Monotonic Latent Manifold Bayesian Inference for Robust Industrial Predictive Maintenance

Oyetayo Oyebisi[1]

[1] *Pan African University Institute for Basic Sciences, Technology and Innovation, Juja, Kiambu, 01001, Kenya*
oyetayooyebisi@pau-au.africa

ABSTRACT

Industrial predictive maintenance systems require not only accurate fault prediction but also reliable uncertainty estimation and robustness under noisy operational conditions. However, conventional deep learning approaches often produce temporally inconsistent failure trajectories and exhibit poor probabilistic calibration when exposed to stochastic sensor disturbances.

This study proposes a Monotonic Latent Manifold Bayesian Neural Network (MM-BNN) framework for robust industrial predictive maintenance. The MM-BNN models degradation progression as a constrained non-decreasing latent probabilistic trajectory, thereby ensuring physically consistent temporal evolution while preserving Bayesian uncertainty quantification. A monotonic latent manifold update mechanism is introduced to enforce structured degradation dynamics across sequential observations. In addition, Bayesian posterior inference is performed using the No-U-Turn Sampler (NUTS) to capture uncertainty in latent degradation states.

The proposed MM-BNN framework is evaluated against multiple baseline models, including multilayer perceptron networks, gated recurrent units, random forests, support vector machines, XGBoost, and standard Bayesian neural networks. Experimental results demonstrate that the MM-BNN achieves superior discriminative performance with an average Area Under the Curve (AUC) of 0.772, outperforming competing approaches, most of which remain close to random classification performance.

Reliability analysis further demonstrates improved probabilistic calibration, with a low Expected Calibration Error (ECE $\approx$ 0.034) and stable reliability behavior under increasing levels of stochastic noise. Ablation studies confirm that the monotonic degradation constraint is the primary factor driving predictive stability and robustness improvements. Furthermore, robustness experiments show that the MM-BNN maintains stable predictive performance even under severe noise perturbations.

Overall, the findings indicate that the proposed MM-BNN framework provides a reliable, interpretable, and uncertainty-aware approach for predictive maintenance in safety-critical industrial environments.

1. INTRODUCTION

The rapid advancement of Industrial Internet of Things (IIoT) technologies and large-scale sensor infrastructures has significantly transformed modern industrial operations and intelligent manufacturing systems(Kordestani, Orchard, Khorasani, Saif & Saif, 2023; Visconti, Rausa, Del-Valle-Soto, Velazquez, Cafagna & De Fazio, 2024). Continuous monitoring through interconnected sensors and cyber-physical systems enables industries to collect large volumes of operational data in real time, creating new opportunities for predictive maintenance, improved system reliability, reduced operational downtime, and intelligent decision-making in Industry 4.0 environments (Zonta, Da Costa, da Rosa Right, De Lima, Da Trindade & Li, 2020 ; Chaudhry, 2024)

As industries increasingly adopt Industry 4.0 technologies, including cyber-physical systems, cloud computing, and edge intelligence, predictive maintenance has emerged as one of the most important applications of intelligent industrial analytics (Kordestani et al., 2023).

Conventional maintenance strategies such as corrective and preventive maintenance often fail to satisfy the operational requirements of modern complex industrial environments (Carvalho, Soares, Vita, Francisco, Basto, & Alcala, 2019). Corrective maintenance is inherently reactive because maintenance actions are initiated only after equipment failure occurs, frequently resulting in unplanned downtime, production disruptions, safety risks, and increased operational costs (Mobley, 2002; Zonta et al., 2020). Preventive maintenance improves this strategy by scheduling maintenance at predefined intervals; however, fixed schedules may not reflect the real-time health condition and degradation state of industrial

https://doi.org/10.36001/IJPHM.2026.v17i2.4874

assets, often leading to either unnecessary maintenance actions or unexpected failures (Carvalho et al., 2019; Chaudhry, 2024). Predictive maintenance addresses these limitations by utilizing operational and sensor data to estimate equipment degradation and predict future failures before they occur (Mobley, 2002; Si, Wang, Hu & Zhou, 2011).

Recent progress in machine learning and deep learning has significantly enhanced the capabilities of predictive maintenance systems in modern industrial environments (Li, Wang & Chen, 2025; Zhang, Wang & Li, 2023). Data-driven approaches can automatically learn complex degradation patterns and nonlinear relationships from high-dimensional industrial sensor streams without extensive reliance on manual feature engineering, thereby improving fault diagnosis, degradation assessment, and remaining useful life prediction (Wang, Cheng, Liu & Jia, 2024; Zonta et al., 2020). Deep learning architectures such as recurrent neural networks (RNNs), long short-term memory networks (LSTMs), gated recurrent units (GRUs), convolutional neural networks (CNNs), and transformer-based models have demonstrated strong performance in fault diagnosis, anomaly detection, and remaining useful life prediction tasks (Wang et al., 2024; Li et al., 2025) . These models are capable of capturing nonlinear temporal dependencies and extracting hidden degradation characteristics from multivariate industrial data.

Despite these advances, many existing predictive maintenance models primarily emphasize prediction accuracy while paying limited attention to uncertainty estimation and probabilistic reliability (Ochella, Dinmohammadi & Shafiee, 2024; Guo, Pleiss, Sun & Weinberger, 2017). In safety-critical industrial environments, maintenance decisions based solely on deterministic predictions may introduce substantial operational and economic risks because model confidence may not accurately reflect true system uncertainty under dynamic operating conditions (Kendall & Gal, 2017; Sajjadi, Dinmohammadi & Shafiee, 2025) . Incorrect confidence estimation can lead either to unnecessary maintenance actions or delayed interventions that may result in catastrophic equipment failure. Therefore, predictive maintenance systems should not only provide accurate predictions but also quantify predictive uncertainty (Kendall & Gal, 2017).

Bayesian deep learning offers a probabilistic alternative to deterministic neural networks by explicitly quantifying predictive uncertainty. Such uncertainty estimation is particularly important for safety-critical predictive maintenance where maintenance decisions must account for confidence as well as prediction accuracy (Ochella et al., 2024; Gal & Ghahramani, 2016).

Although Bayesian inference improves uncertainty estimation, existing prognostic models rarely incorporate structural constraints that reflect the physical characteristics of industrial degradation processes (Ochella et al., 2024; Sajjadi et al., 2025).

In practical industrial systems, degradation is typically cumulative and irreversible due to wear, fatigue, corrosion, vibration, and thermal stress (Xu, Kohtz, Boakye, Gardon & Wang, 2023; Zonta et al., 2020). Equipment operating under such conditions generally exhibits progressively increasing failure risk over time. Many deep learning and Bayesian prognostic models may produce fluctuating or non-monotonic degradation trajectories that contradict known physical behavior (Wehenkel & Louppe, 2019; Runje & Shankaranarayana, 2023). These inconsistencies reduce interpretability and weaken confidence in automated maintenance systems. Consequently, predictive models must align with engineering principles and produce physically meaningful degradation evolution (Xu et al., 2023; Chaudhry, 2024).

