-based failure prediction with mathematical optimization. Sensor data collected from an industrial system are used to develop a machine learning model that estimates failure probabilities and generates risk indicators reflecting machine health conditions. Unlike conventional approaches that rely on deterministic assumptions, the proposed method transforms predictive outputs into probabilistic risk measures and integrates them into a mixed-integer linear programming (MILP) model. This enables dynamic and condition-based decision-making under uncertainty. In addition, machine health-dependent processing times are incorporated into the model to capture performance degradation effects. The optimization framework simultaneously determines production sequencing and maintenance decisions by balancing makespan, maintenance cost, and failure risk. To evaluate the effectiveness of the proposed approach, two benchmark strategies are considered: reactive maintenance and periodic preventive maintenance. The results show that the proposed model reduces makespan by up to 12% and total cost by approximately 30% compared to reactive maintenance, while also achieving significant improvements over periodic maintenance with a 50% reduction in maintenance interventions. The findings demonstrate that integrating machine learning outputs into optimization models provides a more efficient, adaptive, and reliable production planning framework. The proposed approach offers a practical solution for Industry 4.0 applications, where real-time data-driven decision-making is essential.

Keywords

Predictive Maintenance, Remaining Useful Life, Machine Learning, Integrated Production and Maintenance Planning, Mixed Integer Linear Programming

1. Introduction

In industrial manufacturing systems, treating production and maintenance planning as independent decision-making processes is a primary challenge that leads to unplanned downtime, increased operational costs, and significant productivity losses. Traditional reactive (run-to-failure) and time-based (preventive) maintenance strategies often fail

to account for the actual health status of machinery, resulting in either excessive maintenance or unforeseen breakdowns.

With the advent of Industry 4.0, sensor technologies and machine learning methods have enabled the estimation of a machine's Remaining Useful Life (RUL). However, the integration of RUL predictions into production planning—and the reflection of the inherent uncertainty of these predictions within decision models—remains limited in current literature. Specifically, utilizing RUL outputs as deterministic parameters fails to adequately represent risk during the decision-making process.

This study proposes a novel Mixed-Integer Linear Programming (MILP) model that integrates machine learning-based RUL predictions with production and maintenance decisions within a probabilistic framework. Unlike existing approaches, RUL outputs are not used directly; instead, they are transformed into probability distributions representing failure risk. Consequently, the probability of failure within a specific production interval is calculated and incorporated into the decision-making mechanism.

The primary contribution of this study lies in the integration of predictive maintenance outputs into production planning models within a probabilistic framework, moving beyond conventional deterministic approaches.

Numerical analyses conducted using sensor data from a real-world production system demonstrate that the proposed integrated approach reduces unplanned downtime and achieves significant improvements in total operational costs.

1.1. Objectives

The main objectives of this study are as follows:

- I. To develop a machine learning-based framework for estimating machine failure risk using sensor data.
- II. To transform RUL predictions into probabilistic measures and integrate failure risk into a mathematical optimization model.
- III. To incorporate machine health-dependent dynamic processing times into the decision-making framework.
- IV. To optimize production and maintenance decisions under uncertainty using a risk-aware MILP model.
- V. To benchmark the proposed model against deterministic and periodic maintenance strategies.
- VI. To evaluate the impact of the ML-supported approach on production performance in terms of cost, reliability, and operational efficiency.

2. Problem Definition and Literature Review

Recent studies in the field of Predictive Maintenance (PdM) demonstrate that Remaining Useful Life (RUL) and system health status can be predicted with high precision by leveraging machine learning and deep learning methodologies based on high-dimensional sensor data (Shoorkand et al., 2023, Tasci et al., 2023, Shoorkand et al., 2024, Zhuang et al., 2023, Berghout & Benbouzid, 2022, Kang et al., 2021, Wang et al., 2024, Chen et al., 2022, Ayvaz & Alpay, 2021). Deep learning-based approaches, in particular, exhibit superior performance in modeling complex and non-linear degradation processes compared to traditional model-based techniques (Berghout & Benbouzid, 2022, Kang et al., 2021, Chen et al., 2022, Ayvaz & Alpay, 2021, Fernandes et al., 2022). This trend underscores the transformative impact of big data analytics on industrial asset management.

