Computers, Materials & Continua

Tech Science Press

https://doi.org/10.32604/cmc.2026.083649

ARTICLE

Counterfactual Enabled Neuro-Symbolic Digital Twins for Intelligent Industrial Maintenance

Nada Alzaben[1], **Muhammad I. Khan**[2], **Hafeez Ur Rehman Siddiqui**[3], **Abeer Rashad Mirdad**[4] and **Saeed Ali Bahaj**[5,*]

[1]College of Computer and Information Sciences, Princess Nourah Bint Abdulrahman University, Riyadh, Saudi Arabia
[2]Artificial Intelligence & Data Analytics Lab, CCIS Prince Sultan University, Riyadh, Saudi Arabia
[3]Institute of Computer Science, Khwaja Fareed University of Engineering and Information Technology, Abu Dhabi Road, Rahim Yar Khan, Punjab, Pakistan
[4]Department of Information Systems, Prince Sultan University, Riyadh, Saudi Arabia
[5]MIS Department College of Business Administration, Prince Sattam Bin Abdulaziz University, Al-Kharj, Saudi Arabia

*Corresponding Author: Saeed Ali Bahaj. Email: s.bahaj@psau.edu.sa

Received: 08 April 2026; Accepted: 21 May 2026; Published: 23 July 2026

ABSTRACT: Industrial predictive maintenance is a critical challenge in modern manufacturing, where unexpected equipment failures cause significant economic losses through downtime, repair costs, and disrupted production. Conventional maintenance approaches, whether reactive or schedule-based, are becoming inadequate to manage the high-dimensional sensor information of the IoT-enabled machineries. The paper presents a novel hybrid neuro-symbolic digital twin that builds upon Remaining Useful Life (RUL) estimation by combining temporal transformers, physics-informed constraints, and counterfactual reasoning. The model integrates complementary approaches into a single and interpretable predictive system. A temporal transformer backbone is a model of long-range dependencies in multivariate sensor time-series data allowing the detection of gradual patterns of degradation that are frequently overlooked by traditional models of recurrence. The learning objective contains physics-informed constraints that guarantee that predictions are consistent with the principles of thermodynamic, mechanical, and material fatigue, connecting data-driven learning to domain knowledge. A Conditional Variational Autoencoder (CVAE) produces counterfactual failure events, simulating alternative histories of operation, under hypothetical conditions. This mechanism increases data diversity, enables interpretable diagnostics, and reinforces neuro-symbolic reasoning. A method of improving maintenance policies with the help of multi-objective reinforcement learning (MORL) is applied to minimize downtime, unnecessary maintenance, and equipment life. The results of the experiments on 24,042 sensor measurements of CNC machines, pumps, compressors, and robotic arms show good results. The framework has attained RMSE of 21.52 h and a 0.918 score, R^2 which is 25.1 percent better as compared to the baseline models. The accuracy of prediction of failures was 94.2 percent, and the maintenance policies were optimized to achieve a reduction in equipment failures by 51.7 percent compared to the rule-based scheduling, which is based on the fact that CVAE-generated counterfactual state transitions were used to train the Q-learning agent, which indicates the practicality of the framework.

KEYWORDS: Digital twin; predictive maintenance; neuro-symbolic AI; counterfactual reasoning; remaining useful life; temporal transformer

1 Introduction

Accidental equipment breakdowns cost manufacturing businesses as much as 50 billion dollars each year and even single failures at high-throughput manufacturing facilities cost up to 260,000 dollars per hour [1,2]. The proliferation of Industrial Internet of Things (IIoT) sensors and edge computing makes it possible to engage in continuous monitoring [3–5], but converting such sensor streams to sound and actionable maintenance decisions is a consequential and open issue.

The process of degradation is nonlinear, and includes thermal stress, fatigue crack propagation, bearing wear, and lubrication failure, which cannot be fully represented by purely data-driven models without domain knowledge [6,7]. In addition to the quality of prediction, maintenance decisions are multi-objective in nature: unplanned failures will cost resources in terms of repair, loss of production, and safety risk, and overly conservative preventive maintenance will waste resources; the optimum trade-off depends on the type of machine, the mode of operation, the load history, and the business environment [4,8,9].

