Uluslararası İleri Doğa Bilimleri ve Mühendislik Araştırmaları Dergisi
Sayı 9, S. 201-214, 2, 2025
© Telif hakkı IJANSER'e aittir
Araştırma Makalesi

IJANSER

https://as-proceeding.com/index.php/ijanser
ISSN:2980-0811

International Journal of Advanced Natural Sciences and Engineering Researches
Volume 9, pp. 201-214, 2, 2025
Copyright © 2025 IJANSER
Research Article

Deep Learning for Predictive Maintenance: Optimizing Dynamic Time-Dependent Data Streams with Cost Function Analysis

Louaye Mandari*, Mehmet Savsar [2]

[1]*Industrial Engineering Department Uskudar university,Türkiye*
[2]*Industrial Engineering Department Uskudar university,Türkiye*

**Louaye.mandari@st.uskudar.edu.tr*

(2nd International Conference on Pioneer and Innovative Studies ICPIS 2025, January 30-31, 2025)

ATIF/REFERENCE: Mandari, L. & Savsar, M. (2025). Deep Learning for Predictive Maintenance: Optimizing Dynamic Time-Dependent Data Streams with Cost Function Analysis. *International Journal of Advanced Natural Sciences and Engineering Researches*, 9(2), 201-214.

Abstract – This thesis delves into the transformative role of deep learning techniques in predictive maintenance, with a focused investigation on the cost functions used in the evaluation and optimization of predictive models. Both linear and nonlinear forms of the cost function are explored to enhance the performance and precision of machine learning models in predictive maintenance scenarios. The study demonstrates how these cost functions can be tailored to effectively predict equipment failure, whether through binary classification for failure detection or more complex multi-class classification tasks. The research underscores the importance of cost function selection in balancing accuracy and computational efficiency, offering practical insights for industries reliant on continuous operations. By improving early detection of failures, this work aims to minimize downtime and prolong the operational life of machinery, ultimately reducing maintenance costs and increasing overall system reliability.

Keywords – Predictive Maintenance, Deep Learning, Machine learning ,Time series analysis, Classification

I. INTRODUCTION

In the industrial landscape, where operational efficiency and sustainability are critical, predictive maintenance technology has emerged as a transformative solution. This innovative approach harnesses the power of advanced data analytics, machine learning, and IoT to preempt equipment failures, ensuring seamless production and operational longevity. Unlike traditional reactive or preventative maintenance, predictive strategies revolutionize asset management by enabling real-time monitoring and preemptive action based on machine health insights.

By leveraging sensor data and AI algorithms, manufacturers can predict and prevent breakdowns, reducing downtime and operational disruptions. As noted by Oleg Fonarov, founder of Program-Ace, predictive maintenance facilitated by digital twins can increase equipment uptime by 10–20% and decrease maintenance costs by 5–10% (Fonarov) , while enhancing overall efficiency and sustainability .

Moreover, the advent of Industry 4.0 has amplified the potential of predictive maintenance. Sophisticated analytics derived from sensor data not only improve safety and productivity but also align industrial operations with sustainability goals. Predictive strategies significantly reduce waste and environmental impact, creating a more resilient and adaptive industrial ecosystem. As a pivotal aspect of modern industrial innovation, predictive maintenance exemplifies the shift toward a data-driven, proactive approach, fostering long-term economic and ecological benefits.The significance of predictive maintenance extends beyond efficiency and cost reduction. Its applications in diverse industries, from aerospace to water management, highlight its versatility and transformative potential. For instance, Rolls-Royce (Olavsrud) leveraged digital twins for jet engine monitoring, achieving a 50% increase in maintenance intervals while reducing carbon emissions . Similarly, AI-driven predictive maintenance strategies can slash downtime by up to 50% and reduce maintenance costs by 10%, demonstrating its broad applicability and profound impact on industrial operations

This research delves into the principles, advancements, and applications of predictive maintenance technology, aiming to explore its transformative role in reshaping the future of industrial maintenance and operational efficiency.

II. MATERIALS AND METHOD

Precision and Recall : A Comprehensive Guide to Model Evaluation Metrics

In the field of machine learning, especially in classification tasks, evaluating the performance of a model is crucial to ensure that it generalizes well and produces reliable predictions. Common metrics used for evaluating classification models include Precision and Recall. These metrics are particularly important in cases where imbalanced data or high consequences of misclassification exist. In this extended discussion, we'll explore the definitions, importance, and applications of each metric, and also discuss why they are critical to assessing model performance.

