Received: 28 July 2025; Revised: 26 Aug. 2025; Accepted: 29 Sep. 2025
https://doi.org/10.22306/al.v13i1.745

Applying Bayesian Neural Networks to optimize maintenance logistics

Marc Hermans
University of Miskolc, Institute of Logistics, 3515 Miskolc, Egyetem út 1., Hungary, EU,
marc.philip.hermans@student.uni-miskolc.hu (corresponding author)
Peter Tamas
University of Miskolc, Institute of Logistics, 3515 Miskolc, Egyetem út 1., Hungary, EU,
peter.tamas@uni-miskolc.hu

Keywords: maintenance logistics, predictive maintenance, Bayesian Neural Networks, material flow systems, supply chain planning.
Abstract: In maintenance logistics, optimising material and information flows requires predictive methodologies that cope with uncertainty and support decision-making across the maintenance supply chain. This paper proposes a Bayesian Neural Network (BNN) approach to predict maintenance demand and improve logistical planning of spare parts and technician deployment in flow-based production systems. By incorporating uncertainty directly into model outputs, the BNN provides probabilistic predictions that enhance robustness, precision, and responsiveness within logistics operations. The model is trained and validated on historical maintenance data and synthetic scenarios generated to reflect stochastic degradation and repair behaviour typical of industrial flow environments. Compared to deterministic benchmarks, the proposed approach demonstrates superior logistics performance, with improvements in downtime reduction, spare-parts availability, and overall material-flow efficiency. These insights highlight the benefit of probabilistic modelling as a decision-support tool for maintenance logistics in complex manufacturing logistics systems.

1 Introduction

The current global energy crisis triggered by the Russia-Ukraine war, along with the intensifying climate crisis, forces enterprises—irrespective of their industry and operational profiles—to seek solutions that maximize efficiency and sustainability [1]. Maintenance operations are not exempt from these pressures, compelling maintenance teams to optimize resource utilization to ensure the highest possible efficiency.

This study proposes a solution to enhance the efficiency of maintenance logistics in production environments by leveraging Industry 4.0 technologies. However, the chosen approach must remain universally applicable—not only to the latest and most advanced manufacturing systems but also to older production lines. This broader applicability is achieved through state-of-the-art technologies combined with intelligent sensor systems, enabling wider relevance across diverse manufacturing contexts [2].

From an efficiency standpoint, increasing the frequency of preventive interventions may lead to excessive downtime, while delaying interventions can result in sudden failures and production disruptions. Establishing an optimal balance between planned and unplanned events remains a major challenge in maintenance logistics.

Therefore, the objective of this paper is to develop a predictive maintenance logistics approach that integrates probabilistic decision-support into maintenance scheduling and spare-parts provisioning for flow-based production systems. By employing Bayesian Neural Networks (BNNs), we aim to incorporate predictive uncertainty into maintenance decision-making, thereby maximising both production system availability and logistical efficiency.

2 Literature review

2.1 Background and motivation

Efficient maintenance logistics management in complex production environments relies heavily on accurate predictive capabilities to manage uncertainty and enhance operational responsiveness. Traditional maintenance strategies, including preventive maintenance (PM) and condition-based maintenance (CBM), frequently fall short due to their inherent deterministic assumptions, neglecting the stochastic nature of failure occurrences and repairs. Consequently, more sophisticated predictive methods incorporating uncertainty are required.

Within the academic debate on maintenance strategies, Total Productive Maintenance (TPM) plays a significant role and illustrates the challenges of balancing efficiency and downtime. Koch originally described planned maintenance intervals as “potential downtime” — periods in which machinery could be occupied with production but is held idle by maintenance interventions [3]. Other studies, by contrast, position TPM as a valuable concept that reduces system fragility by transforming unplanned failures into pre-emptive, manageable losses [4-6]. These differing perspectives reinforce the

need for uncertainty-aware tools that facilitate sustainable trade-offs between planned and unplanned maintenance activities.

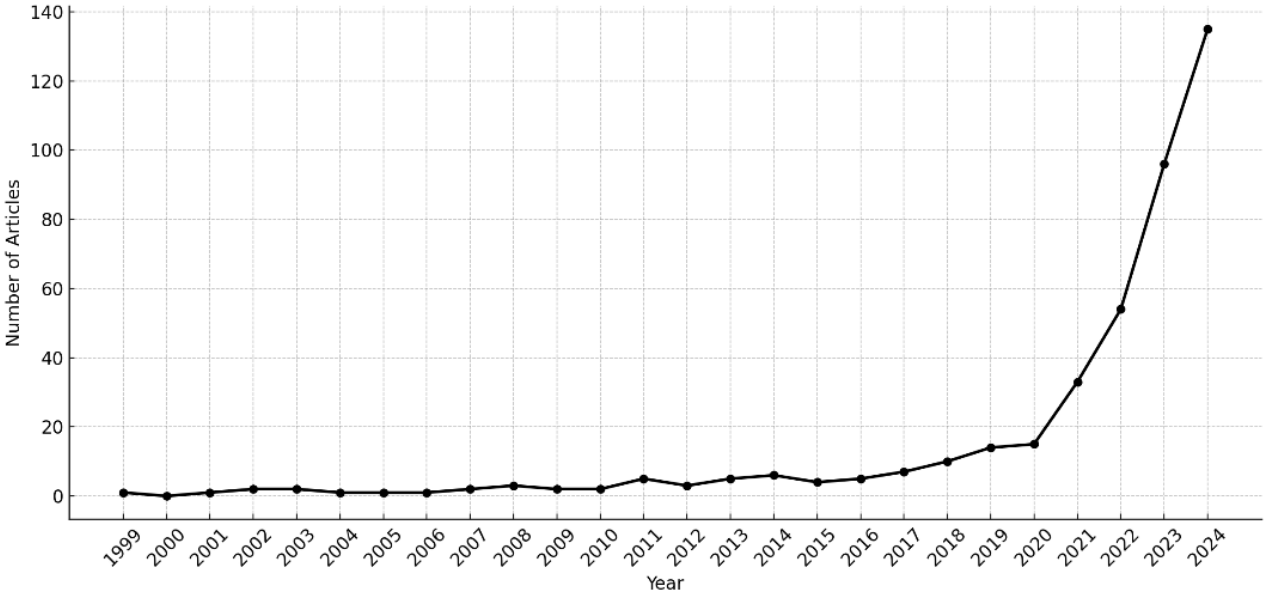

Figure 1 Number of articles published over the years up to 2025, derived from bibliometric analysis of 354 publications identified through systematic literature review

2.2 Bayesian Neural Networks in predictive maintenance

Bayesian Neural Networks (BNNs) have gained increased attention in recent years as promising tools for handling uncertainties in predictive analytics. A comprehensive bibliometric analysis reveals that over the last two decades, scholarly interest in BNNs for maintenance has markedly increased, particularly within Engineering, Computer Science, Decision Science, and Mathematics, with a cumulative total of 354 relevant publications identified. The surge in articles from 38 in 2022 to 92 in 2024 and an estimated 71 in 2025 underscores the growing recognition of BNNs' utility in probabilistic prediction and uncertainty quantification (Figure 1, Figure 2).