To address this challenge, monotonic learning techniques have been explored to incorporate domain knowledge into machine learning models. Monotonic neural networks constrain outputs to follow predefined directional relationships with respect to selected variables, improving interpretability and temporal consistency (Wehenkel & Louppe, 2019; Runje & Shankaranarayana, 2023). Integration of monotonic constraints with Bayesian uncertainty modeling remains limited in predictive maintenance research.

Robustness under noisy operating conditions remains another important challenge because industrial sensor data are often imperfect and highly variable (Chaudhry, 2024; **?**; Zonta et al., 2020; Sajjadi et al., 2025).

Furthermore, explainability and trustworthiness are increasingly important in industrial AI systems (Arrieta, Diaz-Rodrigues, Del Ser, Bennetot, Tabik, Barbado & Others, 2020; Xu et al., 2023). Many deep learning models operate as black-box systems, making interpretation difficult for engineers (Rudin, 2019; Li et al., 2025). This has created demand for predictive frameworks that integrate physical consistency, uncertainty awareness, interpretability, and robustness (Ochella et al., 2024; Chaudhry, 2024).

Motivated by these challenges, this study proposes a Monotonic Latent Manifold Bayesian Neural Network (MM-BNN) framework for robust industrial predictive maintenance. The framework integrates Bayesian uncertainty modeling with monotonic latent degradation learning to enforce physically consistent failure progression over time.

In addition, latent manifold representation learning is used to capture nonlinear degradation dynamics from high-dimensional sensor data. This improves temporal continuity and robustness under noisy conditions. Bayesian posterior inference using the No-U-Turn Sampler further enhances uncertainty estimation and calibration reliability.

The proposed framework is evaluated against multilayer perceptron (MLP), GRU, Random Forest, Support Vector Machine (SVM), XGBoost, and standard Bayesian neural networks. Results show improved accuracy, calibration, and robustness. Ablation studies confirm that monotonic constraints significantly improve stability and generalization.

The major contributions of this study include:

- Development of a monotonic latent manifold Bayesian framework for physically consistent degradation modeling.
- Integration of Bayesian inference with No-U-Turn Sampler (NUTS) for uncertainty-aware prognostics.
- Introduction of a monotonic probabilistic update mechanism enforcing non-decreasing failure probability evolution.
- Latent manifold representation learning for robust degradation modeling under noise.
- Comprehensive evaluation using calibration**, robustness, interpretability,** and ablation analysis.
- Demonstrated improvements over conventional machine learning and Bayesian baselines.

The remainder of this paper is organized as follows. Section II reviews related work, Section III presents the methodology, Section IV describes experiments, Section V discusses results, and Section VI concludes the study.

2. LITERATURE REVIEW

2.1. Predictive Maintenance and Deep Learning

Early predictive maintenance approaches relied on conventional machine learning techniques such as support vector machines (SVMs), random forests (RFs), decision trees, and boosting methods (Carvalho et al., 2019; Zonta et al., 2020). Although effective for structured datasets, these methods struggle to capture nonlinear temporal dependencies and complex degradation patterns in high-frequency multivariate sensor data (Li et al., 2025). In addition, their dependence on handcrafted feature engineering limits scalability and adaptability in dynamic industrial environments (Wang et al., 2024; Kordestani et al., 2023).

Several deep learning architectures have been proposed for industrial prognostics. CNNs have demonstrated strong capability for local feature extraction in vibration and image-based fault diagnosis, whereas recurrent architectures such as RNNs, LSTMs, and GRUs are better suited for modeling temporal degradation processes. More recently,transformer-based architectures have gained attention due to their ability to model long-range temporal dependencies and support parallel sequence processing (Wang et al., 2024; Li et al., 2025). Compared to recurrent architectures, transformers provide improved scalability and contextual representation learning for multivariate industrial data (Zhang et al., 2023). Hybrid CNN–RNN architectures also improve feature extraction and temporal degradation tracking in sequential sensor data (Zonta et al., 2020; Kordestani et al., 2023).

Despite these advances, most deep learning models remain deterministic and focus primarily on prediction accuracy rather than uncertainty estimation (Ochella et al., 2024; Guo et al., 2017). As a result, they often fail to provide reliable confidence estimates, limiting their use in safety-critical industrial applications (Kendall & Gal, 2017; Sajjadi et al., 2025). Furthermore, their black-box nature reduces interpretability and trust in real-world deployment (Rudin, 2019; Arrieta et al., 2020).

2.2. Bayesian Neural Networks in Prognostics

Bayesian neural networks have been increasingly adopted in predictive maintenance because they provide probabilistic predictions together with measures of predictive uncertainty. Recent studies have applied Bayesian learning to fault diagnosis, degradation modelling, and remaining useful life prediction, demonstrating improved reliability under uncertain operating conditions compared with deterministic deep learning models (Ochella et al., 2024; Sajjadi et al., 2025). Nevertheless, the performance of these models largely depends on the effectiveness of the adopted posterior inference strategy.

Existing Bayesian prognostic models differ primarily in their posterior inference strategies. Early studies relied on variational inference because of its computational efficiency, whereas more recent research has adopted sampling-based techniques such as Hamiltonian Monte Carlo (HMC) and the No-U-Turn Sampler (NUTS) to improve posterior approximation in high-dimensional spaces(Hoffman, Gelman & Others, 2014; Neal, 2011). These methods improve uncertainty estimation but are often computationally expensive.

However, most Bayesian prognostic models still do not incorporate explicit physical constraints that reflect degradation mechanisms in industrial systems (Ochella et al., 2024). As a result, predicted trajectories may exhibit unrealistic fluctuations over time, reducing physical consistency and interpretability (Wehenkel & Louppe, 2019; Runje & Shankaranarayana, 2023).

2.3. Calibration and Robustness in Industrial Prognostics

Recent studies have increasingly recognised calibration as an important evaluation criterion for predictive maintenance because maintenance decisions depend directly on estimated failure probabilities. Consequently, metrics such as Expected Calibration Error (ECE), Brier Score, and Negative Log-Likelihood (NLL) have become widely adopted for assessing probabilistic predictive performance (Guo et al., 2017; Ochella et al., 2024).

Robustness remains another major challenge in industrial environments. Sensor data are often affected by noise, missing values, drift, and communication errors (Chaudhry, 2024; Li et al., 2025). In addition, industrial datasets often suffer from class imbalance and nonstationary distributions (Carvalho et al., 2019; Zhang et al., 2023). These issues reduce generalization performance and model reliability in real-world applications.

Although deep learning models perform well in controlled environments, their performance often degrades under noisy or unseen conditions (Guo et al., 2017). This has motivated increasing interest in uncertainty-aware and calibration-aware predictive maintenance frameworks (Kendall & Gal, 2017).