1. Precision

Definition:
Precision is a metric used to measure the accuracy of the positive predictions made by a model. It is calculated as the ratio of correctly predicted positive observations to the total predicted positives. In simple terms, it answers the question: *Of all the instances that the model predicted as positive, how many were actually positive?*

$$Precision = (\,True\ Positives\ (TP))/\ False\ Positives\ (FP) + True\ Positives\ (TP)$$

Importance:

Precision is particularly useful when the cost of false positives is high. False positives occur when a model incorrectly predicts a positive outcome for a negative case. In many applications, this kind of error can have significant negative consequences.

For example, consider a fraud detection system used by a bank. A false positive in this context would mean incorrectly flagging a legitimate transaction as fraudulent. While such an error does not mean a crime was missed, it could cause unnecessary inconvenience to the customer and waste resources on investigating a legitimate transaction. Therefore, in such cases, precision is critical. A model with high precision ensures that when it predicts a transaction to be fraudulent, it is likely to be correct.

Another example could be in the context of email spam detection. If the model incorrectly flags a legitimate email as spam, it might result in important communications being missed. Precision in this case helps in minimizing such errors.

2. Recall

Definition:

Recall, also known as Sensitivity or True Positive Rate, measures the ability of a model to correctly identify all actual positive instances. It is the ratio of correctly predicted positive observations to the total actual positives. In simpler terms, it answers the question: *Of all the actual positives, how many did the model correctly identify?*

$$Recall = (\,True\ Positives\ (TP))/\ False\ Negatives(FN) + True\ Positives\ (TP)$$

Importance:

Recall is important when the cost of false negatives is high. A false negative occurs when a model incorrectly predicts a negative outcome for a positive case. This error can be critical, particularly in fields such as healthcare, where missing a disease diagnosis can be life-threatening, or in security, where failing to detect an intruder can have serious consequences.

For example, in medical diagnostics, imagine a cancer detection model. A false negative would mean the model fails to identify a patient who actually has cancer. This could lead to the patient not receiving timely treatment, which could worsen their condition. In such cases, recall becomes crucial because identifying as many positive cases as possible is a matter of life and death.

Similarly, in the context of a spam filter, recall would be important if the goal is to catch as many spam emails as possible. Missing out on a spam email would mean that unwanted content still reaches the user's inbox.

Why Precision and Recall Matter

The evaluation of a classification model goes beyond just measuring how accurate it is. Precision and recall provide insights into different types of misclassifications and allow you to assess how well the model handles both false positives and false negatives.

Precision: Focus on False Positives

Precision specifically addresses the problem of false positives. It is particularly important when the consequences of falsely identifying a positive case are severe, leading to wasted resources, false alarms, or customer dissatisfaction. By maximizing precision, you reduce the chance of making these costly errors.

Recall: Focus on False Negatives

Recall focuses on false negatives and is most valuable when the cost of missing a true positive is high. In situations like medical diagnoses or fraud detection, failing to identify positive cases can have severe implications. Therefore, maximizing recall ensures that as many true positives as possible are identified, even at the risk of allowing some false positives.

Cost Analysis Using Performance Metrics

A unique aspect of this study is the incorporation of a cost equation to quantify the economic implications of model performance. By integrating recall and precision values into a cost function, the study evaluates the trade-offs between detection rates, false alarms, and the financial investment required per period. Specific values for costs and investments are defined, providing a clear framework for assessing the economic impact of the model's predictions.

The cost equation used in this study is extended to include the summation of all identified failures, accounting for their respective costs and the associated economic consequences. This holistic approach ensures that the analysis captures the broader financial implications of the model's performance, offering valuable insights for decision-making in industrial operations.

The cost Function

The cost function provided is a comprehensive model that calculates the total cost incurred by a machine or system in relation to its failures, detection accuracy, and system maintenance. It takes into account various factors, such as the number of failures detected, the cost associated with fixing these failures, the precision of the machine learning model, and the cost of failures that go undetected. Let's break this cost function down in detail and explore each of its components.

The cost function is given by the equation:

$$Cost = F_d * C_d * P_r + F_d * C_d * (1 + K) * (1 - R_c) + I$$

F_d is the number of failures detected by the machine per period

C_d is the cost of fixing the component or machine of detectable failure

P_r is the cost of precision of the model

k is the percentage of difference between the cost of fixing a detectable failure and undetectable failure

R_c is the recall of the machine learning model

k is the percentage value of how much is the undetectable failure cost is higher than the detectable failure cost

The Relationship Between Detectable and Undetectable Failures

The first key distinction in the cost function is between detectable and undetectable failures. Detectable failures are those that the system is able to identify with its sensors and algorithms. These failures are relatively easier to fix since the system has provided an early warning, allowing for preventive measures to be taken. The cost associated with fixing these failures is represented by C_d, and this is multiplied by F_d, the number of failures detected, and P_r, the precision of the model. The lower the precision, the higher the potential for misclassification, which could lead to additional costs. However, the essential characteristic of detectable failures is that they can be mitigated before they escalate into more severe issues, which is why the cost associated with them, C_d, is generally lower than the cost of fixing undetectable failures.