Recent literature shows BNNs' wide-ranging applicability in predictive scenarios, particularly for maintenance logistics. [7] Xiao et al. (2023) presented a Bayesian Transformer model, "ProFormer," explicitly addressing predictive uncertainties in machinery fault diagnosis and Remaining Useful Life (RUL) predictions. The authors emphasized that Bayesian approaches inherently manage data uncertainty, significantly improving predictive reliability compared to conventional neural networks. Similarly, Ferreira and Gonçalves (2022) provided a comprehensive overview of the advantages of probabilistic methods in handling uncertainties associated with RUL estimations, reinforcing the utility of Bayesian methods in predictive maintenance.

2.3 Trends and thematic focus in current research

A systematic clustering of literature reveals ten prominent research clusters within the selected domains, highlighting key thematic areas and methodological developments:

Machine Learning in Manufacturing (68 articles): Emphasizes data-driven, machine learning-based predictive approaches in industrial contexts, reflecting the core motivation behind applying advanced neural network architectures.

- Models and Uncertainty (62 articles): Focuses on the explicit incorporation of uncertainty into predictive models, highlighting the essential aspect of Bayesian methods for robust predictions.
- RUL Prediction (54 articles): Represents a central theme where Bayesian techniques offer significant advantages in predicting equipment life, essential for proactive maintenance logistics.
- Fault Diagnosis and Detection (41 articles): Illustrates the application of probabilistic models to improve fault detection accuracy, vital for timely maintenance interventions.
- Maintenance Prognostics (23 articles): Closely examines predictive maintenance methodologies, directly aligning with operational logistics and applied diagnostics—a central focus of the present study.

This clustering highlights substantial methodological diversity and indicates a clear interdisciplinary orientation, reinforcing BNNs' extensive applicability and efficacy in handling complex maintenance problems characterized by uncertainty.

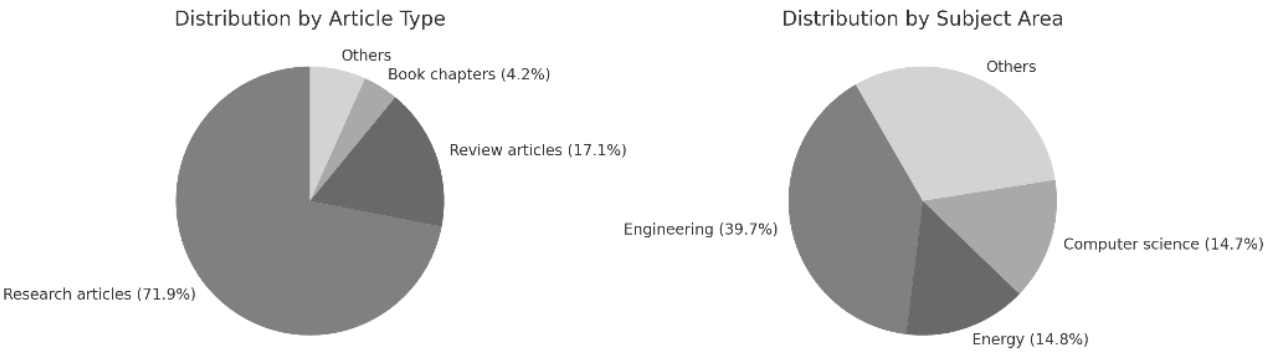

Figure 2 Thematic and typological distribution of publications based on bibliometric analysis of 354 publications across engineering, computer science, decision science, and mathematics domains

2.4 Practical applicability and achievability

The practical feasibility of applying advanced predictive methods, such as BNNs, in industrial maintenance logistics has been supported by recent empirical studies. Notably, Longard et al. (2023) demonstrated the real-time forecasting capability of key performance indicators (KPIs), specifically Overall Equipment Effectiveness (OEE), by leveraging high-frequency sensor and machine data. They identified significant improvements in prediction accuracy through enhanced data granularity and highlighted the importance of high-frequency data labeling for real-time predictive capabilities. Their findings underscore the feasibility and operational benefits of applying sophisticated machine learning frameworks in predictive maintenance.

The successful implementation by Longard et al. further validates the plausibility and potential effectiveness of Bayesian Neural Networks, which inherently offer superior uncertainty quantification compared to traditional deterministic and frequentist approaches. [8] BNNs' probabilistic outputs directly support the critical decisions involved in spare parts provisioning, technician deployment, and overall logistical efficiency, addressing uncertainty explicitly rather than implicitly. The advancement of predictive maintenance maturity models has been recognized as crucial for industrial enterprises seeking to implement sophisticated maintenance strategies. Fidlerová and Chovanová [9] demonstrated that achieving higher levels of predictive maintenance maturity requires comprehensive analysis of machinery and equipment, real-time monitoring capabilities, and proper integration of Industry 4.0 tools. Their findings support the feasibility of implementing advanced predictive methods like BNNs in industrial settings.

2.5 Identified research gap and contribution

Despite extensive research on probabilistic and predictive methods for maintenance logistics, explicit and comprehensive applications of Bayesian Neural Networks remain relatively scarce, particularly in operational contexts involving logistical decisions. Existing literature often limits its scope to specific predictive outcomes (e.g., RUL or fault diagnostics) rather than integrating predictions directly into holistic logistical planning and decision-making processes [10].

This paper addresses this gap by proposing a BNN-driven decision-support framework that embeds probabilistic forecasts directly into maintenance logistics activities. By reframing maintenance scheduling as a probabilistic confidence-threshold problem, our approach explicitly quantifies uncertainty and uses it to trigger logistical actions — such as spare-parts ordering and workforce allocation — thereby enhancing responsiveness and stabilising supply-chain flow performance in complex stochastic production environments.