However, a vast majority of the existing literature treats the prognosis problem independently of the operational decision-making process. RUL estimation is generally regarded as the final output of a study; the systematic integration of this information into production planning, maintenance scheduling, or sequencing decisions is addressed only by a limited number of researchers (Shoorkand et al., 2024, Zhai et al., 2022, Bousdekis et al., 2021, Fernandes et al., 2022, Fernandes et al., 2022). This disconnect between predictive capability and operational action hinders the transition of theoretical achievements into practical manufacturing environments. Consequently, it limits the utilitarian value of PdM applications and leads to the "prediction exists, but decision-making is absent" dilemma.

A review of studies integrating RUL data into maintenance decisions reveals that such integration is predominantly achieved through threshold-based or rule-based heuristics (Wang et al., 2024, Chen et al., 2022, Ayvaz & Alpay, 2021, Pater et al., 2022). While these approaches enhance the dynamism of maintenance policies, they remain insufficient for addressing multi-period and integrated planning problems. Under volatile market conditions and fluctuating production demands, static threshold values fail to maintain system optimality. Furthermore, the inherent uncertainty

in RUL predictions is either entirely overlooked in most studies or handled indirectly via safety coefficients and alarm mechanisms (Chen et al., 2022, Pater et al., 2022).

Recently, emerging research has begun to couple deep learning-based prognostic models with maintenance optimization (Zonta et al., 2022, Zhai et al., 2021, Wesendrup et al., 2023, Chen et al., 2021, Hesabi et al., 2022). These studies utilize RUL or failure probabilities as stochastic inputs for optimization models; however, they typically fail to account for production planning constraints, capacity limitations, demand fulfillment, or the multi-period interactions between production and maintenance. Maintenance optimization performed in isolation from production constraints often results in diminished delivery performance and resource wastage. Most current approaches are restricted to isolated maintenance intervals or mission-based scenarios (Wesendrup et al., 2023, Hesabi et al., 2022).

Conversely, systematic reviews conducted within the *framework* of Industry 4.0 explicitly emphasize that maintenance decisions must be managed simultaneously with production, inventory, and operational planning (Achouch et al., 2022, Zhai et al., 2022, Bousdekis et al., 2021, Fernandes et al., 2022, Fernandes et al., 2022). These findings highlight that multi-period, decision-oriented optimization models—particularly those with uncertainty awareness—constitute a significant gap in the literature. The absence of an integrated approach often causes maintenance activities to be perceived as an impediment to production flow rather than a strategic enabler.

In this context, the fundamental gap in the literature is the lack of a multi-period, mathematical optimization-based planning framework that simultaneously addresses production and maintenance decisions while incorporating probabilistic or uncertainty-aware RUL predictions. Aiming to bridge this gap, this study proposes a decision-support framework that:

- I. utilizes RUL information derived from data-driven methodologies,
- II. incorporates this information with uncertainty awareness, providing robust decision mechanisms even under high predictive variance, and
- III. optimizes production and maintenance decisions in a fully integrated manner.

3. Methodology

This study proposes an integrated approach where machine learning (ML) outputs are directly incorporated into production and maintenance planning decisions. The proposed methodology consists of three main stages: (i) data preprocessing and feature extraction, (ii) failure probability and risk score estimation, and (iii) risk-based production and maintenance optimization. Unlike traditional deterministic approaches, this framework explicitly reflects health-dependent uncertainties within the decision model.