The existing predictive maintenance methods do not sufficiently complement each other. Models based on physics, inspired by Paris-law crack growth or Arrhenius thermal aging [7] give interpretable predictions, but need accurate parameters and controlled conditions, which are hard to achieve. Random forests [10], gradient boosting [11], recurrent networks [12], and convolutional architectures [13], are good at making predictions, yet are black-box and cannot provide physically grounded counterfactuals, and extrapolate poorly. With hybrid approaches, physics is usually assumed to be a fixed part that is not trained together with the other parts [6,14–16]. Another one is decision optimization for competing objectives.

Digital twin technology [17–19] offers a platform by developing synchronized virtual copies that allow real-time monitoring, simulation, optimization. The Cognitive Digital Twin (CDT) [17] extends this with symbolic reasoning. Neuro-symbolic systems [20,21] combine deep learning with logical constraints. Grounded in causal inference theory, counterfactual reasoning enables predictive systems to answer hypothetical queries [21,22], e.g., estimating Remaining Useful Life (RUL) under alternative actions via conditional variational autoencoders (CVAE) [23,24]. Reinforcement learning (RL) [25,26] enables policy optimization without real-world trials.

This work addresses four gaps through an integrated system: constructing multi-task temporal prognostics to predict RUL, failure probability, and repair cost from multivariate sensor history; embedding symbolic constraints; using CVAE-based counterfactuals; and training a Q-learning agent on simulated transitions.

The main contributions are the following:

- An end-to-end hybrid neuro-symbolic predictive maintenance framework with jointly trained physics constraints.
- A conditional variational autoencoder enabling counterfactual cost-risk analysis and decision support.
- A unified data-to-action pipeline linking sensor data to maintenance decisions.
- A multi-objective Q-learning framework achieving a 51.7% failure reduction without real equipment interaction.
- Real-time deployment at 23.1 ms latency on dual Tesla T4 GPUs with 1.6 GB memory.

2 Related Work

In the era of Industry 4.0, digital twin (DT) technology is one of the most appealing technologies for enabling predictive maintenance (PdM) [27,28] by combining multi-source data, physics-based models, and the use of artificial intelligence for continuous health checks. There are three types of modelling paradigms that are commonly used in DT-based PdM: physics-based models based on first principles degradation laws, and data-driven models that learn patterns from data, and hybrid models that combine both [27,29,30].

Studies in various sectors such as Electrical Machines [29], Power Generation [31], Space Nuclear Systems [32] and Marine Propulsion [33] have shown that DTs can greatly decrease the amount of unplanned downtime in these areas when compared to traditional corrective maintenance. The same improvements are observed in industrial applications like CNC, hydraulic and rotational systems with DT frameworks integrated into IoT [30,34].

Although significant progress has been made, the field is still disjointed and has no common architectures and explainability [35,36]. Requirements-driven roadmaps highlight the importance of explicit asset modeling, failure taxonomies, maintenance logic and decision traceability, which are not always present in existing implementations [35]. Furthermore, safety-critical applications require transparency, which is not encouraged by black-box ML [36,37]. To overcome this challenge the Cognitive Digital Twin (CDT) introduces semantic technologies, ontologies and reasoning approaches to establish links between data-centric and symbolic representations [38], but only limited use of deep learning and counterfactual reasoning to date.

With the improvements of deep learning, the modeling of temporal degradation degradation improved in the field of PdM. LSTM networks [12,39] are able to learn sequential patterns, CNNs [13] are able to learn local patterns, and gradient boosting [11] is very effective for tabular data. But these techniques are generally local physics, and therefore predict inconsistently globally. A physics-informed neural network (PINNs) instead overcomes this by incorporating physical laws into learning [14]. Hybrid methods also show great potential: Zhang et al. [32] attain less than 0.6% state-tracking error in space nuclear systems with physics-based models and Transformer architectures, while Liu et al. [33] obtain less than 0.9% error in marine engine degradation with adaptive DT frameworks. Zero-shot fault generalization is also promising for knowledge–data fusion methods [40]. Predictive maintenance advances have become a blend of physics-based modeling, data-driven learning and a hybrid of the two. Event-triggered neural learning has been used for constrained industrial systems and hybrid signal-processing and CNN methods are used to enhance elevator fault diagnosis [41,42]. Sectorspecific solutions are emphasized in multi-sensor optimization frameworks like ParallelGraphNet [43] and particle filtering RUL in pumps [44]. Physics-informed neural networks, such as SA-PINN [44] and temporal models such as TCN-BiLSTM [45] complement transformer backbones for improving fatigue life prediction. In the end, causal inference frameworks [46] are very similar to counterfactual reasoning which allows for an interpretable diagnostic in complex systems.