Undetectable failures pose a significant challenge in various industries, particularly in high-risk environments such as aerospace and manufacturing. These failures are not immediately detectable by conventional monitoring systems, making them particularly dangerous because they can cause major disruptions to operations, production, or even catastrophic events. Unlike detectable failures, which are identified and rectified before they escalate, undetectable failures often go unnoticed until they result in unplanned downtimes, operational halts, and costly repairs.

For example, in the aerospace industry, undetectable failures can occur in critical components like turbine engines or hydraulic systems, which are designed to function under extreme conditions. These failures are often not visible or audible to sensors during regular maintenance checks, and they may only become apparent after causing significant damage. For instance, a slight crack in a turbine blade, which remains undetected during routine inspections, could lead to engine failure during flight, potentially causing catastrophic consequences. The repair cost of such undetectable failures often surpasses the cost of easily detectable failures, such as a malfunctioning sensor or an unbalanced rotor.

To quantify this difference in cost, let's assume that fixing a detectable failure in an aerospace engine, such as a malfunctioning sensor, costs C_d. However, if an undetectable failure occurs, such as a microscopic crack in a critical component that goes unnoticed, the repair cost may be C_d×(1+k), where k represents the percentage by which the undetectable failure repair cost exceeds the cost of fixing a detectable failure. The value of k could vary depending on the severity of the failure and the complexity of the repair, but it typically ranges between 30% to 100%, or even more in extreme cases where extensive overhaul is required.

Example:

Consider a scenario in an aerospace setting where a turbine blade develops a minor crack that is not detected during routine maintenance. While a detectable issue, such as a broken sensor, might cost around $50,000 to fix, the undetected crack in the turbine could result in a complete engine overhaul costing $150,000, which includes unplanned downtimes, extensive diagnostics, and the replacement of the damaged parts. In this case, k = 2, meaning the cost of fixing the undetectable failure is 200% higher than the detectable failure.

Undetectable failures not only increase repair costs but also extend downtime, leading to missed production schedules, especially in high-stakes industries like aerospace. These failures require more time and resources to identify, typically involving more sophisticated diagnostic tools and procedures. Moreover, undetected issues can cause a cascading effect on other systems, leading to further failures and increased operational risks.

In industries such as aerospace, early detection of failures is vital for minimizing these costs. Technologies like predictive maintenance, utilizing sensors, machine learning algorithms, and real-time monitoring systems, are increasingly employed to detect signs of undetectable failures before they escalate into catastrophic events. However, the challenge remains to effectively reduce the occurrence of undetectable failures, particularly in complex systems where numerous variables are at play.

Precision and Recall in the Model

The model also incorporates two important metrics used to evaluate the performance of failure detection systems: precision and recall.

- Precision (P_r) measures how accurate the model is when it identifies a failure. A high precision means that when the system predicts a failure, it is likely to be correct. This is important because, with high precision, we incur fewer costs in fixing failures that were wrongly identified as problematic. Precision is a direct cost factor in the equation since it impacts the total cost of dealing with detectable failures.
- Recall (R_c), on the other hand, measures how well the system identifies actual failures. A high recall means that most of the actual failures are detected by the system. If recall is low, it indicates that many failures are going undetected, leading to higher costs for undetectable failures. The recall value is used in the cost function to adjust the cost of fixing undetectable failures. Specifically, the cost of undetectable failures is increased when recall is lower, as the system misses more failures, thereby incurring higher costs.

The logic behind precision and recall in this equation is rooted in the performance of the detection system. A high precision means that the system is more accurate when it predicts failures, reducing unnecessary fixes and associated costs. On the other hand, a high recall means that fewer failures go undetected, minimizing the risk of catastrophic failure that would result in substantial unplanned costs.

The Role of Initial Investment

In addition to the ongoing costs of failure detection and repair, the cost function also includes an initial investment (I), which represents the cost of implementing the predictive model and maintaining the system. This investment can be broken down into several components, including the cost of purchasing the machine learning model, operator man-hours for running and maintaining the system, and any other associated operational costs. This initial investment is incurred periodically, such as on a monthly basis, and is a fixed cost that must be considered when calculating the total cost of maintaining the system. Even if the machine does not experience failures during a particular period, the operational and maintenance costs still need to be accounted for.

The Impact of Undetectable Failures

A significant portion of the cost in this model comes from undetectable failures. When the system fails to detect an actual problem, the consequences can be severe. These failures often result in unanticipated downtime, which directly impacts the production process and leads to lost revenue. For example, if a machine experiences a failure that the system does not detect, it may continue to operate in a compromised state, causing further damage or reducing the efficiency of the system. This increases the likelihood of requiring expensive repairs, potentially including the replacement of components that could have been serviced earlier if the failure had been detected.