3.1. Data Preprocessing and Feature Extraction

Industrial sensor datasets frequently contain missing observations, inconsistent timestamps, and non-informative measurements that may negatively affect machine learning performance. Therefore, a systematic preprocessing pipeline was developed prior to model construction. The preprocessing procedure consists of the following steps:

- Construction and synchronization of timestamp variables across all sensor streams
- Chronological sorting of observations according to sensor identifier and timestamp
- Removal of invalid and inconsistent timestamp records
- Missing value imputation using forward-fill and backward-fill techniques
- Elimination of constant and non-informative variables

Following these operations, the dataset was transformed into a wide-format structure where each sensor measurement corresponds to a separate feature. This representation enables efficient feature extraction and machine learning model training while preserving temporal information.

Feature engineering plays a critical role in predictive maintenance applications since raw sensor measurements may not directly reveal degradation patterns. Therefore, statistical, temporal, and vibration-based indicators were extracted from each sensor stream. Lag variables and change-based features were generated to capture temporal dependencies and degradation trends. Rolling statistics, including moving averages and standard deviations, were used to characterize local behavior and fluctuations. Furthermore, vibration-specific metrics such as RMS, kurtosis, crest factor, skewness, and peak-to-peak values were computed to represent abnormal machine conditions and early degradation signatures. These engineered features allow the machine learning model to capture both short-term anomalies and long-term degradation behavior.

Additionally, z-score normalization was applied to each sensor stream:

$$z = \frac{x - \mu}{\sigma}$$

3.2. Health Index (HI)

The Health Index (HI) was introduced to provide a single aggregated measure representing the overall degradation state of the machine. Instead of analyzing dozens of sensors signals separately, the proposed HI combines normalized deviations from all available sensors into a unified indicator. Consequently, increasing HI values indicate increasing deviation from normal operating conditions and therefore a higher degradation level.

To represent the overall degradation of the machine, the following Health Index was defined:

$$HI_t = \sum_k |z_{k,t}|$$

This index represents the sum of deviations across all sensors, allowing the machine's degradation level to be monitored via a single metric.

3.3. Definition of Failure Events

Since explicit failure labels were not available in the raw dataset, failure events were constructed using domain knowledge and operational alarm thresholds. Two critical indicators were considered: unbalance alarms and frequency deviation alarms. An observation was labeled as a failure event whenever at least one of these indicators exceeded its corresponding alarm threshold.

This approach enables the generation of meaningful target labels while preserving the operational interpretation of machine failures. Failure events are defined based on two critical sensor signals. First, specific thresholds were applied to the unbalanced signal; exceeding these thresholds triggers a system alarm. Second, frequency deviation was calculated as the absolute deviation from the signal's mean and evaluated against statistical thresholds. The failure event is defined as follows:

$$event_t = \begin{cases} 1, & \text{if any alarm occurs} \\ 0, & \text{otherwise} \end{cases}$$

This structure serves as the target variable for the machine learning model.

3.4. Failure Probability Estimation via Machine Learning

The failure prediction problem was treated as a classification task and modeled using the XGBoost algorithm. XGBoost was selected due to its ability to model nonlinear relationships, handle high-dimensional feature spaces, and provide robust performance in noisy industrial datasets. Compared with conventional ensemble methods such as Random Forest, XGBoost sequentially minimizes prediction errors through gradient boosting, allowing the model to capture complex degradation patterns more effectively.

The model estimates failure probabilities for multiple time horizons , $t + 2, \dots, t + 10$, enabling the inclusion of both short-term and mid-term risks.