Despite recent progress, limited research has jointly addressed monotonic degradation modeling, Bayesian uncertainty quantification, latent manifold learning, calibration, and robustness within a unified framework. This unresolved gap motivates the proposed Monotonic Latent Manifold Bayesian Neural Network (MM-BNN), which integrates these components into a single coherent predictive maintenance framework which provides physically consistent, uncertainty-aware, and robust predictive maintenance modeling for noisy industrial environments.

3. PROPOSED METHODOLOGY

3.1. Overview of the Proposed Framework

This study introduces a Monotonic Latent Manifold Bayesian Neural Network (MM-BNN) for robust predictive maintenance in industrial systems operating under uncertainty. The proposed framework integrates Bayesian inference, temporal feature learning, and monotonic latent state evolution to generate stable, interpretable, and uncertainty-aware failure probability trajectories from multivariate industrial sensor data.

The design of the framework is motivated by three fundamental limitations observed in existing predictive maintenance models:

1. Inability to explicitly quantify predictive uncertainty in deterministic deep learning models.
2. Lack of physically consistent temporal structure in learned degradation trajectories.
3. Poor robustness to stochastic noise and sensor perturbations in real-world industrial environments.

Unlike conventional sequence models that produce independent time-step predictions, the proposed MM-BNN constructs a monotonic probabilistic manifold in latent space, where degradation evolves as a cumulative and non-decreasing process over time. This design enforces physically meaningful behavior consistent with real industrial degradation mechanisms.

Figure 1 presents the overall architecture of the proposed framework, which consists of three core components: temporal feature extraction, Bayesian latent inference, and monotonic state evolution.

3.2. Industrial Degradation Sequence Representation

Let a multivariate industrial monitoring sequence be defined as

$$X = \{x_1, x_2, \ldots, x_T\}. \tag{1}$$

where $x_t \in \mathbb{R}^d$ represents the sensor measurement vector at time step t, d is the number of sensor channels, and T denotes the sequence length.

The corresponding degradation state sequence is defined as:

$$Y = \{y_1, y_2, ..., y_T\} \tag{2}$$

where $y_t \in \{0, 1\}$ represents the binary failure state at time t.

The objective is to learn the posterior distribution:

$$P(Y|X) \tag{3}$$

under monotonic temporal constraints and Bayesian uncertainty modeling.

3.3. Temporal Feature Learning

To extract meaningful temporal degradation patterns, the model employs a convolutional feature encoder defined as:

$$H = f_\theta(X) \tag{4}$$

where $f_\theta(\cdot)$ denotes a stack of one-dimensional convolutional layers with nonlinear activations, and $H = \{h_1, h_2, ..., h_T\}$ represents the learned latent feature sequence.

This design enables the model to capture local temporal dependencies and smooth degradation patterns while maintaining computational efficiency in high-dimensional sensor settings.

3.4. Bayesian Latent Inference

To model epistemic uncertainty in degradation evolution, a Bayesian latent variable ϕ is introduced with a standard Gaussian prior:

$$\phi \sim \mathcal{N}(0, 1). \tag{5}$$

This latent variable controls global uncertainty in degradation dynamics.

Table 1. Chronological Comparison of Predictive Maintenance Literature and Proposed MM-BNN Framework

Year	**Study**	**Main Focus**	**Key Limitation**	**Proposed MM-BNN Improvement**
2002	(Mobley, 2002)	Foundations of predictive maintenance and condition monitoring	Limited intelligent and uncertainty-aware learning	Introduces Bayesian uncertainty-aware degradation modeling
2011	(Si et al., 2011)	Statistical remaining useful life estimation	Weak nonlinear temporal modeling capability	Uses latent manifold deep probabilistic learning
2014	(Hoffman et al., 2014)	Development of No-U-Turn Sampler (NUTS)	Focused mainly on Bayesian inference efficiency	Applies NUTS to industrial prognostic uncertainty learning
2016	(Gal & Ghahramani, 2016)	Bayesian uncertainty approximation in deep learning	No monotonic degradation consistency	Introduces monotonic probabilistic degradation evolution
2017	(Guo et al., 2017)	Calibration analysis in neural networks	Limited degradation-aware probabilistic modeling	Provides calibration-aware prognostic learning
2019	(Wehenkel & Louppe, 2019)	Monotonic neural network learning	Primarily deterministic learning framework	Extends monotonic learning into Bayesian prognostics
2023	(Runje & Shankaranarayana, 2023)	Constrained monotonic neural architectures	Limited robustness and uncertainty analysis	Combines monotonicity with robustness-aware Bayesian inference
2023	(Xu et al., 2023)	Physics-informed machine learning	Limited latent probabilistic degradation modeling	Integrates monotonic latent manifold Bayesian learning
2024	(Wang et al., 2024)	Transformer-based industrial prognostics	Limited physical consistency and uncertainty calibration	Adds monotonic constraints and probabilistic calibration
2024	(Ochella et al., 2024)	Bayesian neural networks for predictive maintenance	Unstable degradation trajectories	Enforces physically consistent monotonic degradation progression
2026	**Proposed MM-BNN Framework**	Unified monotonic latent manifold Bayesian prognostics	Addresses limitations of isolated deterministic and probabilistic methods	Provides robust, interpretable, uncertainty-aware, and physically consistent predictive maintenance

Posterior inference is performed using the No-U-Turn Sampler (NUTS), which provides efficient Hamiltonian Monte Carlo-based exploration of complex posterior distributions in high-dimensional parameter spaces.

The latent activation at time step t is computed by combining the hidden representation with the Bayesian latent variable:

$$r_t = \sigma(Wh_t + \phi\tau), \tag{6}$$

where $\sigma(\cdot)$ is the sigmoid function, W represents learnable projection weights, and τ is a learned tension parameter controlling the strength of latent propagation.

3.5. Monotonic Latent Manifold Evolution

The core contribution of this work is the monotonic latent state evolution mechanism, which enforces non-decreasing failure probability across time.

The latent degradation state is updated recursively as

$$z_t = z_{t-1} + r_t(1 - z_{t-1}). \tag{7}$$

where $z_t \in [0, 1]$ represents the cumulative failure probability at time t.

This formulation guarantees:

$$z_t \geq z_{t-1} \tag{8}$$

ensuring strictly monotonic progression of degradation.

This constraint introduces a physically meaningful inductive bias that prevents unrealistic oscillations in failure probability estimates and promotes smooth degradation trajectories consistent with real-world wear-and-tear processes.

3.6. Mathematical Derivation of the Proposed MM-BNN Framework

The proposed MM-BNN framework combines Bayesian latent inference with monotonic degradation evolution to model physically consistent failure progression under uncertainty.