In contrast to deterministic methods, the proposed framework explicitly incorporates uncertainty into the decision logic itself, enabling confidence-based maintenance actions rather than fixed-threshold planning [11].

3 Methodology

3.1 Maintenance as a Bayesian Decision Problem — from synthetic data to strategic confidence

In predictive maintenance, we often ask: “What is the risk of failure if I don’t act?” But perhaps the more powerful question is:

"How confident are we that we can safely do nothing?"

This question reframes maintenance from being merely reactive to being a measured act of confidence. Instead of waiting for signs of failure or responding to predefined thresholds, we build a system that estimates — with statistical rigor — the probability that our process will continue running smoothly even if we postpone maintenance.

To formalize this, we use Bayesian logic. We ask: what is the probability that OEE (Overall Equipment Effectiveness) will drop below a critical threshold, such as 80%, in the next time step, given the current machine state and a decision not to perform maintenance? If this probability is low, we confidently do nothing. If it is high, we take action. This forms the basis of our Bayesian decision rule.

To estimate this probability, we use a Bayesian Neural Network (BNN). The BNN is trained on a set of features that reflect both performance and degradation — cycle-time-based OEE, time since last maintenance, and the gap between expected and actual performance. Its output is a probability that feeds directly into our decision rule (1):

$$P(OEE\_(t+1) < OEE\_planned \mid evidence, \neg M_t) \quad (1)$$

The BNN allows us to set a confidence threshold, creating a rule that is both interpretable and flexible:

- If the probability is greater than, we perform maintenance.
- If the probability is less than or equal to, we postpone maintenance.

This logic empowers decision-making not only based on risk, but on the confidence in continued performance. To train such a model, we first need relevant and expressive data.

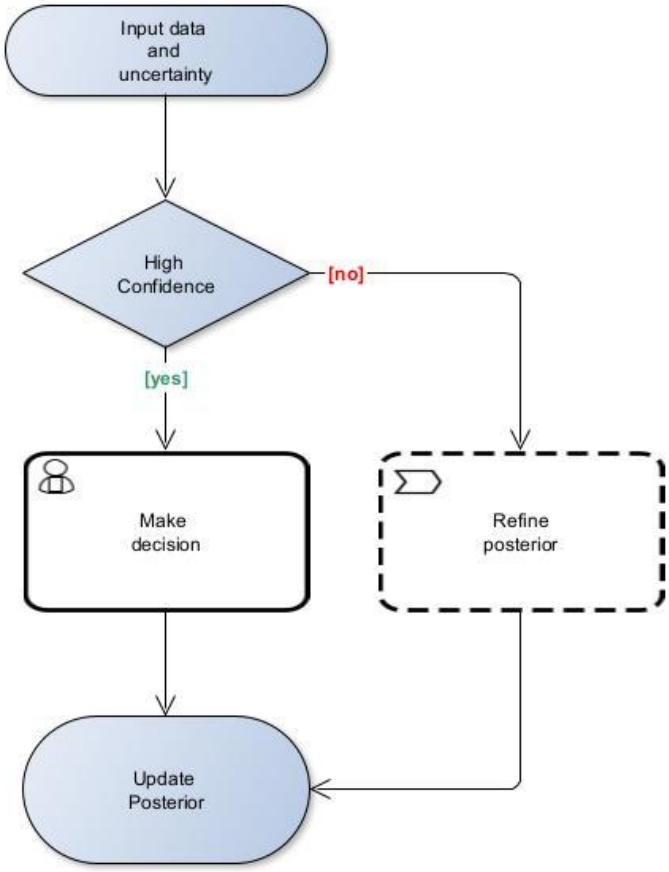

Figure 3 BNN confidence decision flow

Figure 3 illustrates the Bayesian decision framework where the BNN evaluates current machine state (OEE, time since maintenance, performance gap) to calculate P (failure | no maintenance). If this probability exceeds the confidence threshold β, maintenance is triggered; otherwise, operation continues. The feedback loop shows how each decision updates the BNN's posterior beliefs.

We generate synthetic data using a discrete event simulation (DES) environment such as Plant Simulation. Two degradation profiles are simulated: a Weibull-based sawtooth hazard for periodic maintenance, and a Negative Exponential decay for continuous deterioration. This produces data rich in behavioral contrast — from healthy operation to gradual breakdown.

Crucially, OEE is computed from raw timestamped logs [12]. Each cycle time is derived from the difference between consecutive part exit timestamps. Values within the interquartile range (IQR) are considered performance losses, while

outliers are interpreted as availability losses. This minimal data export approach allows for autonomous OEE calculation using only time-based measurements.

Building on this, we introduce the concept of prior and posterior OEE. We measure performance before the machine, assuming ideal availability and quality — this gives us the prior. After the machine, we measure actual performance and availability — this becomes the posterior. The difference between the two, is a signal of unanticipated degradation. This gap acts as a likelihood signal in Bayesian reasoning (Figure 4). The greater the gap, the stronger the evidence that the system is deviating from ideal behaviour.

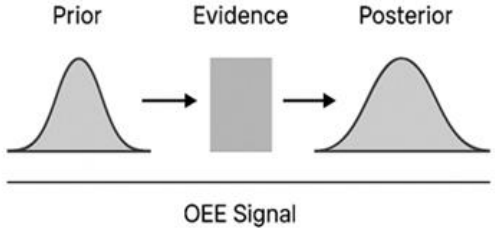

Figure 4 Bayesian interpretation of the OEE signal, conceptual diagram showing the relationship between prior and posterior OEE measurements in the Bayesian framework

The BNN interprets this signal along with other features to estimate future risk. This transforms our approach from failure prediction to confidence estimation. In this framework, maintenance becomes a fully probabilistic decision process. The prior is what we expect, the evidence is what we observe, and the posterior is our updated belief. [13] The action — whether to intervene or not — is taken based on a well-defined confidence threshold. Ultimately, this represents a philosophical shift. Maintenance is no longer simply about avoiding failure. It becomes an expression of trust in the system's resilience, guided by measurable signals and probabilistic reasoning (Figure 5). This framework does not merely predict failure — it quantifies how confident we are that we can safely do nothing.