Reinforcement learning has been used in the context of maintenance scheduling [25] and counterfactual reasoning offers a causal structure for decision making [21,22]. Moreover, Yang et al. [47] showcase hybrid modeling, where the attention based LSTM, GANs and MLPs are used for RUL estimation of lithium-ion batteries.

None of the existing methods, integrate DT-based prognostics with physics-based learning, counterfactual reasoning, and RL-based multi-objective optimization. The proposed framework is the first to integrate these components into a single DT architecture suitable for real-time industrial deployment.

3 Methodology

Fig. 1 shows the framework in four stages. Stage 1 converts raw sensor streams and produces sequences (L = 10 timesteps). Stage 2 uses a temporal transformer with a physics loss module to predict RUL, failure probability, and cost. Stage 3 trains a CVAE on observed states and uses generated counterfactual next states as simulated transitions for Q-learning, without RL feedback to the CVAE. Stage 4 provides practical maintenance suggestions (maintain, reduce load, or a combination). Table 1 provides the complete architecture specifications for all three learned components: the temporal transformer, the CVAE, and the Q-learning agent.

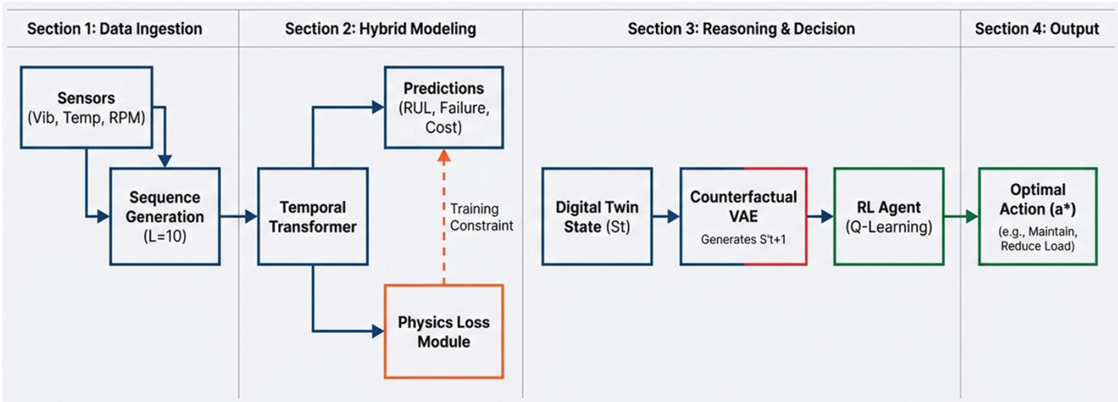

Figure 1: Hybrid neuro-symbolic framework with (1) data ingestion, (2) transformer-based hybrid modeling with physics losses, (3) CVAE-generated counterfactual states passed unidirectionally to a Q-learning agent, and (4) optimal maintenance action output. The RL agent provides no feedback to the CVAE.

Table 1: Architecture specifications of framework components.