In this context, undetectable failures are often the most costly because they are unpredictable. The ability to prevent these failures through early detection is the main benefit of predictive maintenance systems, which is why improving recall is essential. A failure that is detected early on can be addressed in a timely manner, reducing the potential for catastrophic damage. On the other hand, if a failure is not detected, it may escalate into a system breakdown, leading to a much higher cost due to the need for emergency repairs, downtime, and lost production opportunities.

Understanding the Economics of Detectable vs. Undetectable Failures

The fundamental distinction between detectable and undetectable failures forms the backbone of the cost model. This separation allows for a clearer understanding of the costs associated with different types of failures and provides a basis for decision-making around system improvements, investments, and operational strategies. As mentioned, detectable failures are those that can be identified and mitigated by the predictive maintenance system, while undetectable failures are the ones that the system misses, leading to unexpected and costly consequences.

Detectable failures typically have lower costs associated with them because they are identified early on. Early detection allows for preventive maintenance, which involves addressing the failure before it results in significant damage. This early intervention ensures that repairs can be planned, minimizing system downtime and preventing the failure from escalating. Furthermore, detectable failures often lead to relatively low repair costs since the components or machines affected are typically still operable to some extent, and only a small intervention (such as a component replacement or calibration) is needed.

However, the undetectable failures introduce substantial risk and higher costs. These failures are more challenging to detect because the system is either unaware of the impending problem or has missed signals that would indicate failure. The impact of undetectable failures can range from minor to catastrophic. In many industrial settings, the failure of a single critical component can result in significant production delays, costly system overhauls, or even irreversible damage to the machinery or system. These consequences contribute to the much higher cost associated with undetectable failures, represented by the variable $C_d * (1 + K)$ in the equation.

Moreover, the system's failure to detect these problems also leads to additional, often unseen, costs. For example, undetected failures can lead to unplanned downtime, where the machine or system is out of operation due to the failure's unexpected nature. This downtime often results in significant productivity losses, which is not only a direct loss in revenue but also can have a cascading effect on the supply chain, delaying deliveries, and affecting customer satisfaction. For businesses relying on continuous production, this can translate into penalties, reputational damage, and even the loss of customers.

The Trade-Off Between Precision and Recall

One of the most challenging aspects of predictive maintenance systems is the trade-off between precision and recall. Improving recall often leads to a decrease in precision, and vice versa. This trade-off arises because, to increase recall, the system may become more sensitive to potential failures, leading to more failures being flagged, including some that don't require intervention. This improves recall but may reduce precision. On the other hand, improving precision often involves making the system more conservative in its failure detection, which may result in fewer failures being detected but with higher accuracy.

In practice, organizations need to balance these two metrics based on the specific needs of their operations. For example, in high-risk industries like aviation or oil and gas, where undetectable failures can lead to catastrophic consequences, the priority may be to improve recall, even at the cost of reducing precision. In such cases, the financial impact of missing a failure (and potentially leading to system failure or disaster) is far greater than the cost of flagging a few non-critical failures.

Conversely, in environments where downtime and maintenance costs are a critical concern but the consequences of undetectable failures are less severe, businesses may prioritize precision. By focusing on reducing false positives, these organizations can save on unnecessary repairs, prevent over-maintenance, and optimize their resources. However, even with a focus on precision, some level of recall must still be maintained to prevent catastrophic system failures.

Thus, the cost function is dynamic, as it responds to the balance between precision and recall, which in turn affects the number of failures detected and the costs associated with those failures. The ultimate goal is to achieve an optimal point where the system detects a sufficient number of failures without flagging too many false alarms, reducing the overall operational cost while minimizing the risk of catastrophic failures.

Initial Investment and Long-Term Savings

Another critical aspect of this cost model is the inclusion of the initial investment (I), which accounts for the costs associated with setting up the predictive maintenance system. These costs typically include purchasing the necessary sensors, software, and hardware, as well as the manpower required to implement, operate, and maintain the system.

While the initial investment may seem like a significant expense, it is important to recognize that this investment is aimed at preventing much higher costs in the long run. Predictive maintenance systems are designed to reduce the overall cost of failure detection and repair by improving the precision and recall of the system. In turn, this minimizes the occurrence of undetectable failures and unnecessary repairs. The initial investment, therefore, pays for itself over time through savings from reduced downtime, fewer catastrophic failures, and optimized resource utilization.

However, the effectiveness of the system in delivering these savings depends heavily on how well the system is maintained and calibrated. Regular updates and proper maintenance of the predictive model are crucial to ensuring that the system remains accurate and reliable. The operator man hours cost included in the equation represents the ongoing efforts needed to maintain the system's functionality, update its models, and ensure that it continues to deliver value over time.