3.4.1. Conditional Failure Probabilities

For each time horizon h , the following probability is defined:

$$p_{fail}^{(h)} = P(\text{breakdown at } t+h)$$

3.4.2. Threshold Exceedance Probability

The probability of a failure occurring within a specific window is calculated as: r :

$$p_{thr} = P(\text{Probability of at least one failure within a specific time window})$$

3.4.3. Risk Score

The probabilities from different horizons are weighted based on model performance to derive a composite risk score :

$$risk_score = 0.5 \cdot \sum_h w_h p_{fail}^{(h)} + 0.5 \cdot p_{thr}$$

The proposed risk score serves as an intermediary layer between machine learning predictions and optimization decisions. The first component aggregates horizon-specific failure probabilities and captures temporal degradation dynamics. The second component represents the probability of observing at least one failure within a predefined operational window. By combining these two sources of information, the risk score provides a comprehensive representation of both immediate and future operational risks. This transformation allows machine learning outputs to be directly incorporated into optimization-based decision-making.

3.5. Risk Classification

Using the risk scores, the system is categorized into three levels :

$$s_i = \begin{cases} 0, & \text{low risk} \\ 1, & \text{medium risk} \\ 2, & \text{high risk} \end{cases}$$

The continuous risk score is transformed into discrete health states to facilitate maintenance decision-making. This classification simplifies the optimization model while preserving the essential information contained in machine health indicators. Furthermore, it enables different maintenance actions to be associated with different risk levels.

3.6. Expected Processing Time

To model the impact of machine health on performance, the expected processing time is defined as :

$$E[T_{j,i}] = t_j(1 - p_{thr,i}) + t_j^{deg} \cdot p_{thr,i}$$

This expression accounts for the nominal time when the machine is healthy and the prolonged processing time when degradation is present. Traditional scheduling models assume fixed processing times regardless of machine condition. However, degraded machines often operate at reduced efficiency, resulting in increased processing durations. To capture this phenomenon, the expected processing time is modeled as a weighted combination of nominal and degraded processing times. The weighting factor is determined by the estimated failure probability. Consequently, machine health information directly affects production planning decisions, allowing the optimization model to anticipate productivity losses caused by degradation.

3.7. Integrated Production and Maintenance Model

In this study, a single-machine integrated production and maintenance planning model has been developed, utilizing risk parameters derived from machine learning outputs. The model simultaneously optimizes production sequencing and maintenance decisions. Notation of mathematical model were represented in Table 1.

Table 1. Notation of Mathematical Model

<i>Symbol</i>	<i>Description</i>
<i>I</i>	Set of production positions, indexed by <i>i</i>
<i>J</i>	Set of jobs, indexed by <i>j</i>
α_i	Degradation indicator at position <i>i</i>
r_i	Risk score at position <i>i</i>
s_i	Risk class at position <i>i</i> , where $s_i \in \{0,1,2\}$
<i>T</i>	Nominal processing time
γ	Degradation coefficient affecting processing time
p_i	Processing time at position <i>i</i>

t^{PM}	Preventive maintenance duration
c^{PM}	Preventive maintenance cost
λ	Risk penalty coefficient
x_{ji}	Binary variable equal to 1 if job j is assigned to position i , 0 otherwise
y_i	Binary variable equal to 1 if preventive maintenance is performed at position i , 0 otherwise
w_{ji}	Auxiliary binary variable for the linearization of $x_{ji}(1 - y_i)$
C_i	Completion time at position i
C_{max}	Maximum completion time (makespan)
Z	Objective function value

3.7.1. Mathematical Model

$$p_i = T + \gamma \alpha_i \quad (1)$$

$$\min Z = C_{max} + c^{PM} \sum_{i \in I} y_i + \lambda \sum_{i \in I} \sum_{j \in J: s_i=1} r_i w_{ji} \quad (2)$$