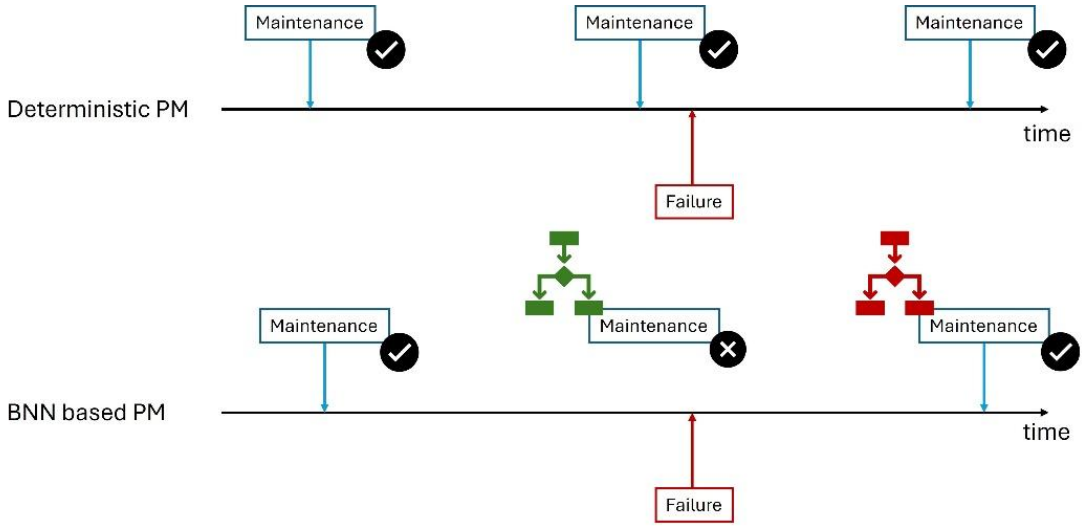

Figure 5 Timeline comparison Deterministic Preventive Maintenance vs. BBN-Based Decision, theoretical comparison illustrating the difference between fixed-interval and probabilistic maintenance scheduling

3.2 Why Bayesian Neural Networks

Bayesian Neural Networks (BNNs) offer a powerful and elegant approach to predictive modelling in manufacturing and maintenance. Unlike traditional neural networks, BNNs are designed to manage uncertainty, incorporate prior knowledge, and adapt over time — making them especially well-suited for complex, evolving environments like serial production. Here's why:

1. They Encode Prior Knowledge: Bayesian methods allow you to embed domain-specific beliefs directly into the model. For instance, you may know from experience or simulation that machine failures are unlikely within the first few minutes of operation. This prior can be reflected in the model before real data is even observed, guiding initial predictions toward realistic behaviour.

2. They Prevent Overfitting: Unlike standard neural networks that may memorize synthetic patterns, BNNs inherently regularize their parameters. By treating weights as distributions, BNNs avoid extreme values unless strongly supported by the data, resulting in models that generalize better — especially important when training begins with synthetic data.
3. They Support Online Learning: BNNs can be incrementally updated as new data arrives, without full retraining. In production, this means your model can improve continuously as each new cycle, temperature reading, or tool failure is observed — keeping the model always aligned with the system's current behaviour.
4. They Use a Distribution of Models — Not Just One: BNNs don't make predictions based on a single set of parameters. Instead, they consider many possible models (parameter sets), each weighted by its probability. This leads to more robust predictions, especially in edge cases or under noisy conditions.
5. They Quantify Their Own Confidence: Because every parameter is treated as a distribution, the model inherently understands which relationships it's confident about and which it's uncertain about. For example, if the influence of cycle time on failure is clear but temperature data is noisy, the BNN will reflect higher confidence in the former.
6. They Provide Probabilistic Outputs: BNNs do not produce a single point prediction, but a distribution over the output — for example, a probability distribution over time to failure or OEE. This allows for risk-aware decision-making, such as: What's the probability that the machine fails before the next scheduled maintenance?"
7. They Are Easily Extendable to New Inputs: As your system evolves, you may begin measuring new factors like temperature or vibration. While a BNN cannot use these signals instantly, it can be easily adapted. You simply expand the model to include the new input, and fine-tune or retrain it. The Bayesian framework will naturally begin incorporating the new information, learning its relevance — along with associated uncertainty — over time.

Together, these properties make Bayesian Neural Networks not merely predictive models, but adaptive decision-support systems that evolve with the production environment (Figure 6). By continuously updating their internal beliefs as new data is observed, BNNs transition from assumptions to reality — transforming prior expectations into confident, data-driven decisions. This makes them ideally suited for preventive and predictive maintenance in high-stakes industrial settings, where uncertainty must not only be managed but actively quantified and acted upon.

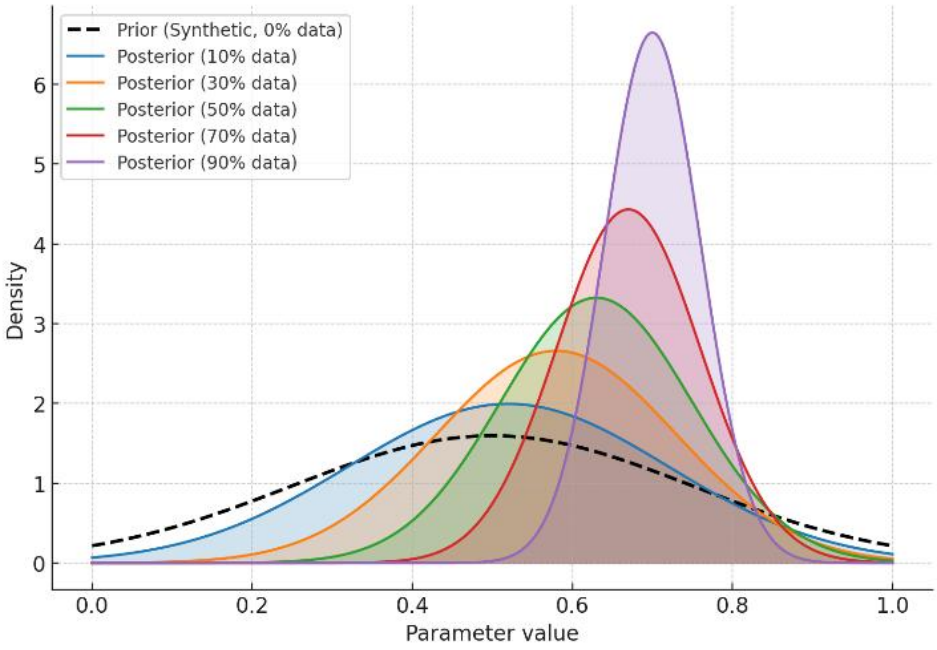

Figure 6 Bayesian Neural Network: posterior evolution with real data, conceptual illustration of how BNN parameters evolve as real operational data is incorporated into the model

4 Results

4.1 Theoretical behaviour of the BNN under varying maintenance regimes

The Bayesian Neural Network (BNN) acts as a real-time "confidence barometer," quantifying both the probability of surviving until the next planned service window and the uncertainty surrounding that estimate. Its guidance naturally diverges across three archetypal scheduling regimes:

4.1.1 Overly aggressive maintenance

When machines are serviced far more often than their condition warrants, unplanned downtime is virtually eliminated—yet at the expense of cost and production flow.