Cost failure analysis when there are multiple failure classes

We can expand our cost model further since we have multiple types of failures into the following expansion

$$Cost \;=\; F_d * C_d * P_r + F_d * C_d * (1 + k) * (1 - R_c) \;+\; I$$

$$Cost_{TOTAL} = \sum_i^n \quad Cost_i \;=\; I + \sum_i^n \quad F_{di} * C_{di} * P_{ri} + F_{di} * C_d * (1 + k_i) * (1 - R_{ci})$$

$Cost_{TOTAL}$ is assumed to be the total cost of the sumatıon of all the costs of the predıctıve maintenance program given the period . This form is used when there are multiple failure types in a multi-classification problem in order to take into account the different precions and recalls of each failure type since we may have multiple failure types.

$$Cost_{TOTAL} = \sum_{i}^{n} \quad Cost_i$$

$$Cost_{TOTAL} = Cost_{Power\ Failure} + Cost_{Tool\ Wear\ Failure} + Cost_{Overstrain\ Failure} + Cost_{Heat\ Dissipation\ Failure} + Cost_{Random\ Failures}$$

Final Thoughts on Total Cost Calculation:

By plugging in the precision, recall, and failure detection information for each failure type, we can compute the cost of predictive maintenance for each failure type and then sum them to get the total cost. The precision and recall values significantly influence the cost, particularly the costs associated with undetected failures, which can be high in failure types with lower recall.

The extended cost model provides a comprehensive framework for evaluating and optimizing a predictive maintenance program. It incorporates both the cost of detecting failures and the cost of failing to detect them, ensuring a balanced and informed approach to maintenance strategies.

The initial investment, I, is assumed to be unified due to its singular nature as an upfront cost. This includes the cost of a software subscription or maintenance service, which remains consistent regardless of the types of failures considered. However, investments may vary depending on different SaaS payment structures, which are often piecewise and not fixed.

Nonlinearization of the cost function to mimic real-world behavior

In real-world scenarios, failures often lead to a drastic increase in costs, especially when these failures are compounded over a given period. The more frequently an event occurs, the higher the likelihood that it could lead to a catastrophic failure, which not only affects the operation but can also have a significant financial impact. This is particularly important in industries where safety, regulations, and the nature of the equipment amplify the consequences of a failure. For instance, a failure in an airliner or a rocket is much more expensive compared to a failure in a textile machine. The cost discrepancy stems from several factors, including the complexity of the machinery, the safety standards required in those industries, and the far-reaching effects of a failure, such as loss of life or environmental damage.

Recognizing this complexity, I have devised a method to account for the non-linear escalation of failure costs over time. This method incorporates two parameters—alpha and beta—that are designed to reflect the real-world behavior of cost accumulation due to failures.

$$Cost = \boldsymbol{\alpha}[F_d * C_d * P_r]^{\boldsymbol{\beta}} + \boldsymbol{\alpha}[F_d * C_d * (1 + k) * (1 - R_c)\,]^{\boldsymbol{\beta}} + I$$

Alpha is an industry-specific parameter. In certain industries, a failure may carry a higher financial burden due to industry-specific factors such as stringent safety regulations, legal requirements, and public safety concerns. For example, the aerospace industry deals with incredibly high costs in the event of a failure because of safety protocols, insurance liabilities, and the potential for catastrophic outcomes. In contrast, industries like textiles might not face the same level of scrutiny or high-risk outcomes, so the failure costs are generally lower. By incorporating alpha, I can model how different industries will be affected differently by failures, allowing for a more accurate reflection of their specific cost structures.

Beta, on the other hand, is a machine-specific parameter that accounts for the likelihood of failure based on the complexity of the machinery in question. More complex machines—such as aircraft or industrial robots—tend to have a higher likelihood of failure simply due to the intricacies of their design, the number of components involved, and the operational environments they function in. As complexity increases, the failure rate tends to grow non-linearly, meaning that as machines age or accumulate more failures, the cost of failure compounds, often leading to a steeper increase in maintenance and repair costs. By introducing beta, I can model this compounding effect, which is crucial for predicting and managing long-term maintenance expenses, particularly for high-complexity machines.

To determine the values for alpha and beta, a predictive maintenance operator can approach the task in two ways. One option is to choose arbitrary values for these parameters based on industry standards, available knowledge, or estimates of potential failure costs. This method can provide a quick approximation of failure costs, though it may not be as precise. The second, more accurate method involves fitting these parameters using historical cost data through a nonlinear regression function. By analyzing past data, the operator can identify trends in failure costs over time and adjust the values of alpha and beta to better match the specific cost dynamics of their equipment and industry. This approach ensures that the cost function is as accurate as possible, making the predictive maintenance model more reliable and effective in forecasting future costs and maintenance needs.