Subject to

$$\sum_{i \in I} x_{ji} = 1 \forall j \in J \quad (3)$$

$$\sum_{j \in J} x_{ji} + y_i \leq 1 \forall i \in I \quad (4)$$

$$w_{ji} \leq x_{ji} \forall j \in J, i \in I \quad (5)$$

$$w_{ji} + y_i \leq 1 \forall j \in J, i \in I \quad (6)$$

$$w_{ji} \geq x_{ji} - y_i \forall j \in J, i \in I \quad (7)$$

$$C_1 = \sum_{j \in J} x_{j1} p_1 + t^{PM} y_1 \quad (8)$$

$$C_i = C_{i-1} + \sum_{j \in J} x_{ji} p_i + t^{PM} y_i \forall i \in I, i > 1 \quad (9)$$

$$C_{max} \geq C_i \forall i \in I \quad (10)$$

$$y_i = 0 \forall i \in I: s_i = 0 \quad (11)$$

$$y_i = 1 \forall i \in I: s_i = 2 \quad (12)$$

$$x_{ji} \in \{0,1\} \forall j \in J, i \in I \quad (13)$$

$$y_i \in \{0,1\} \forall i \in I \quad (14)$$

$$w_{ji} \in \{0,1\} \forall j \in J, i \in I \quad (15)$$

The objective function in Equation (2) minimizes the overall system cost by jointly considering makespan, preventive maintenance cost, and risk cost. Equation (3) ensures that each job is assigned to exactly one position, while Equation (4) guarantees that each position can be used either for production or for maintenance. Equations (5)–(7) linearize the interaction between job assignment and maintenance decisions. Equations (8) and (9) define the completion times along the single-machine flow, and Equation (10) determines the makespan. Finally, Equations (11) and (12) impose

risk-based maintenance rules: maintenance is prohibited in low-risk states and mandatory in high-risk states, whereas the medium-risk state is left to the optimization model.

4. Data Collection

The dataset used in this study consists of real-world time-series sensor data collected from an industrial production system operating under continuous conditions. Multiple sensors are installed on the machine to monitor its operational state, including vibration (unbalance) and frequency measurements.

Sensor data from different sources and sheets are merged into a unified structure based on timestamp alignment to ensure consistency across all measurements .

4.1. Dataset Characteristics

The dataset used in this study consists of multivariate time-series sensor measurements collected from an industrial machine operating under real production conditions. Following preprocessing and feature engineering procedures, the resulting dataset contains 58,236 observations representing the machine's operational behavior over time. Initially, 359 features were generated from raw sensor measurements. To improve model robustness and reduce redundancy, a variance filtering procedure was applied, eliminating low-information variables and reducing the feature space to 322 features. Subsequently, correlation-based filtering was performed to remove highly correlated variables and mitigate multicollinearity effects. After this process, the final dataset consisted of 206 features. The resulting high-dimensional feature space captures both short-term operational variations and long-term degradation behavior, providing a comprehensive representation of machine health dynamics. Detailed characteristics of dataset were represented in Table 2.

Table 2. Dataset Characteristics.

Metric	Value
Total Observations	58,236
Initial Features	359
Features After Variance Filtering	322
Features After Correlation Filtering	205
Final Feature Count	206

4.2. Event Definition and Statistics

Failure events are not directly available in the dataset and are constructed using domain knowledge and signal thresholds .

Two main indicators are used:

4.2.1. Unbalance signal

- Warning threshold: 129 occurrences
- Alarm threshold: 73 occurrences
- Event count: 12

4.2.2. Frequency deviation

- Warning events: 2,913
- Alarm events: 2,912
- Event count: 416

4.2.3. Combined failure events

A combined failure event is defined as:

$$event_t = \begin{cases} 1, & \text{if any alarm occurs} \\ 0, & \text{otherwise} \end{cases}$$

Total number of combined events: 428 failure events

4.3. Data Visualization

Figure 1 illustrates the frequency deviation signal over time along with the detected warning and alarm thresholds. It can be observed that the signal exhibits increasing variability during certain periods, indicating potential degradation phases.




Figure 1. Frequency Deviation and Detected Alarm Events of Processing Data.

Figure 2 presents the health index together with the detected combined alarm events. The results show that higher health index values are strongly associated with failure events, validating the effectiveness of the constructed features.