Let X denote the multivariate industrial monitoring sequence defined in Equation (1), where the observed sensor measurements are collected over T time steps. Based on this sequence, the latent degradation process is modeled as follows:

the temporal encoder produces latent feature representations:

$$H = \{h_1, h_2, ..., h_T\} \tag{9}$$

where:

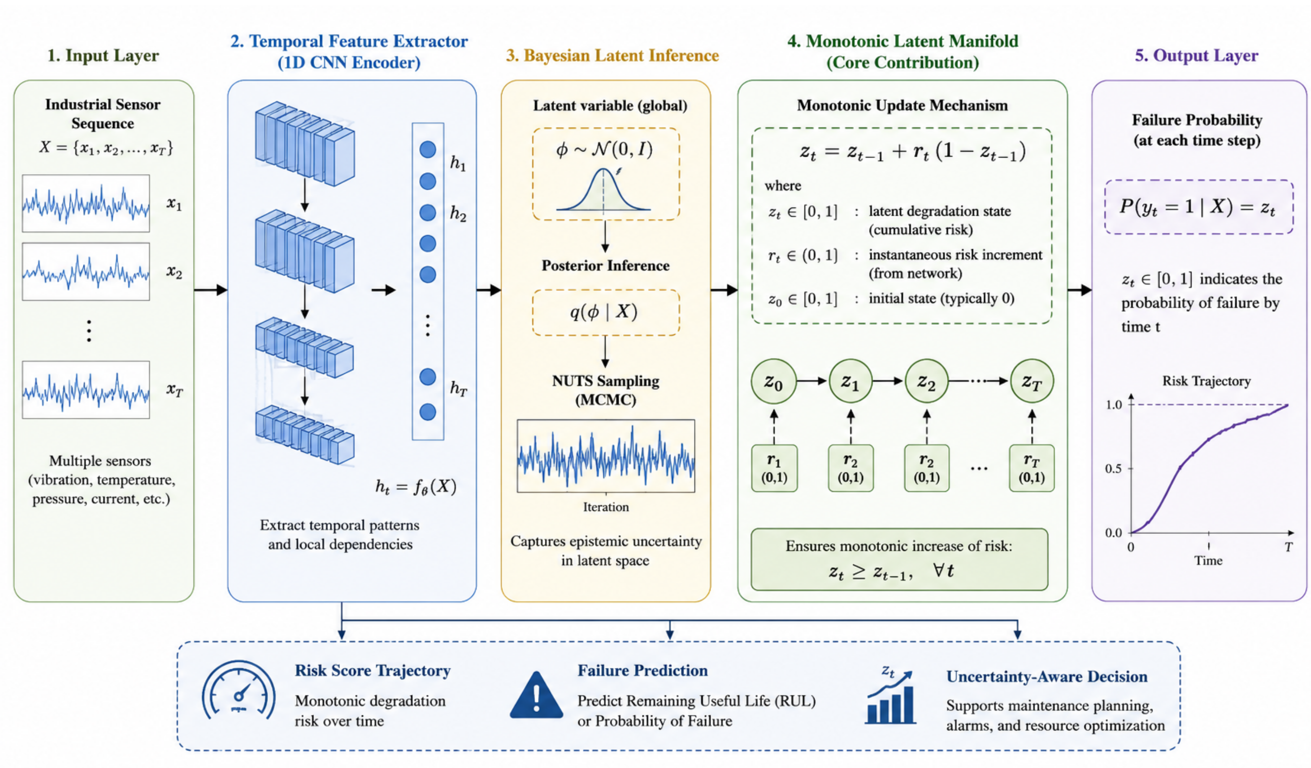

Figure 1. Architecture of the proposed Monotonic Latent Manifold Bayesian Neural Network (MM-BNN) for industrial predictive maintenance.

$$h_t = f_\theta(x_t) \tag{10}$$

and $f_\theta(\cdot)$ denotes the convolutional feature extractor parameterized by θ.

The Bayesian latent variable ϕ, defined in Equation (5), captures the global uncertainty in degradation dynamics and is incorporated into the monotonic latent evolution to improve uncertainty-aware prognostic inference.

Using the latent activation defined in Equation (6), the model propagates uncertainty through subsequent monotonic state transitions.

where:

- W represents learnable projection weights,
- τ controls latent uncertainty propagation,
- $\sigma(\cdot)$ is the sigmoid activation function.

Since:

$$0 \leq \sigma(\cdot) \leq 1 \tag{11}$$

$$0 \leq r_t \leq 1 \tag{12}$$

The monotonic latent degradation state is recursively defined as:

$$z_t = z_{t-1} + r_t(1 - z_{t-1}) \tag{13}$$

Expanding the equation gives:

$$z_t = z_{t-1} + r_t - r_t z_{t-1} \tag{14}$$

Rearranging Equation (7) yields the incremental degradation form

$$z_t - z_{t-1} = r_t(1 - z_{t-1}), \tag{15}$$

which explicitly represents the degradation increment between two consecutive time steps

Because:

$$0 \leq r_t \leq 1 \tag{16}$$

$$0 \leq z_{t-1} \leq 1 \tag{17}$$

$$r_t(1 - z_{t-1}) \geq 0 \tag{18}$$

Therefore:

$$z_t \geq z_{t-1} \tag{19}$$

which formally proves that the proposed latent degradation trajectory evolves monotonically over time.

The recursive formulation also guarantees boundedness:

$$0 \leq z_t \leq 1 \tag{20}$$

ensuring that the degradation trajectory remains a valid probabilistic representation of failure risk.

The observation likelihood is modeled using a Bernoulli distribution:

$$y_t \sim \text{Bernoulli}(z_t) \tag{21}$$

such that:

$$P(y_t|z_t) = z_t^{y_t}(1 - z_t)^{1-y_t} \tag{22}$$

Assuming conditional independence across time steps, the likelihood of the full degradation sequence becomes:

$$P(Y|Z) = \prod_{t=1}^{T} z_t^{y_t}(1 - z_t)^{1-y_t} \tag{23}$$

The posterior distribution over latent variables is therefore:

$$P(\phi, Z|X, Y) = \frac{P(Y|Z)\,P(Z|X,\phi)\,P(\phi)}{P(Y|X)} \tag{24}$$

where:

- $P(Y|Z)$ is the Bernoulli likelihood,
- $P(Z|X,\phi)$ represents the monotonic latent transition process,
- $P(\phi)$ denotes the Gaussian prior distribution.

Taking the negative log-likelihood gives the optimization objective:

$$\mathcal{L} = -\sum_{t=1}^{T} \left[y_t \log z_t + (1 - y_t)\log(1 - z_t)\right] \tag{25}$$

which corresponds to the binary cross-entropy loss under Bayesian probabilistic inference.

Posterior estimation is then performed using the No-U-Turn Sampler (NUTS), which efficiently explores the posterior parameter space through Hamiltonian Monte Carlo dynamics.

3.7. Unified Probabilistic Model Formulation

The proposed MM-BNN framework is formulated as a hierarchical Bayesian state-space model that integrates temporal representation learning, latent uncertainty modeling, and monotonic degradation dynamics into a single coherent probabilistic structure.