- Expected probability: Almost always near 100%, signalling that the asset could easily keep running until the following service slot.
- Expected uncertainty: Very low, because the asset's state is refreshed so frequently.
- BNN verdict: "Safe to skip." The model persistently recommends deferring the next intervention since the safety margin remains comfortably wide.

4.1.2 Insufficient maintenance

Here, preventive maintenance is insufficiently scheduled, allowing wear to accumulate and failure risk to spike.

- Expected probability: Noticeably below the company's risk threshold—sometimes dramatically so.
- Expected uncertainty: High, reflecting patchy condition updates and growing process variability.
- BNN verdict: "Service now." The network flags an urgent need for corrective action, warning that the likelihood of failure before the next planned stop is unacceptably high.

4.1.3 Stable (optimal) maintenance

An optimal schedule strikes the delicate balance between cost and risk, matching interventions to actual degradation.

- Expected probability: Oscillates around the target safety band, neither excessively high nor dangerously low.
- Expected uncertainty: Moderate and steady, characteristic of a well-monitored line.
- BNN verdict: Condition-specific.
  - If the probability clearly exceeds the safety band, the BNN advises postponement.
  - If it drifts towards or below the lower limit, it advises immediate maintenance.
  - In most cases only minor timing tweaks are suggested, confirming that the existing interval is close to ideal.

4.1.4 Transition scenario

For example, switching to a Cheaper but Less-Reliable Spare-Parts Supplier. Supplier changes are common levers for cost reduction [14], yet they often introduce hidden reliability risk. Moving from an expensive, highly reliable supplier to a cheaper, more failure-prone one typically unfolds in three recognisable phases, each of which elicits a characteristic BNN response:

- Phase A – the honeymoon (first few weeks). At first, nothing changes on the shop floor. Existing inventories still use premium parts, so failure rates and headline KPIs remain exactly where they were. The BNN therefore shows the same healthy picture it always did: a very high probability of surviving to the next service window and minimal uncertainty. Its verdict matches the stable baseline—keep the current maintenance interval and carry on as normal.
- Phase B – divergence (weeks to months). As lower-grade components begin to work their way into production, small warning signs appear: minor stoppages, quality rework and scrap inch upward; throughput variance creeps into the OEE trend. The BNN senses this drift quickly. Its probability estimate of "making it to the next interval" starts to fall, and the associated uncertainty band widens because the data distribution is no longer familiar. In practical terms the model flips from routinely advising "postpone" to urging "maintain." It often recommends shorter, more frequent inspections, early part replacement or tighter incoming-quality checks to prevent those seemingly minor issues from snowballing into major downtime.
- Phase C – the new baseline (or reversion) over the long term. From here the plant can head in one of two directions.
  - Adjusted process: engineering counter-measures—tighter tolerances, extra QA gates—stabilise performance, albeit at higher internal cost. The BNN's probability and uncertainty curves flatten again, but at a less favourable level. It permanently settles on a shorter recommended maintenance interval to keep risk in check.
  - Supplier rollback: management decides the hidden cost outweighs the savings and returns to the premium supplier. The BNN's metrics steadily recover, uncertainty contracts and its advice migrate back toward the original, longer interval.

The BNN acts as an early-warning radar whenever the underlying failure distribution drifts [15] —be it from a supplier switch, a process tweak, or environmental change. By tracking both probability (risk level) and uncertainty (confidence in that risk), it guides maintenance planners smoothly from a Stable regime, through a Divergence period of heightened vigilance, and—once countermeasures are in place—into a new equilibrium, whether that equilibrium is a tighter schedule or a return to the old one.

4.2 Cost- and Logistics-Effects of a BNN-Driven, probabilistic maintenance strategy

4.2.1 Production-cost view

- Eliminating needless interventions:
  A Bayesian Neural Network supplies a confidence score for every planned service window. If that score shows a very low likelihood of failure, the job is simply skipped. Each skipped task saves direct labour [16], avoids set-up losses and keeps spare parts in stock. Plants that switch from rigid time-based PM to this "BNN gate" typically see the count of preventive jobs fall first; the resulting savings snowball because every cancelled stop also erases its associated start-up scrap and lost throughput.
- A live cost-versus-risk surface:
  With OEE now expressed probabilistically, every incremental dip in availability, performance or quality already carries a euro value—lost contribution per minute, scrap cost per part, and so on. Multiply that impact by and you obtain a continuously updated expected-loss curve. Using the saw-tooth hazard-rate model, you can sweep that curve across the planning horizon and pinpoint the moment when the penalty of running on outweighs the cost of a scheduled stop. Figure 7 illustrates the evolution of the expected failure cost over time, calculated using a standard exponential failure model (2). As the risk of failure accumulates, the product of failure probability and economic impact (downtime, repair) surpasses the fixed cost of preventive maintenance. The intersection point indicates the cost-optimal intervention threshold, where BNN-guided policies can interrupt the process before failure becomes economically dominant.

$$Expected\ loss\ loss\ (€) = P(failure\ before\ next\ slot) \times failure - impact\ (€) \quad (2)$$

- Throughput sensitivity in buffered lines:
  Minor speed losses do not stay minor; they ripple through conveyors and buffers in nonlinear ways. In a discrete-event model the BNN can inject stochastic slow-downs so you can watch buffers drain, downstream machines starve, and finished-goods volume disappear. By repeating the run under different maintenance cadences, you find the point where marginal OEE loss meets marginal maintenance cost—the true economic optimum.