Figure 2. Health Index with Combined Alarm Events of Processing Data

5. Results and Discussion

5.1 Numerical Results

To evaluate the effectiveness and solution capacity of the proposed machine learning (ML)-driven optimization model, three fundamental benchmark scenarios were considered:

- S1: Reactive Maintenance Strategy (A "run-to-failure" approach where an unplanned downtime occurs every 3 jobs).
- S2: Periodic Preventive Maintenance Strategy (A conventional approach where maintenance is performed every 3 jobs, regardless of the machine's actual condition).
- S3: Proposed ML-Based Adaptive Model (An integrated approach dynamically shaped by real-time risk scores).

The comparative performance results are summarized in Table 3.

Table 3. Performance comparison of different maintenance strategies

Model	Makespan (Cmax)	Total Cost	Number of Maintenance Actions
Reactive Maintenance (S1)	256,000	416,000	8 (CM - Corrective)
Periodic Maintenance (S2)	240,000	320,000	8 (PM - Preventive)
ML-Based Model (S3)	225,751	291,699	4 (PM - Preventive)

5.2. Performance Improvement Analysis

The empirical results clearly demonstrate that the proposed ML-based model significantly outperforms both traditional strategies from both statistical and operational perspectives.

5.2.1. Compared to Reactive Maintenance (S1):

- Makespan was reduced from 256,000 to 225,751 (~12% improvement).
- Total cost decreased from 416,000 to 291,699 (~30% reduction). These data explicitly highlight the operational blindness of reactive strategies and the high-cost disruptions caused by unplanned downtimes. Notably, the 30% cost advantage underscores the system's success in managing unexpected failure risks.

5.2.2. Compared to Periodic Maintenance (S2):

- Makespan was reduced from 240,000 to 225,751 (~6% improvement).
- Total cost decreased from 320,000 to 291,699 (~9% reduction).
- Maintenance actions were reduced from 8 to 4 (50% reduction). These findings prove that periodic maintenance, by ignoring machine health, leads to redundant interventions and halts the system without providing proportional benefits. The 50% decrease in maintenance frequency indicates a dramatic increase in resource utilization efficiency.

5.3. Maintenance Strategy Comparison

The three strategies exhibit fundamentally different behaviors in terms of system management:

1. Reactive maintenance (S1) maintains the lowest efficiency due to unplanned failures and high-cost corrective interventions.
2. Periodic maintenance (S2) improves system reliability but applies maintenance according to a fixed rule, disregarding the actual degradation state of the machine.
3. ML-based strategy (S3) dynamically schedules maintenance based on real-time risk signals derived from sensor data, thereby optimizing the balance between cost and reliability.

In the proposed model, maintenance actions were focused only on critical positions (13, 18, 21, and 25) identified by ML predictions, maximizing the operational added value of each intervention.

5.4. Impact on Production Efficiency

The systematic integration of machine health information into the decision-making process led to a significant smoothing of the production flow:

- The reactive strategy frequently interrupts the flow due to failures.
- The periodic strategy forces the system into unnecessary downtimes due to fixed intervals.
- The ML-based strategy optimizes flow continuity by minimizing both failure-induced and redundant interruptions. Consequently, the model offers a smoother and more realistic production schedule with improved time management.

5.5. Discussion

The findings clearly expose the inadequacy of traditional maintenance approaches within the Industry 4.0 ecosystem. The high operational costs of reactive maintenance and the lack of flexibility in periodic maintenance emphasize the increasing importance of integrated models.

The ML-driven approach proposed in this study provides a data-driven decision-support mechanism where maintenance is performed only when economically and technically rational. This integration offers two primary strategic advantages:

1. Risk Awareness: Uncertainty and degradation dynamics are successfully reflected in the mathematical model through probabilistic measures.
2. Adaptive Planning: The system gains a flexibility independent of static rules, allowing for instantaneous responses to changes in the machine's actual condition.