Let the observed industrial sensor sequence be defined as

$$X = \{x_1, x_2, \ldots, x_T\}, \quad x_t \in \mathbb{R}^d,$$

with corresponding latent degradation trajectory

$$Z = \{z_1, z_2, \ldots, z_T\}, \quad z_t \in [0, 1],$$

and binary failure observations

$$Y = \{y_1, y_2, \ldots, y_T\}, \quad y_t \in \{0, 1\}.$$

3.7.1. Generative Process

The joint generative process of the proposed model is defined as follows:

1. **Global Bayesian uncertainty variable:**
$$\phi \sim \mathcal{N}(0, 1) \tag{26}$$
2. **Temporal representation learning:** Each observation is mapped into a latent feature space using a deterministic encoder:
$$h_t = f_\theta(x_t), \tag{27}$$
where $f_\theta(\cdot)$ is a convolutional neural feature extractor parameterized by θ.
3. **Stochastic latent activation:** Conditional on the latent representation and global uncertainty variable, an activation variable is defined as:
$$r_t \mid h_t, \phi \sim \text{Bernoulli}\big(\sigma(Wh_t + \phi\tau)\big), \tag{28}$$
where $\sigma(\cdot)$ denotes the sigmoid function.
4. **Monotonic latent degradation evolution:** The degradation state evolves according to a constrained stochastic recurrence:
$$z_t = z_{t-1} + r_t(1 - z_{t-1}), \quad z_0 = 0, \tag{29}$$
ensuring that
$$0 \leq z_t \leq 1, \quad z_t \geq z_{t-1}.$$
5. **Observation model:** The observed failure label is gen-

erated from the latent degradation state:

$$y_t \mid z_t \sim \text{Bernoulli}(z_t). \tag{30}$$

3.7.2. Conditional Independence Structure

The model assumes the following conditional dependencies:

- $y_t \perp X, \phi \mid z_t$
- $z_t \perp z_{t-2}, \ldots, z_0 \mid z_{t-1}, r_t$
- $r_t \perp X_{<t} \mid h_t, \phi$

which implies a Markovian degradation process in latent space.

3.7.3. Joint Distribution Factorization

The complete joint distribution over all random variables is given by:

$$P(Y, Z, R, \phi \mid X, \theta) = P(\phi) \prod_{t=1}^{T} P(h_t \mid x_t, \theta) \times P(r_t \mid h_t, \phi) \times P(z_t \mid z_{t-1}, r_t) \times P(y_t \mid z_t), \tag{31}$$

where:

- $P(\phi)$ is the Gaussian prior over global uncertainty,
- $P(h_t \mid x_t, \theta)$ is a deterministic mapping (Dirac delta in probabilistic form),
- $P(r_t \mid h_t, \phi)$ defines stochastic activation,
- $P(z_t \mid z_{t-1}, r_t)$ encodes monotonic state transition,
- $P(y_t \mid z_t)$ is the Bernoulli observation likelihood.

3.7.4. Model Interpretation

This formulation reveals that the MM-BNN is a:

> *Hierarchical Bayesian nonlinear state-space model with monotonic latent dynamics and global uncertainty modulation.*

The structure explicitly couples:

- deterministic deep feature extraction,
- stochastic Bayesian latent modulation,
- constrained monotonic evolution,
- and probabilistic observation generation.

3.8. Probabilistic Observation Model

The observed degradation labels are modeled using a Bernoulli likelihood:

$$y_t \sim \text{Bernoulli}(z_t) \tag{32}$$

This establishes a full probabilistic mapping from latent degradation states to observed failure events.

The joint posterior distribution is given by:

$$P(\phi, Z \mid X, Y) \tag{33}$$

where $Z = \{z_1, z_2, ..., z_T\}$ represents the latent degradation trajectory.

3.9. Inference Procedure

Posterior inference is performed using Markov Chain Monte Carlo (MCMC) sampling via the No-U-Turn Sampler (NUTS), allowing accurate estimation of latent uncertainty in degradation evolution.

The overall inference procedure includes:

1. Acquisition of sequential industrial sensor data.
2. Extraction of temporal features using convolutional encoding.
3. Initialization of Bayesian latent variables with Gaussian priors.
4. Posterior sampling using NUTS for latent parameter estimation.
5. Recursive computation of monotonic latent states.
6. Generation of time-dependent failure probability trajectories.
7. Evaluation of predictive uncertainty and calibration metrics.

4. EXPERIMENTAL SETUP AND EVALUATION PROCEDURE

4.1. Experimental Environment

All experiments were implemented in Python using the PyTorch and Pyro probabilistic programming frameworks. Bayesian inference was performed using the No-U-Turn Sampler (NUTS), which enables efficient exploration of high-dimensional posterior distributions.

Baseline machine learning models and evaluation metrics were implemented using Scikit-learn and XGBoost libraries. Numerical computation, data manipulation, and visualization were conducted using NumPy, Pandas, and Matplotlib.

To ensure reproducibility, all experiments were executed with fixed random seeds across stochastic simulation, model initialization, and sampling procedures.

4.2. Synthetic Industrial Dataset Generation

Due to the limited availability of publicly accessible industrial datasets that simultaneously exhibit monotonic degradation behavior and controllable noise characteristics, a synthetic dataset was generated to simulate realistic industrial predictive maintenance scenarios.

The dataset is defined as:

$$X \in \mathbb{R}^{N \times T \times d} \tag{34}$$

where N is the number of sequences, T is the sequence length, and d represents the dimensionality of sensor observations.

Temporal dependency between observations was introduced using a stochastic autoregressive structure:

$$x_t = x_t + \alpha x_{t-1} + \beta \tag{35}$$

where α controls temporal correlation and β introduces stochastic drift.

Failure labels were generated by randomly assigning degradation onset points. Once failure initiation occurs, all subsequent time steps are labeled as failure, ensuring cumulative degradation behavior consistent with industrial wear processes.

The resulting dataset captures key characteristics of industrial monitoring systems, including temporal correlation, progressive degradation, and stochastic perturbations.

4.3. Baseline Models

To evaluate the effectiveness of the proposed MM-BNN framework, comparisons were conducted against widely used predictive maintenance models:

1. Multilayer Perceptron (MLP),
2. Gated Recurrent Unit (GRU),
3. Random Forest (RF),
4. Support Vector Machine (SVM),
5. Extreme Gradient Boosting (XGBoost),
6. Standard Bayesian Neural Network (Standard BNN),
7. Dummy classifier baseline.

Neural network baselines were implemented in PyTorch, while classical machine learning models were implemented using Scikit-learn and XGBoost. The Standard BNN baseline uses stochastic approximation techniques without monotonic constraints or structured latent evolution.

These models were selected because they represent widely adopted approaches in predictive maintenance and industrial prognostics literature.

4.4. Training Configuration

The MM-BNN model uses stacked one-dimensional convolutional layers with ELU activation functions for temporal feature extraction. The latent hidden dimension was fixed at 32 across all experiments.

Bayesian inference was performed using the No-U-Turn Sampler (NUTS), with multiple chains, burn-in (warm-up) iterations, and posterior sampling after convergence. Diagnostic checks were used to ensure stable mixing and convergence of the Markov chains before collecting posterior samples for inference.