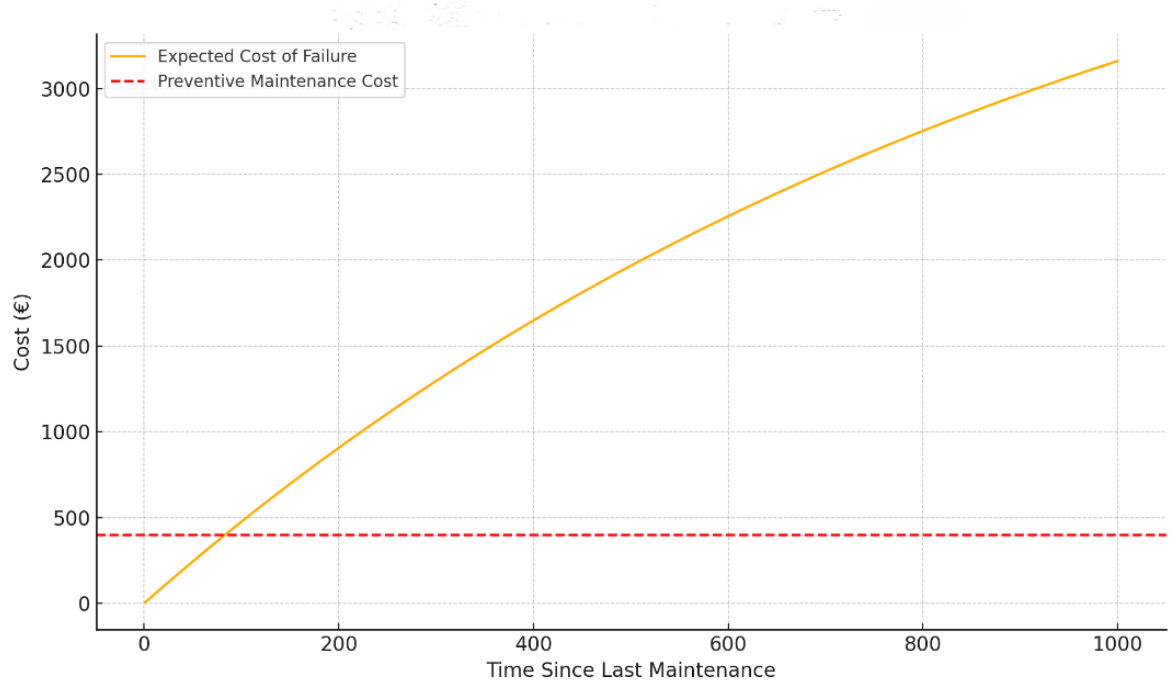

Figure 7 Expected cost of failure versus fixed-cost preventive maintenance across increasing time intervals, the BNN selects the cost-optimal intervention point based on real-time cost–risk trade-offs, calculated using analytical approximation with Python-based cost modelling and exponential failure probability distributions

4.2.2 Logistics-disruption view

- Probabilistic spare-parts pipeline
  Because the BNN delivers a distribution of time-to-failure—not a calendar date—you can replace "just-in-case" inventory with probability-based reorder points. Safety stock falls, freeing working capital, but the supply chain

now relies more heavily on vendor reliability—an exposure that became clear when the plant switched to lower-grade parts. To evaluate the logistical implications of BNN-guided maintenance, Figure 8 compares safety stock requirements across policies with varying confidence thresholds. The BNN's probabilistic time-to-failure forecasts enable a shift from fixed reordering to confidence-based inventory control. This results in a measurable reduction in safety stock while preserving service level targets.

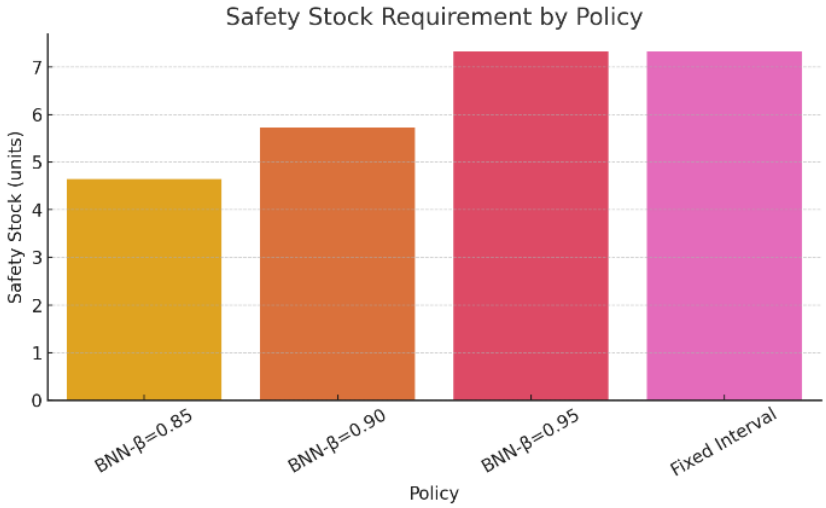

Figure 8 Safety stock requirements under different maintenance triggering policies, BNN-driven thresholds (β = 0.85–0.95) significantly reduce inventory compared to fixed-cycle reordering, derived from analytical approximation with Python-based inventory modelling and probabilistic calculations

- Smarter technician dispatch
  Failure-likelihood windows generated by the BNN allow planners to cluster work orders by geography and time. Wrench-time utilisation rises, call-out fees fall, and truly urgent missions shrink to the thin tail of the probability curve. Figure 9 summarizes improvements in technician dispatch efficiency under BNN control. Emergency interventions are cut by two-thirds, travel time per job drops, and wrench-time utilization rises significantly. These gains highlight the value of probabilistic service windows for clustering geographically similar jobs with overlapping maintenance windows.