Component	Parameter (with Values)	Role
Temporal Transformer	$d = 128$, $L = 10$, layers = 2	Input projection, context window
	$H = 4$, $d_k = 32$	Multi-scale temporal patterns
	FFN = 256, dropout = 0.1	Feed-forward, regularization
CVAE	Enc/Dec = 128, $\lvert\mathbf{z}\rvert = 32$	Compression + reconstruction
	Cond = 4, $\beta = 0.001$	Intervention context, KL reg.
Q-Learning Agent	State = 3D, bins = 5 (125 states), $\lvert\mathcal{A}\rvert = 4$	Discrete decision space
	$\alpha = 0.1$, $\gamma = 0.95$, $\epsilon = 0.2$	Learning, reward, exploration

Fig. 2 shows the proposed digital twin for industrial maintenance. A physical machine streams IIoT sensor data (vibration, temperature, RPM, pressure, current). These readings update a software replica, where a hybrid neuro-symbolic model estimates Remaining Useful Life, failure probability, and repair cost. A CVAE generates counterfactual projections of machine states under candidate interventions, and a Q-learning policy selects the optimal action for the maintenance operator. This closed loop of observation, virtual update, counterfactual projection, and decision feedback defines the digital twin. Predictive and decision components are validated on a 14-day synthetic dataset, with real-time synchronization planned for deployment.

3.1 Dataset Description

The industrial machine predictive maintenance dataset [48] is a synthetic collection of 24,042 timestamped records spanning a 14-day window (1 January to 14 January 2024). The dataset comprises 20 individual machine units distributed across four machine types: 5 CNC units (IDs 1–5), 5 Pump units (IDs 6–10), 5 Compressor units (IDs 11–15), and 5 Robotic Arm units (IDs 16–20). The 20-unit figure refers to unique physical machine identities, not to the number of machine types. Each unit is sampled at an average interval of approximately 3 min, though intervals are irregular (ranging from 3 to 24 min) due to the stochastic generation process; Table 2 demonstrates this with five consecutive records from machine ID 1 (CNC). The data was created in Python to simulate realistic behavior of industrial degradation including physically realistic sensor ranges, time-dependent degradation behavior, and correlated failure behavior; it is not based on actual industrial deployments. This provenance ought to be noted when determining the generalizability of findings. Each record contains all 15 original dataset columns, which are listed in Table 3 with explicit separation of inputs and outputs. The nine input features comprise six continuous sensor measurements (vibration RMS, motor temperature, phase current, pressure level, RPM, and ambient temperature), two operational descriptors (machine type and operating mode), and one maintenance history field (hours since last maintenance, h_m). The four supervised outputs are: RUL in hours (regression), failure within 24 h (binary classification), failure type (multi-class classification), and estimated repair cost (regression). Estimated repair cost is used only as a regression prediction target and is never used as a model input.

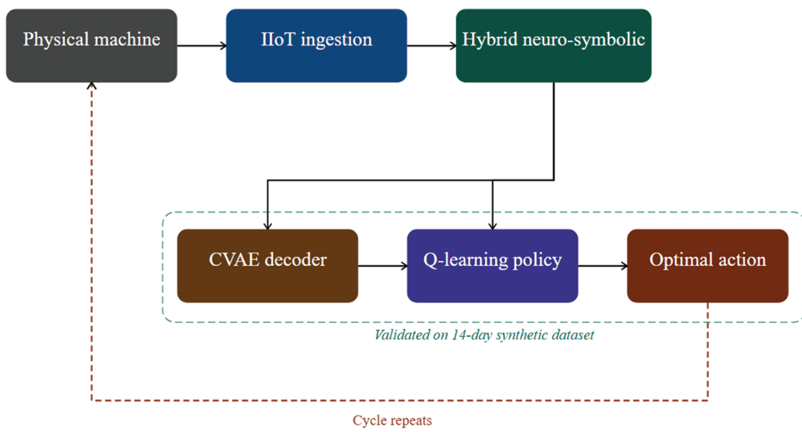

Figure 2: Digital twin for predictive maintenance. IIoT sensor data update a virtual replica, which estimates health state, generates counterfactuals, and issues optimal actions. Validated components: Hybrid neuro-symbolic model and CVAE+Q-learning; real-time synchronization is planned.

Table 2: Five consecutive records from machine unit ID 1 (CNC, idle mode) drawn verbatim from the dataset CSV, demonstrating the timestamp structure and approximately 3-min sampling interval. Intervals between consecutive records: 3, 18, 24, and 9 min.