Specifically, the achieved savings suggest that implementing this model in high-volume production facilities can create a scalable and sustainable competitive advantage. Overall, the proposed framework introduces a far more efficient perspective to production and maintenance planning, aligning closely with real-world shop floor dynamics.

6. Conclusion and Key Contributions

This study proposes an integrated production and maintenance planning framework that synergizes machine learning-based failure forecasting with mathematical optimization. The following points summarize how the research objectives were achieved:

Primarily, a robust machine learning-driven framework was developed to quantify machine failure risks using high-dimensional sensor data, successfully fulfilling the first objective. The model effectively transforms latent raw signals into interpretable degradation indicators and probabilistic risk measures.

Secondly, rather than relying on deterministic Remaining Useful Life (RUL) estimates, this research converts failure predictions into stochastic risk metrics and integrates them into a Mixed-Integer Linear Programming (MILP) model. This directly addresses the second objective by enabling risk-aware decision-making within a rigorous optimization framework.

Thirdly, the proposed approach incorporates dynamic processing times that are functionally dependent on machine health, allowing the model to encapsulate performance degradation effects. This satisfies the third objective and ensures a high-fidelity representation of real-world production systems.

Fourthly, production sequencing and maintenance scheduling are optimized simultaneously under uncertainty through a risk-informed MILP formulation. The empirical results confirm that the model effectively balances operational throughput with system reliability, thereby attaining the fourth objective.

Fifthly, the efficacy of the proposed model was rigorously benchmarked against conventional reactive and periodic maintenance strategies. The comparative analysis reveals significant improvements across all key performance indicators (KPIs), fulfilling the fifth objective.

Finally, the impact of the ML-supported paradigm was evaluated in terms of total cost, makespan, and maintenance frequency. The findings demonstrate a substantial reduction in both operational expenditures and maintenance interventions, while simultaneously enhancing efficiency, thus addressing the sixth objective.

In conclusion, the results substantiate that the integration of machine learning outputs into optimization models facilitates a more resilient, adaptive, and data-driven production planning environment. The proposed framework offers a practical and scalable solution for Industry 4.0 ecosystems, where the synthesis of real-time data and intelligent decision-making is paramount for competitive advantage.

References

Achouch, M., Dimitrova, M., Ziane, K., Karganroudi, S. S., Dhoubib, R., Ibrahim, H., and Adda, M., On Predictive Maintenance in Industry 4.0: Overview, Models, and Challenges, *Applied Sciences*, vol. 12, p. 8081, 2022.