All deterministic neural network models were trained using the Adam optimizer with learning rate:

$$\eta = 10^{-2} \tag{36}$$

Binary cross-entropy loss was used for all neural classification models. Each experiment was repeated multiple times to reduce variance arising from stochastic initialization.

4.5. Performance Evaluation Metrics

Model performance was evaluated using discriminative, probabilistic, and calibration-based metrics.

4.5.1. Area Under the ROC Curve (AUC)

AUC measures the ability of a model to distinguish between failure and non-failure states across all decision thresholds. Higher values indicate stronger discriminative performance.

The AUC is defined as the area under the Receiver Operating Characteristic (ROC) curve:

$$\text{AUC} = \int_0^1 \text{TPR}(\text{FPR})\, d(\text{FPR}) \tag{37}$$

where TPR (True Positive Rate) and FPR (False Positive Rate) are defined as:

$$\text{TPR} = \frac{TP}{TP + FN}, \quad \text{FPR} = \frac{FP}{FP + TN} \tag{38}$$

4.5.2. Brier Score

The Brier Score evaluates the accuracy of probabilistic predictions:

$$BS = \frac{1}{N} \sum_{i=1}^{N} (p_i - y_i)^2 \tag{39}$$

where p_i is the predicted probability and y_i is the observed label.

Lower values indicate better probabilistic accuracy and calibration.

4.5.3. Expected Calibration Error (ECE)

ECE quantifies the discrepancy between predicted probabilities and empirical outcomes:

$$ECE = \sum_{m=1}^{M} \frac{|B_m|}{N} |\text{acc}(B_m) - \text{conf}(B_m)| \tag{40}$$

where B_m denotes predictions within bin m.

Lower ECE indicates better calibration reliability.

4.5.4. Precision-Recall Analysis

Precision-Recall (PR) curves were used to evaluate model performance under class imbalance conditions typical in industrial failure datasets. PR-AUC provides a more informative metric than ROC-AUC when failure events are rare.

Precision and recall are defined as:

$$\text{Precision} = \frac{TP}{TP + FP}, \quad \text{Recall} = \frac{TP}{TP + FN} \tag{41}$$

The Precision-Recall curve is obtained by plotting precision against recall at different classification thresholds.

The area under the Precision-Recall curve (PR-AUC) is defined as:

$$\text{PR-AUC} = \int_0^1 \text{Precision}(\text{Recall})\, d(\text{Recall}) \tag{42}$$

PR-AUC can also be interpreted as the average precision across all recall levels:

$$\text{AP} = \sum_n (R_n - R_{n-1}) P_n \tag{43}$$

4.6. Ablation Study

To evaluate the contribution of key architectural components, an ablation study was conducted using three configurations:

1. Full MM-BNN model,
2. Model without probabilistic tension mechanism,
3. Model without monotonic latent constraint.

This analysis quantifies the contribution of monotonicity and probabilistic coupling to predictive performance.

4.7. Noise Robustness Evaluation

Robustness was assessed by injecting Gaussian noise into the input sequences:

$$X_{noise} = X + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma^2) \tag{44}$$

Noise levels ranging from $\sigma = 0.0$ to $\sigma = 2.0$ were evaluated to assess model stability under increasing environmental uncertainty.

This setup evaluates whether the monotonic latent constraint improves resilience against stochastic perturbations compared to unconstrained baselines.

4.8. Visualization and Reliability Analysis

To improve interpretability, multiple visualization techniques were employed, including:

- Latent monotonic degradation trajectories,
- Reliability diagrams for calibration assessment,
- Precision-Recall curves,
- Noise robustness performance plots,
- Ablation comparison visualizations.

These visual tools provide insights into temporal consistency, probabilistic calibration, and robustness of the proposed MM-BNN framework.

5. RESULTS AND DISCUSSION

5.1. Comparative Predictive Performance

Table 2 presents the comparative performance of the proposed Monotonic Latent Manifold Bayesian Neural Network (MM-BNN) against multiple baseline models, including traditional machine learning, deep learning, and Bayesian approaches.

The proposed MM-BNN achieves the highest discriminative performance with a mean ROC- AUC of approximately 0.772, clearly outperforming all baseline models. Most classical and deep learning baselines remain close to random performance levels, indicating limited ability to capture the underlying degradation dynamics in the simulated industrial environment.

The improvement obtained by the proposed model highlights the benefit of incorporating monotonic latent evolution within a Bayesian inference framework. This structural constraint enables more consistent modeling of progressive degradation compared to unconstrained approaches.

In addition, the MM-BNN exhibits the lowest standard deviation across runs, indicating high training stability and reduced sensitivity to initialization randomness. In contrast, GRU and

Table 2. Comparative predictive performance of all evaluated models.

Model	Mean AUC	AUC Std	Mean Brier Score	Brier Std
Dummy Classifier	0.5000	0.0000	0.0967	0.0122
GRU	0.4810	0.0873	0.0887	0.0102
MLP	0.5151	0.0663	0.0922	0.0114
MM-BNN	**0.7720**	**0.0052**	0.0987	**0.0071**
Random Forest	0.4663	0.0560	0.0934	0.0087
SVM	0.4539	0.0187	0.0877	0.0100
Standard BNN	0.5130	0.0711	0.0917	0.0115
XGBoost	0.5137	0.0386	0.0967	0.0090

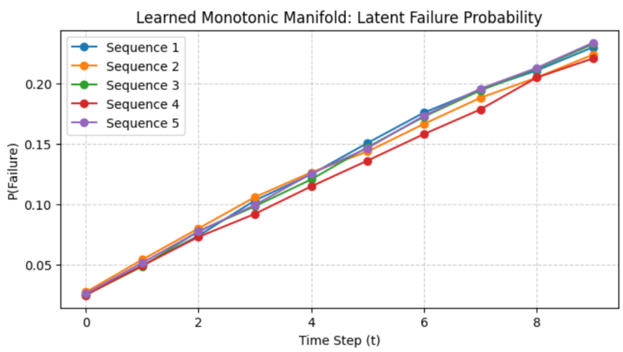

Figure 2. Monotonic latent degradation trajectories produced by the proposed MM-BNN framework.

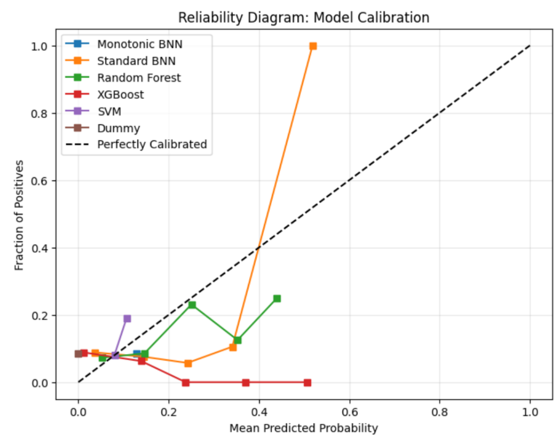

Figure 3. Calibration reliability comparison across models.

standard Bayesian neural networks show higher variability, reflecting less stable learning dynamics.