In summary, by introducing alpha and beta, I am able to model and adjust for the unique characteristics of both the industry and the machinery, providing a more nuanced and realistic understanding of the long-term financial impact of equipment failures. This approach helps organizations to make more informed decisions about maintenance strategies, budgeting, and risk management.

Results

The C-MAPSS dataset (Saxena et al. 2008), serves as a benchmark for developing predictive maintenance models, particularly in the domain of aircraft engine health monitoring. This dataset simulates realistic engine degradation by capturing 26 sensor readings, operational conditions, and cycle times, thereby reflecting the complexities of real-world aviation maintenance challenges. It consists of four subsets (FD001–FD004), each representing distinct fault scenarios and varying operational conditions. The dataset enables the development of machine learning models that account for sensor noise, operational variability, and nonlinear degradation patterns, making it invaluable for predictive maintenance research. A Bidirectional GRU (Gated Recurrent Unit) model was employed, with preprocessing steps including data normalization, noise reduction, and feature extraction via computed rolling statistics. The Remaining Useful Life (RUL) threshold was set at 35 cycles, with failures classified as binary events (1 for failure, 0 for normal operation). This methodological approach enhances the model's ability to discern complex temporal dependencies within the degradation patterns.

The model demonstrated 86.25% accuracy, suggesting a high-level capacity for distinguishing between functional and failing engine states. However, accuracy alone is an inadequate measure in predictive maintenance, particularly given the asymmetrical costs associated with misclassification. The model achieved a precision of 82.46%, indicating a low rate of false positives and thereby minimizing unnecessary maintenance actions. Conversely, its recall stood at 61.13%, meaning that a substantial proportion of actual failures were not detected. This discrepancy implies that the model is conservative in failure predictions, favoring precision over recall and consequently overlooking critical failure instances. Such a trade-off is problematic in industrial settings, where undetected failures can lead to catastrophic system breakdowns, prolonged downtimes, and financial losses.

A fundamental consideration in predictive maintenance is the cost asymmetry governed by α (false positives) and β (false negatives). While α influences incremental increases in maintenance expenditures, the effect of β is highly nonlinear, leading to exponential cost escalation as missed failures result in severe system disruptions. Unlike α, where minor adjustments lead to marginal cost variations, β has a multiplicative effect on total operational costs due to the cascading consequences of unplanned outages. In complex industrial systems, unexpected failures necessitate prolonged repair times, disrupt supply chains, and significantly erode revenue streams. Thus, a model prioritizing high precision at the cost of recall might not be economically optimal, as the risk of underestimating failure events outweighs the burden of excessive maintenance.