- Ayvaz, S. and Alpay, K., Predictive maintenance system for production lines in manufacturing: A machine learning approach using IoT data in real-time, *Expert Systems with Applications*, vol. 173, p. 114598, 2021.
- Berghout, T. and Benbouzid, M., A Systematic Guide for Predicting Remaining Useful Life with Machine Learning, *Electronics*, vol. 11, p. 1125, 2022.
- Bousdekis, A., Lepenioti, K., Apostolou, D., and Mentzas, G., A Review of Data-Driven Decision-Making Methods for Industry 4.0 Maintenance Applications, *Electronics*, vol. 10, p. 828, 2021.
- Chen, C., Shi, J., Lu, N., Zhu, Z. H., and Jiang, B., Data-driven predictive maintenance strategy considering the uncertainty in remaining useful life prediction, *Neurocomputing*, vol. 494, pp. 79–88, 2022.
- Chen, C., Zhu, Z. H., Shi, J., Lu, N., and Jiang, B., Dynamic Predictive Maintenance Scheduling Using Deep Learning Ensemble for System Health Prognostics, *IEEE Sensors Journal*, vol. 21, no. 23, 2021.
- Fernandes, M., Corchado, J. M., and Marreiros, G., Machine learning techniques applied to mechanical fault diagnosis and fault prognosis in the context of real industrial manufacturing use-cases: a systematic literature review, *Applied Intelligence*, vol. 52, pp. 14246–14280, 2022.
- Hesabi, H., Nourelfath, M., and Hajji, A., A deep learning predictive model for selective maintenance optimization, *Reliability Engineering and System Safety*, no. 219, 2022.
- Kang, Z., Catal, C., and Tekinerdogan, B., Remaining Useful Life (RUL) Prediction of Equipment in Production Lines Using Artificial Neural Networks, *Sensors*, vol. 21, p. 932, 2021.
- Pater, I., Reijns, A., and Mitici, M., Alarm-based predictive maintenance scheduling for aircraft engines with imperfect Remaining Useful Life prognostics, *Reliability Engineering and System Safety*, no. 221, 2022.
- Shoorkand, H. D., Nourelfath, M., and Hajji, A., A deep learning approach for integrated production planning and predictive maintenance, *International Journal of Production Research*, vol. 61, no. 23, pp. 7972–7991, 2023.
- Shoorkand, H. D., Nourelfath, M., and Hajji, A., A hybrid CNN-LSTM model for joint optimization of production and imperfect predictive maintenance planning, *Reliability Engineering and System Safety*, no. 241, 2024.
- Shoorkand, H. D., Nourelfath, M., and Hajji, A., A hybrid deep learning approach to integrate predictive maintenance and production planning for multi-state systems, *Journal of Manufacturing Systems*, pp. 397–410, 2024.
- Taşcı, B., Omar, A., and Ayvaz, S., Remaining useful lifetime prediction for predictive maintenance in manufacturing, *Computers & Industrial Engineering*, vol. 184, 2023.
- Wang, L., Chen, Y., Zhao, X., and Xiang, J., Predictive Maintenance Scheduling for Aircraft Engines Based on Remaining Useful Life Prediction, *IEEE Internet of Things Journal*, vol. 11, no. 13, 2024.
- Wesendrup, K. and Hellingrath, B., ML2-enabled Condition-based Demand, Production, Inventory, and Maintenance Planning, *IFAC-PapersOnLine*, vol. 56, no. 2, pp. 6600–6605, 2023.
- Zhai, S., Gehring, B., and Reinhart, G., Enabling predictive maintenance integrated production scheduling by operation-specific health prognostics with generative deep learning, *Journal of Manufacturing Systems*, no. 61, pp. 830–855, 2021.
- Zhai, S., Kandemir, M. G. and Reinhart, G., Predictive maintenance integrated production scheduling by applying deep generative prognostics models: approach, formulation and solution, *Production Engineering*, vol. 16, pp. 65–88, 2022.
- Zhuang, L., Xu, A., and Wang, X., A prognostic driven predictive maintenance framework based on Bayesian deep learning, *Reliability Engineering and System Safety*, no. 234, 2023.
- Zonta, T., Costa, C. A., Zeiser, F. A., Ramos, G. O., Kunst, R., and Righi, R. R., A predictive maintenance model for optimizing production schedule using deep neural networks, *Journal of Manufacturing Systems*, no. 62, pp. 450–462, 2022.

Biographies

Selin Özdemir is a student in the Master Program of Logistics and Financing Management/ Industrial Engineering Dept. in Galatasaray University. Bachelor's degree in Industrial Engineering, Industrial Engineering Department, University of Turkish Aeronautical Association. She is currently production planning engineer at Leading Defense Industry. Her research interests are operation research, simulation, optimization and machine learning.

Ebru Angün is Associate Professor of Operations Research in the Department of Industrial Engineering at Galatasaray University in Istanbul. Her research interests include simulation optimization, stochastic optimization, and applications in disaster management. She has been an author and reviewer for many international journals, and a project reviewer for the National Science Foundation of Turkey and the Research Foundation Flanders (FWO) in Belgium. She spent one year at the Georgia Institute of Technology, as a visiting scholar. Also see <https://avesis.gsu.edu.tr/eangun> and <https://www.researchgate.net/profile/Ebru-Angun>.