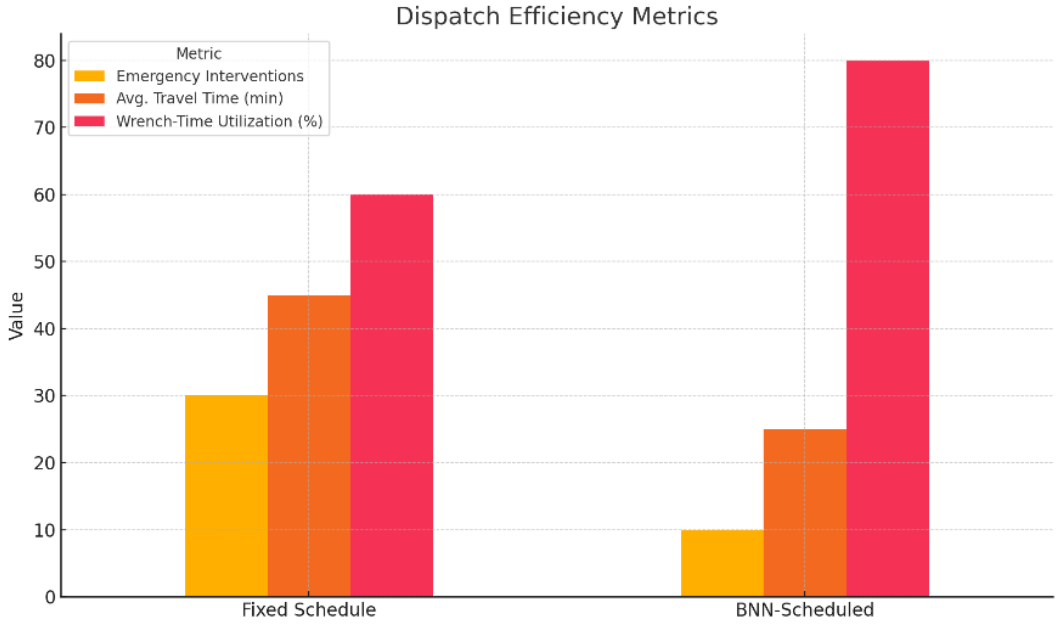

Figure 9 Technician dispatch key performance indicators (KPIs) under fixed scheduling versus BNN-driven dynamic maintenance windows, results obtained from analytical approximation using Python-based scheduling optimisation and probability calculations

- Buffer robustness assessment
  If postponing a job would drain an intermediate buffer before the next slot, the BNN raises a flag. A quick DES run then tells you whether to hold more WIP or invest in another parallel machine as cheaper insurance [17].
- Soft-sensor feedback loop
  High-frequency sensor data, converted to intraday OEE forecasts, spots deterioration hours before the end-of-shift report. Feeding those early warnings back into the BNN tightens its posterior distribution, which in turn sharpens parts ordering, technician routing and short-term production scheduling.

5 Discussion

5.1.1 Logistical implications of BNN-Guided Maintenance

The probabilistic guidance provided by Bayesian Neural Networks (BNNs) does not only reduce unnecessary maintenance but also reshapes the structure and responsiveness of logistics systems that support maintenance operations. Traditionally, logistics planning—covering spare parts provisioning, technician routing, and buffer sizing—relies on fixed-interval service assumptions. These rigid frameworks are prone to overstocking and reactive deployments. By contrast, a BNN-driven maintenance policy can introduce a stochastic trigger mechanism, allowing maintenance logistics to shift from static anticipation to dynamic preparedness.

From a supply chain standpoint, each BNN forecast introduces a time-dependent probability distribution over failure likelihood [18]. This allows spare part reorder points to be recalibrated based on forecasted degradation curves. Instead of reordering "just-in-case" after a predefined cycle, spare part replenishment can now follow a probability threshold, such as (3):

$$P(failure\ before\ part\ arrival) < \alpha \tag{3}$$

where the lead time to procure and deliver the spare part becomes an explicit constraint in the BNN decision logic. This approach reduces safety stock requirements and can lower working capital while still maintaining service levels. Technician dispatching also benefits: the BNN outputs define probable service windows rather than deterministic ones. By clustering jobs with overlapping high-risk intervals and geographically proximate assets, the dispatching problem becomes a constraint-aware routing task with probabilistic priority scoring. This leads to improved wrench-time utilization and a reduction in urgent, high-cost interventions.

5.1.2 Integration into supply chain and operational planning

To model this interaction, a discrete-event simulation (DES) can be extended to incorporate logistical elements such as vehicle travel times, stock depletion, and restocking delays. Scenarios with and without BNN guidance can be compared to quantify improvements in logistics KPIs, including:

- Expedited shipment frequency,
- Average technician travel time per job,
- Fill rate versus stockout probability for key components.

Furthermore, the BNN's confidence threshold β can be adaptively adjusted based on logistics constraints, forming a composite threshold $\beta(\tau)$, where τ includes spare part lead time, technician availability, and warehouse access. In situations of poor logistical flexibility, the model can be tuned to recommend earlier interventions; conversely, when logistics are robust, deferrals become more attractive.

This integration establishes a feedback loop between predictive analytics and operational logistics, closing the gap between data-driven maintenance and real-world execution. By jointly optimizing risk, readiness, and resource consumption, the BNN-guided system evolves into a logistics-aware decision-support tool — bridging predictive maintenance with responsive supply chain management.

6 Conclusion and outlook

This article has introduced a conceptual framework for embedding Bayesian Neural Network (BNN)–driven uncertainty quantification directly into maintenance-logistics decisions. By reframing the classic Remaining Useful Life (RUL) question as a confidence-based decision rule — "How confident are we that we can safely do nothing?" — the study bridges the long-standing gap between predictive-maintenance analytics and the day-to-day realities of spare-parts planning and technician dispatching.

6.1 Key contributions

- Unified decision logic – We formalise a single likelihood-ratio threshold that transforms raw BNN predictive distributions into binary defer/act decisions and, in turn, into tangible logistics actions.

- Maintenance-cost perspective – The approach explicitly links prediction uncertainty to cost elements (downtime, stock, overtime), enabling managers to reason in monetary terms rather than abstract accuracy metrics.
- Generalisability – Because the framework is model-agnostic apart from requiring predictive probabilities, it can incorporate future advances in probabilistic forecasting with minimal process changes.

6.2 Limitations

Although the theoretical underpinnings are laid out in full, the framework has not yet been validated on real-world or large-scale synthetic datasets. The present work therefore stops short of providing empirical evidence of cost savings or service-level improvements. Furthermore, computational considerations (latency, hardware footprint) are analysed only qualitatively at this stage.

6.3 Future work

- Prototype implementation – We are developing a minimally viable BNN service layer and decision API to be tested on a live assembly line; initial latency targets are <200 ms per decision on industrial PCs.
- Simulation-based stress tests – A discrete-event simulation seeded with actual MTBF/MTTR distributions will be used to quantify inventory reduction and uptime gains across a year-long horizon.
- Sensitivity analysis of the confidence threshold (β) – ROC-style curves will help translate risk appetite into numerical β values for different asset classes.
- Field deployment & benchmarking – Partner plants in the automotive and FMCG sectors have agreed to run A/B pilots comparing standard rule-based scheduling with the BNN-enabled logic.
- The synthetic results in Figures 7 to 9 provide preliminary evidence of the BNN's potential to reduce maintenance cost, spare parts inventory, and dispatch overhead. These findings will guide the simulation-based stress tests described in our future work, validating both the economic and logistical performance under realistic failure and demand scenarios.