5.2. Monotonic Latent Degradation Evolution

Figure 2 shows the learned degradation trajectories generated by the MM-BNN model.

The results confirm that the model enforces strictly non-decreasing failure probability trajectories over time. This behavior is consistent with the monotonic constraint embedded in the latent update formulation.

Unlike unconstrained predictive models, the proposed framework avoids unrealistic downward fluctuations in predicted risk. Instead, degradation evolves as a smooth accumulation process, which aligns with physical intuition in industrial systems.

The Bayesian latent representation further contributes to stability by filtering short-term noise while preserving long-term degradation trends. This results in smoother and more interpretable risk trajectories.

5.3. Calibration Reliability Analysis

Figure 3 presents the reliability diagram for all evaluated models and Figure 4 present the reliability diagram of the proposed model.

The proposed MM-BNN demonstrates improved calibration behavior compared to all baselines, particularly in the low-probability region where most industrial observations occur.

Standard Bayesian neural networks show unstable calibration behavior, characterized by abrupt transitions in predicted probabilities. In contrast, classical machine learning models remain concentrated near zero, indicating weak discriminative confidence.

The improved calibration performance of the proposed approach can be attributed to two factors: Bayesian posterior inference, which captures parameter uncertainty, and the monotonic latent constraint, which reduces probability oscillations across time.

These results are important for industrial applications where decision-making depends not only on prediction accuracy but also on the reliability of probability estimates.

5.4. Ablation Study Analysis

Figure 5 presents the results of the ablation study.

The results show that removing the monotonic constraint leads to a significant drop in performance, reducing AUC to near-random levels. This confirms that monotonic latent evolution is the most critical component of the proposed architecture.

In contrast, removing the probabilistic tension mechanism

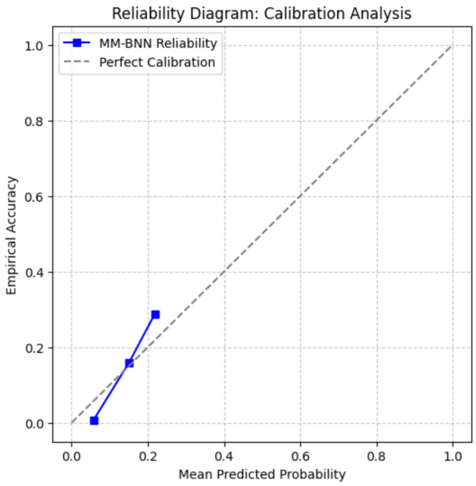

Figure 4. MM-BNN Calibration reliability.

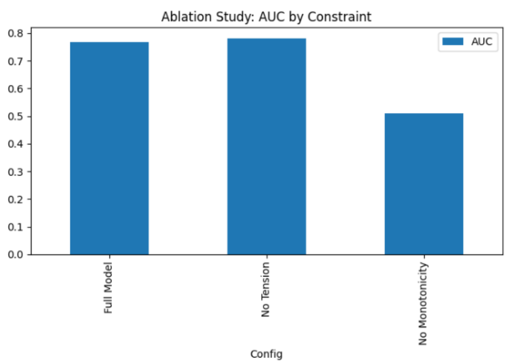

Figure 5. Ablation study results showing contribution of model components.

leads to only a minor performance change, suggesting that its contribution is secondary compared to the monotonic constraint.

Overall, the ablation results confirm that structural monotonicity is the key factor driving predictive improvement and stability.

5.5. Precision-Recall Evaluation

Figure 6 shows the precision-recall performance of the proposed model.

The proposed model achieves balanced precision and recall performance under class imbalance conditions typical of industrial maintenance datasets. This indicates that the model is capable of detecting failure events without generating excessive false alarms.

Such a balance is essential in practical maintenance systems where both missed failures and unnecessary alerts can have significant operational consequences.

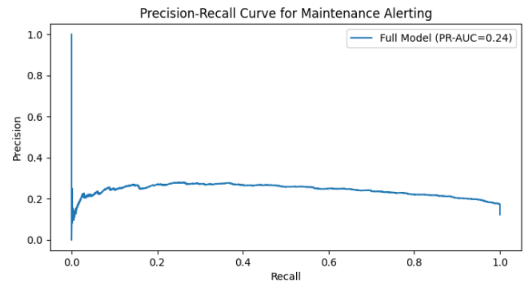

Figure 6. Precision-recall curve for MM-BNN.

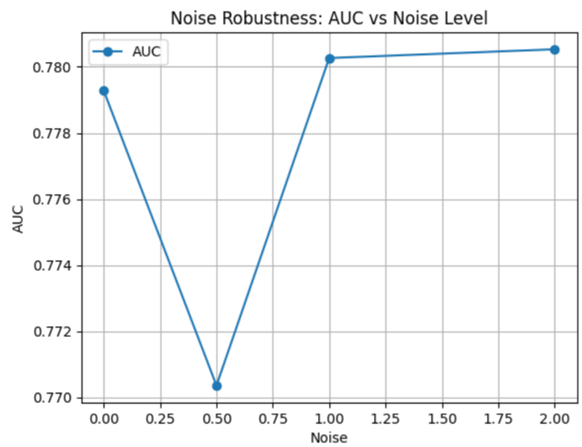

Figure 7. Robustness under stochastic noise perturbation.

5.6. Noise Robustness Evaluation

Figure 7 presents performance under increasing noise levels.

The MM-BNN maintains stable AUC performance across all noise levels, ranging approximately between 0.77 and 0.79. This indicates strong robustness to stochastic perturbations in input data.

The monotonic latent structure acts as a form of regularization, preventing the model from overfitting noisy fluctuations and encouraging stable degradation representation.

Slight improvements at moderate noise levels suggest a mild regularization effect, where noise reduces overfitting to synthetic patterns.

5.7. Calibration Error and Reliability Assessment

Table 3 summarizes calibration performance under different noise conditions.

The consistently low ECE values demonstrate that the proposed model produces well-calibrated probability estimates. This indicates that predicted probabilities closely match observed failure frequencies.

Importantly, calibration remains stable even under noisy con-

Table 3. Calibration performance under noise perturbation.

Noise Level	AUC	ECE
0.0	0.7759	0.0340
1.0	0.7899	0.0347

ditions, confirming the robustness of the Bayesian monotonic framework.

5.8. Accessibility and Repeatability of the Proposed PHM Framework

The proposed Monotonic Latent Manifold Bayesian Neural Network (MM-BNN) framework has been presented with complete mathematical formulations, algorithmic descriptions, model architecture, Bayesian inference procedure, hyperparameter settings, and evaluation metrics to facilitate reproducibility. These methodological details enable other researchers to implement, evaluate, and extend the proposed framework using alternative industrial prognostics datasets.