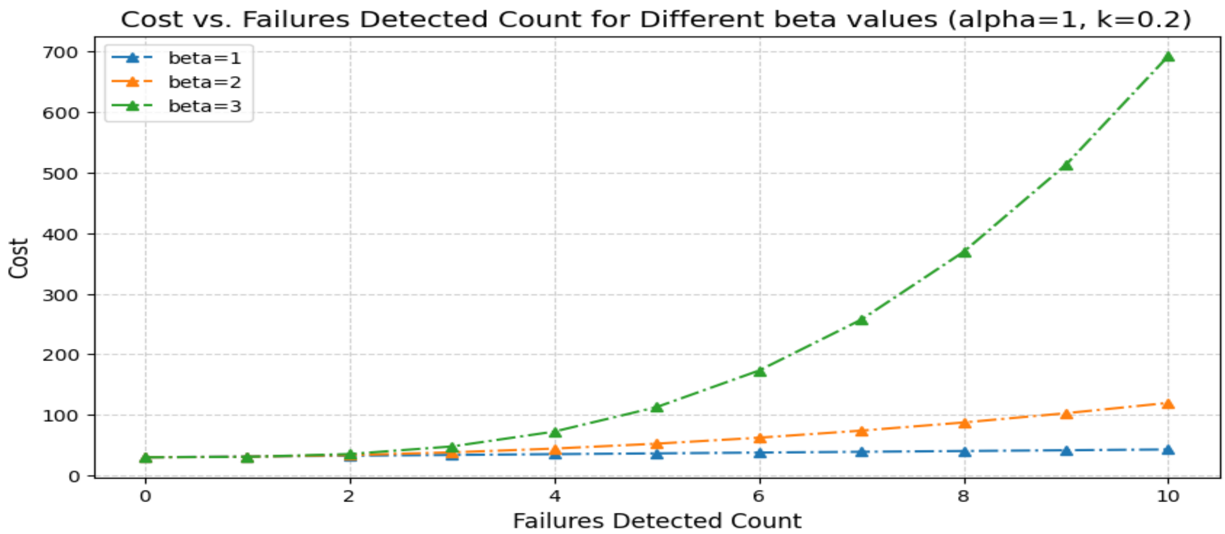

Chart 1: Cost varying with the change in alpha

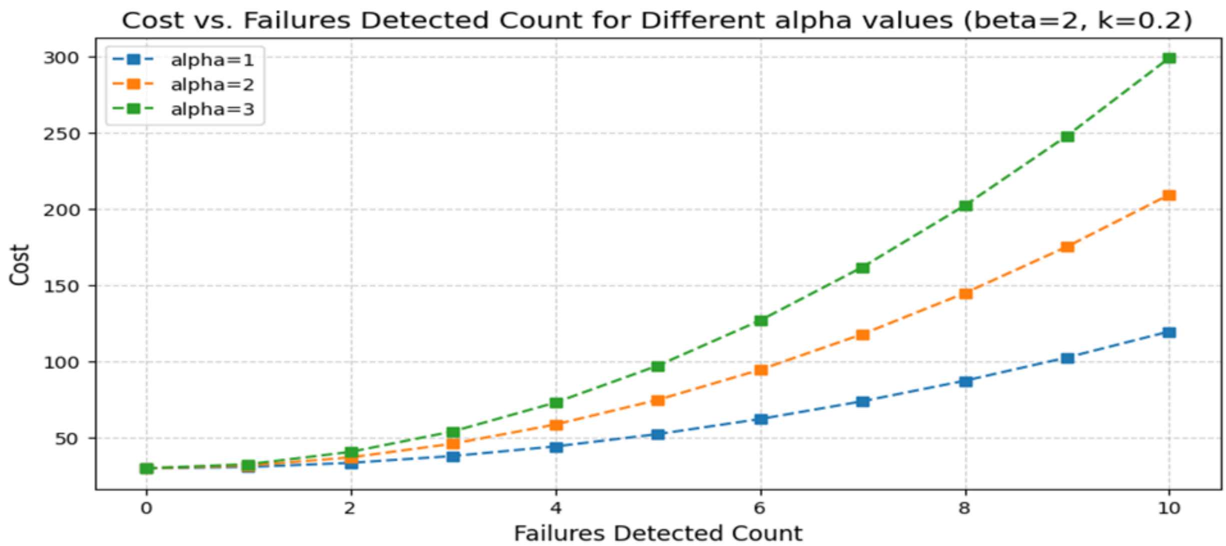

chart 2: cost varying with change in beta

While the model achieves strong accuracy and precision, its limited recall constrains its applicability in high-stakes predictive maintenance scenarios. Future refinements should focus on adjusting classification thresholds, leveraging cost-sensitive learning approaches, and enhancing feature engineering techniques to mitigate the economic risks associated with undetected failures. A more balanced trade-off between α and β would ensure that predictive maintenance models are not only statistically robust but also aligned with the financial realities of industrial asset management.

III. DISCUSSION

The results from the extended cost model provide valuable insights into optimizing predictive maintenance programs by weighing both the costs of detecting failures and the costs associated with undetected failures. By minimizing these costs, companies can improve operational efficiency, reduce downtime, and enhance resource allocation. However, the effectiveness of this model heavily depends on the accuracy of the input data, including historical failure rates and operational costs, as well as the specific predictive maintenance strategy employed by each company.

Potential Expansion: Operational Research (OR) with Failure Costs

One key expansion of the model is the incorporation of Operational Research (OR) techniques to further minimize failure-related costs. By considering various failure scenarios and their associated costs, it’s possible to refine the decision-making process for when and how maintenance activities should occur. This involves exploring optimal maintenance scheduling and leveraging cost-benefit analysis to strike the best balance between preventive and corrective maintenance actions, thus improving overall cost efficiency.

Non Linearization of Terms for More Accurate Results

The current model assumes linear relationships between the costs and maintenance strategies. However, nonlinear regression offers the potential to capture more complex interactions and provide better-fitting models for certain industrial scenarios. Many factors, such as the age of machinery, wear-and-tear dynamics, and machine learning algorithm parameters, may exhibit nonlinear behavior that significantly impacts maintenance costs. Nonlinear models could more accurately predict costs across different stages of equipment life cycles, thus allowing for better-tailored predictive maintenance plans.

Mass Survey for Market-Wide Standardization

To enhance the applicability and scalability of this model, a mass survey across multiple industries could help gather feedback and refine the model. By collecting data from a wide range of companies operating with diverse machines and maintenance strategies, it would be possible to standardize the alpha and beta parameters of the nonlinear cost equation. These parameters would account for various failure rates, detection times, and cost factors, ensuring that the predictive maintenance model remains relevant and accurate across a broad spectrum of industries.

Impact on Predicting Costs for Any Factory Using Predictive Maintenance

With standardized parameters, the model could be used universally, allowing any factory to predict the costs of implementing predictive maintenance using machine learning. This would help in forecasting the required investment, operational expenses, and potential savings from reduced downtime and failures.