6.4 Closing remarks

Predictive-maintenance research has advanced rapidly, yet many factories still rely on static service intervals because the operational impact of uncertainty remains opaque. By converting probabilistic predictions into economically meaningful actions, the proposed framework offers a pragmatic path to break that impasse [19]. The forthcoming implementation and validation phases will determine the extent to which the anticipated inventory and downtime benefits—often quoted but rarely measured—can be realised in practice.

References

[1] FRANCIOSI, CH., IUNG, B., MIRANDA, S., RIEMMA, S.: Maintenance for Sustainability in the Industry 4.0 context: a Scoping Literature Review, *IFAC-PapersOnLine*, Vol. 51, No. 11, pp. 903-908, 2018. https://doi.org/10.1016/j.ifacol.2018.08.459

[2] ABDELILLAH, F.M., NORA, H., SAMIR, O., BENSLIMANE, S.M.: *Predictive maintenance approaches in industry 4.0: A systematic literature review*, 2023 IEEE International Conference on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE), Paris, france, pp. 1-6, 2023. https://doi.org/10.1109/wetice57085.2023.10477842

[3] NAKAJIMA, S.: *Introduction to TPM: Total Productive Maintenance*, Productivity Press Inc., University of Minnesota, USA, 1988.

[4] BUDAI, G., DEKKER, R., NICOLAI, R.P.: *Maintenance and Production: A Review of Planning Models,* In: Complex System Maintenance Handbook, Springer Series in Reliability Engineering, Springer, London, pp. 321-344, 2008. https://doi.org/10.1007/978-1-84800-011-7_13

[5] DEMARTINI, M., TONELLI, F., PACELLA, M., PAPADIA, G.: A review of Production Planning Models: Emerging features and limitations compared to practical implementation, *Procedia CIRP*, Vol. 104, pp. 588-593, 2021. https://doi.org/10.1016/j.procir.2021.11.099

[6] AGHEZZAF, E.H., JAMALI, M.A., AIT-KADI, D.: An integrated production and preventive maintenance planning model, *European journal of operational research*, Vol. 181, No. 2, pp. 679-685, 2007. https://doi.org/10.1016/j.ejor.2006.06.032

[7] IFTIKHAR, N., LIN, Y.-C., NORDBJERG, F.: *Machine learning based predictive maintenance in manufacturing industry*, Proceedings of the 3rd International Conference on Innovative Intelligent Industrial Production and Logistics IN4PL - Volume 1, Valletta, Malta, pp. 85-93, 2022. https://doi.org/10.5220/0011537300003329

[8] JASKÓ, S., SKROP, A., HOLCZINGER, T., CHOVÁN, T., ABONYI, J: Development of manufacturing execution systems in accordance with Industry 4.0 requirements, *Computers in Industry*, Vol. 123, 103300, pp. 1-18, 2020. https://doi.org/10.1016/j.compind.2020.103300

[9] MESAROSOVA, J., MARTINOVICOVA, K., FIDLEROVA, H., HRABLIK CHOVANOVA, H., BABCANOVA, D., SAMAKOVA, J.: Improving the Level of Predictive Maintenance Maturity Matrix in Industrial Enterprise, *Acta logistica*, Vol. 9, No. 2, pp. 183-193, 2022. https://doi.org/10.22306/al.v9i2.292

[10] ONG, K.S., WANG, W., NIYATO, D., FRIEDRICHS, T.: Deep-reinforcement-learning-based predictive maintenance model for effective resource management in industrial IOT, *IEEE Internet of Things Journal*, Vol. 9, No. 7, pp. 5173-5188, 2022. https://doi.org/10.1109/jiot.2021.3109955

[11] LIU, X.-W.: *Fuzzy inference based aggregation method for multiobjective decision making problems*, Proceedings, International Conference on Machine Learning and Cybernetics, Beijing, China, Vol. 3, pp. 1296-1300, 2002. https://doi.org/10.1109/icmlc.2002.1167413

[12] GONZÁLEZ, I., SÁNCHEZ, I.: Variable selection for Multivariate Statistical Process Control, *Journal of Quality Technology*, Vol. 42, No. 3, pp. 242-259, 2010. https://doi.org/10.1080/00224065.2010.11917822

[13] QIU, S., MING, X.: Explicit and implicit valuation-based system methods for the risk assessment of systems subject to common-cause failures under uncertainty, *Knowledge-Based Systems*, Vol. 214, 106665, pp. 1-14, 2021. https://doi.org/10.1016/j.knosys.2020.106665

[14] LIU, M., QIN, J., LU, D.-G., ZHANG, W.-H., ZHU, J.-S., FABER, M.H.: Towards resilience of offshore wind farms: A framework and application to asset integrity management, *Applied Energy*, Vol. 322, 119429, pp. 1-18, 2022. https://doi.org/10.1016/j.apenergy.2022.119429

[15] ABU-SAMAH, A., SHAHZAD, M.K., ZAMAI, E., SAID, A.B.: Failure prediction methodology for improved proactive maintenance using Bayesian approach, *IFAC-PapersOnLine*, Vol. 48, No. 3, pp. 844-851, 2015. https://doi.org/10.1016/j.ifacol.2015.09.632

[16] BASTOS, N., LOURES, E.R., SANTOS, E.A.P., DE PAULA, M.A.B.: Production process efficiency analysis: an approach based on colored Petri Nets, *IFAC Proceedings* Volumes, Vol. 41, No. 2, pp. 1863-1868, 2008. https://doi.org/10.3182/20080706-5-KR-1001.00318

[17] RUANE, P., WALSH, P., COSGROVE, J.: Development of a digital model and metamodel to improve the performance of an automated manufacturing line, *Journal of Manufacturing Systems*, Vol. 65, pp. 538-549, 2022. https://doi.org/10.1016/j.jmsy.2022.10.011

[18] ETXEBERRIA, I., BASTIDA, S., ALBERDI, S.: Digital Twin for maintenance: a literature review, *Computers in Industry*, Vol. 123, 103316, pp. 1-14, 2020. https://doi.org/10.1016/j.compind.2020.103316

[19] LANDSTRÖM, A., ALMSTRÖM, P., WINROTH, M., ANDERSSON, C., ÖBERG, A.E., KURDVE, M., SHAHBAZI, S., WIKTORSSON, M., WINDMARK, C., ZACKRISSON, M.: A life cycle approach to business performance measurement systems, *Procedia Manufacturing*, Vol. 25, pp. 126-133, 2018. https://doi.org/10.1016/j.promfg.2018.06.066

Review process
Single-blind peer review process.