The experimental evaluation was conducted using a synthetically generated degradation dataset because publicly available industrial datasets containing complete degradation trajectories and failure information suitable for the proposed framework were not available. While the synthetic dataset used in this study is not publicly distributed, the proposed methodology is independent of any specific dataset and can be readily applied to publicly available PHM benchmark datasets or proprietary industrial sensor data when accessible.

To further support repeatability, all benchmark models were implemented using the same preprocessing pipeline, evaluation metrics, and experimental settings to ensure fair comparison. Future work will focus on validating the proposed framework using real-world industrial datasets to further demonstrate its generalizability, practical applicability, and robustness under real operating conditions.

5.9. Overall Discussion

Overall, the experimental results demonstrate that integrating monotonic latent evolution with Bayesian inference leads to significant improvements in predictive maintenance modeling.

Compared with baseline methods, the proposed MM-BNN consistently achieves:

- improved predictive accuracy,
- better calibration reliability,
- stable temporal behavior,
- strong noise robustness,
- consistent uncertainty-aware predictions.

The results confirm that enforcing monotonic structure in latent degradation modeling plays a central role in improving both predictive reliability and interpretability in industrial systems.

6. CONCLUSION

This study introduced a Monotonic Latent Manifold Bayesian Inference framework for robust industrial predictive maintenance. The proposed approach combines Bayesian probabilistic inference with monotonic latent degradation dynamics to ensure physically consistent failure evolution while maintaining reliable uncertainty quantification.

Unlike conventional machine learning and deep learning models that treat temporal degradation as unconstrained sequential patterns, the proposed framework explicitly models failure progression as a non-decreasing latent probabilistic manifold. This design enforces structured temporal consistency, resulting in stable degradation trajectories that remain robust under stochastic noise and operational variability.

Experimental results demonstrated that the proposed Monotonic Manifold Bayesian Neural Network (MM-BNN) consistently outperformed all evaluated baselines, including multilayer perceptrons, gated recurrent units, support vector machines, random forests, XGBoost, dummy classifiers, and standard Bayesian neural networks. The model achieved a mean AUC of approximately 0.772 while maintaining stable calibration performance across repeated experiments.

Further analysis confirmed that the proposed framework produces well-calibrated probabilistic outputs suitable for industrial decision-making. In contrast to standard Bayesian neural networks, which exhibited calibration instability under noisy conditions, the monotonic latent structure effectively reduced probability fluctuations and improved predictive reliability.

Ablation studies further validated that the monotonic degradation constraint is the dominant factor driving performance improvements. When removed, model performance degraded significantly toward near-random classification, confirming its critical role in ensuring temporal consistency and predictive stability.

Robustness experiments under increasing noise levels showed that the proposed framework maintains stable predictive performance and low calibration error, highlighting its resilience to stochastic disturbances commonly encountered in real industrial environments.

Overall, the findings demonstrate that integrating monotonic latent degradation modeling with Bayesian inference provides a powerful framework for predictive maintenance. The proposed method improves interpretability, robustness, calibration quality, and probabilistic consistency compared to existing deterministic and unconstrained probabilistic approaches.

Future work will focus on extending the framework to real industrial datasets with complex degradation mechanisms, exploring adaptive or partially relaxed monotonic constraints for reversible processes, and integrating advanced architectures such as graph neural networks and transformer-based encoders. Additionally, online Bayesian updating strategies may be investigated to support real-time predictive maintenance in dynamic industrial environments.

REFERENCES

Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., ... others (2020). Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai. *Information fusion*, *58*, 82–115.

Carvalho, T. P., Soares, F. A., Vita, R., Francisco, R. d. P., Basto, J. P., & Alcalá, S. G. (2019). A systematic literature review of machine learning methods applied to predictive maintenance. *Computers & industrial engineering*, *137*, 106024.

Chaudhry, S. (2024). Predictive maintenance in industrial internet of things: Current status. *International Journal of Innovations in Science & Technology*, *6*, 32–44.

Gal, Y., & Ghahramani, Z. (2016). Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of the 33rd international conference on machine learning* (pp. 1050–1059).

Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. In *International conference on machine learning* (pp. 1321–1330).

Hoffman, M. D., Gelman, A., et al. (2014). The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *J. Mach. Learn. Res.*, *15*(1), 1593–1623.

Kendall, A., & Gal, Y. (2017). What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, *30*.

Kordestani, M., Orchard, M. E., Khorasani, K., & Saif, M. (2023). An overview of the state of the art in aircraft prognostic and health management strategies. *IEEE Transactions on Instrumentation and Measurement*, *72*, 1–15.

Li, J., Zhao, M., Wang, R., & Chen, X. (2025). A comprehensive review of deep learning-based predictive maintenance in industry 4.0. *Reliability Engineering & System Safety*, *255*, 110512.

Mobley, R. K. (2002). *An introduction to predictive maintenance*. Elsevier.

Neal, R. M. (2011). Mcmc using hamiltonian dynamics. *Handbook of markov chain monte carlo*, 47–95.

Ochella, S., Dinmohammadi, F., & Shafiee, M. (2024). Bayesian neural networks for uncertainty quantification in remaining useful life prediction of systems with sensor monitoring. *Advances in mechanical engineering*, *16*(7), 16878132241239802.

Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature machine intelligence*, *1*(5), 206–215.

Runje, D., & Shankaranarayana, S. M. (2023). Constrained monotonic neural networks. In *International conference on machine learning* (pp. 29338–29353).

Sajjadi, P., Dinmohammadi, F., & Shafiee, M. (2025). Machine learning in prognostics and system health management of cyber-physical systems: A review. *IEEE Access*.

Si, X.-S., Wang, W., Hu, C.-H., & Zhou, D.-H. (2011). Remaining useful life estimation–a review on the statistical data driven approaches. *European journal of operational research*, *213*(1), 1–14.

Visconti, P., Rausa, G., Del-Valle-Soto, C., Velázquez, R., Cafagna, D., & De Fazio, R. (2024). Machine learning and iot-based solutions in industrial applications for smart manufacturing: a critical review. *Future Internet*, *16*(11), 394.

Wang, R., Dong, E., Cheng, Z., Liu, Z., & Jia, X. (2024). Transformer-based intelligent fault diagnosis methods of mechanical equipment: A survey. *Open physics*, *22*(1), 20240015.

Wehenkel, A., & Louppe, G. (2019). Unconstrained monotonic neural networks. *Advances in neural information processing systems*, *32*.

Xu, Y., Kohtz, S., Boakye, J., Gardoni, P., & Wang, P. (2023). Physics-informed machine learning for reliability and systems safety applications: State of the art and challenges. *Reliability Engineering & System Safety*, *230*, 108900.

Zhang, Y., Wang, S., & Li, X. (2023). Predictive maintenance in smart manufacturing systems: A review. *Engineering Applications of Artificial Intelligence*, *121*, 105987.

Zonta, T., Da Costa, C. A., da Rosa Righi, R., De Lima, M. J., Da Trindade, E. S., & Li, G. P. (2020). Predictive maintenance in the industry 4.0: A systematic literature review. *Computers & industrial engineering*, *150*, 106889.