It would also enable factories to perform what-if analyses, allowing them to evaluate different predictive maintenance strategies and select the one that maximizes cost efficiency based on their specific operational context.

IV. CONCLUSION

By integrating real-world data through surveys and continuous feedback, the predictive maintenance framework can evolve into a robust, industry-wide tool for cost prediction and optimization. This research presents a cost function and framework to assess the financial implications of adopting a predictive maintenance strategy, leveraging machine learning parameters such as recall and precision. Additionally, the use of nonlinear equations enables a more realistic representation of cost dynamics, capturing complex relationships that traditional linear models may overlook.

Furthermore, this approach holds significant potential for multi-failure scenarios, where multiple failures occur within the same system. By optimizing costs across various failure types, the framework can enhance efficiency and decision-making in maintenance strategies. Finally, industry-wide standardization of key parameters, such as alpha and β, through comprehensive surveys could further refine predictive maintenance models, fostering greater adoption and consistency across industries.

REFERENCES

Ali, Mohamed Iyad, Nai Shyan Lai, and Raed Abdulla. "Predictive Maintenance of Rotational Machinery Using Deep Learning." School of Engineering, Asia Pacific University Technology and Innovation, Kuala Lumpur, Malaysia.

Chawla, N. V., et al. "SMOTE: Synthetic Minority Over-Sampling Technique." *Journal of Artificial Intelligence Research*, vol. 16, 2002, pp. 321-357. *arXiv*, https://doi.org/10.48550/arXiv.1106.1813.

Chung, Junyoung, et al. *Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling.* Presented in *NIPS 2014 Deep Learning and Representation Learning Workshop*, *arXiv*, 12 Dec. 2014, https://doi.org/10.48550/arXiv.1412.3555.

Deloitte. *Predictive Maintenance: Taking Proactive Measures Based on Data-Driven Insights.* Deloitte, www2.deloitte.com/content/dam/Deloitte/us/Documents/process-and-operations/us-predictive-maintenance.pdf. Accessed 30 Jan. 2025.

Fonarov, Oleg. "Digital Twins and Their Role in Shaping the Future of Manufacturing." *Forbes Technology Council*, *Forbes Councils Member*, 1 Mar. 2024, www.forbes.com/councils/forbestechcouncil/2024/03/01/digital-twins-and-their-role-in-shaping-the-future-of-manufacturing/.

Heaton, Jeff. "An Empirical Analysis of Feature Engineering for Predictive Modeling."

Liu, Yang, et al. "Natural-Logarithm-Rectified Activation Function in Convolutional Neural Networks." National Digital Switching System Engineering and Technological R&D Center, Zhengzhou, China. Supported by the National Natural Science Foundation of China, Grants 61521003, 61601513, and 61803384.

Loncarski, Jelena. Department of Engineering Sciences, Division for Electricity Research, Uppsala University, Sweden.

McKinsey & Company. *Prediction at Scale: How Industry Can Get More Value Out of Maintenance.* 22 July 2021, www.mckinsey.com/capabilities/operations/our-insights/prediction-at-scale-how-industry-can-get-more-value-out-of-maintenance. Accessed 30 Jan. 2025.

Olavsrud, Thor. "Rolls-Royce Turns to Digital Twins to Improve Jet Engine Efficiency." *CIO*, 10 June 2021, https://www.cio.com/article/188765/rolls-royce-turns-to-digital-twins-to-improve-jet-engine-efficiency.html. Accessed 3 Feb. 2025.

Paolanti, Marina, Luca Romeo, Andrea Felicetti, Adriano Mancini, and Emanuele Frontoni. "Machine Learning Approach for Predictive Maintenance in Industry 4.0." *Department of Information Engineering, Università Politecnica delle Marche*, Ancona, Italy.

Potter, Kaledio, and Lucas Doris. "Predictive Maintenance for Electric Vehicles: Enhancing Reliability and Efficiency."

Prytz, Rune, Sławomir Nowaczyk, Thorsteinn Rögnvaldsson, and Stefan Byttner. "Predicting the Need for Vehicle Compressor Repairs Using Maintenance Records and Logged Vehicle Data." *Volvo Group Trucks Technology, Advanced Technology & Research*, Göteborg, Sweden; *Center for Applied Intelligent Systems Research, Halmstad University, Sweden.*

Razaa, Ahmed, and Vladimir Ulansky. "Modelling of Predictive Maintenance for a Periodically Inspected System." *Department of the President's Affairs, Overseas Projects and Maintenance*, Abu Dhabi, UAE; *National Aviation University, Kiev, Ukraine.*

Saxena, A., Goebel, K., Simon, D., and Eklund, N. "Damage Propagation Modeling for Aircraft Engine Run-to-Failure Simulation." *PHM08*, Denver, CO, 2008.