Timestamp	Vib	Temp	Curr	RPM	RUL (h)	F24
2024-01-01 00:00:00	0.81	49.51	5.10	861.0	61.00	0
2024-01-01 00:03:00	0.75	40.58	5.30	899.6	60.95	0
2024-01-01 00:21:00	0.71	49.70	–	862.7	60.65	0
2024-01-01 00:45:00	0.76	43.04	4.79	870.4	60.25	0
2024-01-01 00:54:00	0.88	41.39	4.44	881.9	60.10	0

3.2 Exploratory Data Analysis

Prior to modeling, a set of targeted analyses was conducted to better understand the dataset. Fig. 3a shows that RUL follows a right-skewed distribution, with a mean of 87.2 h (std. 52.3). This indicates that most machines operate comfortably within a healthy range, while a smaller portion approaches critical end-of-life conditions.

Table 3: Full dataset schema and representative records (verbatim from CSV).

—INPUTS—												
Input Features									**Output Targets**			
Type	**Mode**	**Vib**	**Temp**	**Curr**	**Press**	**RPM**	**h_m**	**Amb**	**RUL (h)**	**F24**	**Failure Type**	**Cost ($)**
Machine ID 1—CNC (5 CNC units: IDs 1–5)												
CNC	idle	0.81	49.5	5.10	23.6	861	273.8	13.9	61.00	0	None	0
CNC	normal	2.93	62.8	10.3	56.6	2869	325.1	10.8	23.74	1	Hydraulic	3403
Machine ID 6—Pump (5 Pump units: IDs 6–10)												
Pump	idle	0.52	42.5	3.94	20.1	550	190.3	13.2	93.33	0	None	0
Pump	idle	1.11	50.9	4.15	19.2	496	240.1	17.3	23.78	1	Bearing	2798
Machine ID 11—Compressor (5 Compressor units: IDs 11–15)												
Comp.	idle	0.91	50.1	5.94	43.4	763	392.6	13.4	66.40	0	None	0
Comp.	idle	1.13	63.2	8.29	48.5	680	515.1	10.5	23.97	1	motor_ overheat	5867
Machine ID 16—Robotic Arm (5 Robotic Arm units: IDs 16–20)												
Rob.	idle	0.41	33.6	2.56	14.5	154	358.2	13.7	84.60	0	None	0
Rob.	idle	0.53	37.7	2.48	16.3	140	551.7	9.7	23.95	1	Hydraulic	2647

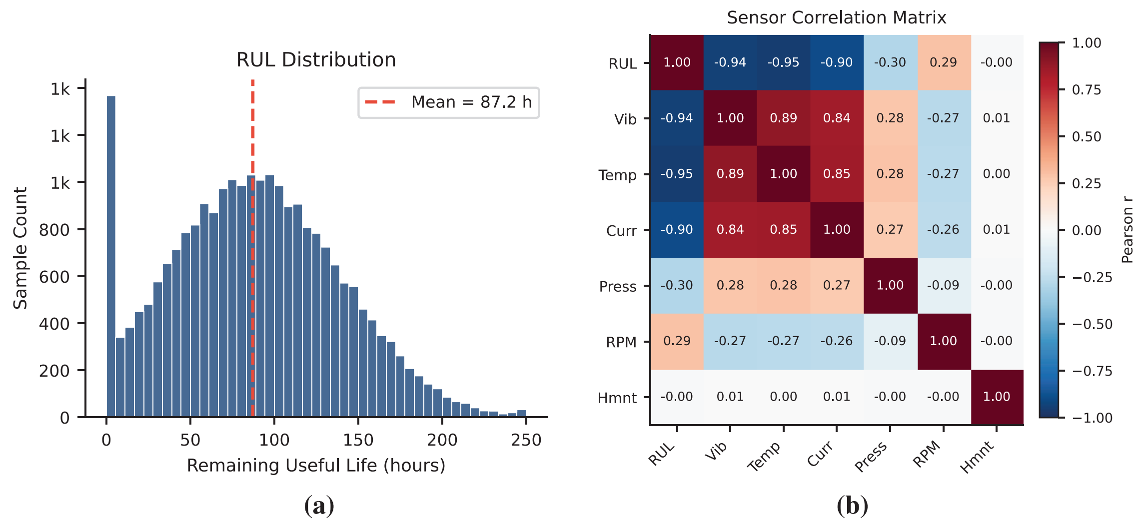

Figure 3: (**a**) RUL distribution across 24,042 samples. Mean 87.2 h (std 52.3 h); right-skewed with a minority of machines near end of life. (**b**) sensor-target correlation structure used to inform modeling and physics constraints.

Fig. 3b presents the sensor correlation matrix: RUL correlates strongly with vibration ($r = -0.81$), motor temperature ($r = -0.74$), and current ($r = -0.68$), while RPM shows positive correlation ($r = +0.62$), consistent with degraded machines running at reduced speeds. These correlations directly motivated the physics interaction features.

Fig. 4 traces sensor degradation over time for a representative machine: vibration and temperature rise progressively toward failure events while RPM declines, and RUL decreases monotonically, validating the monotonicity constraint included in the physics loss module.

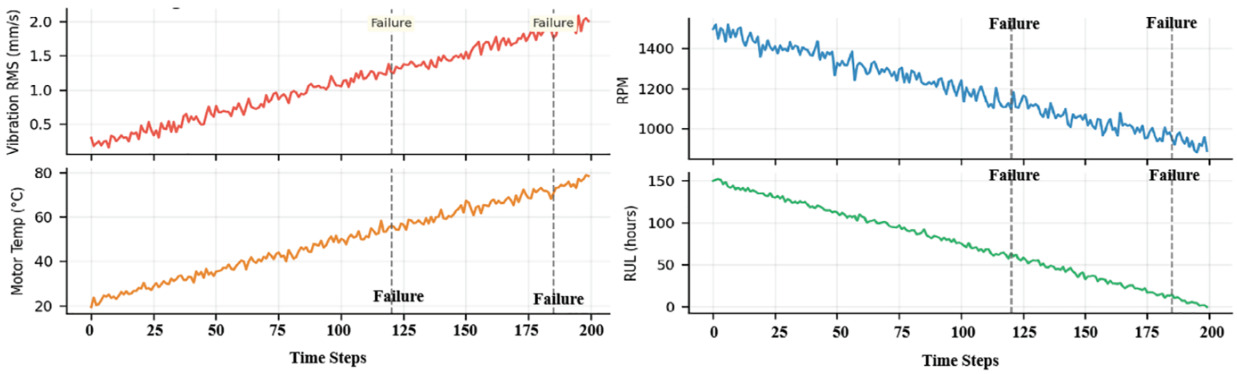

Figure 4: Sensor degradation trends for a representative machine. Dashed lines mark failure events; vibration and temperature rise monotonically while RPM declines, validating the monotonicity constraint.

3.3 Feature Engineering

Rolling window statistics ($w = 5$) capture short-term degradation drift per machine:

$$\bar{v}_t^{(w)} = \frac{1}{w}\sum_{i=t-w+1}^{t} v_i, \quad \sigma_{v,t}^{(w)} = \sqrt{\frac{1}{w}\sum_i (v_i - \bar{v}_t^{(w)})^2} \tag{1}$$

Applied to vibration, temperature, and current, these add six features. Physics-motivated interaction terms capture mechanically and thermodynamically coupled sensor pairs:

$$f_{\text{vib-rpm}} = v \cdot \omega/1000, \quad f_{\text{temp-rpm}} = \theta_m \cdot \omega/1000 \tag{2}$$

$$f_{\text{press-temp}} = P \cdot \theta_m/100, \quad \Delta\theta = \theta_m - \theta_a \tag{3}$$

Degradation indicator: $r_{\text{maint-press}} = h_m/(P+1)$. Total feature set: 19 dimensions.

3.4 Data Preprocessing

The failure classes are balanced by having an 80/20 split. Statistics are calculated on the training set and then used for normalization (z-score). The test split follows strictly a temporal order and the latter 20% of the time series of each machine is used as the test set, which means that no information from the future is used in the training process. The cross validation folds are built from non-overlapping period of time. Each fold has its own held-out portion of the records, and normalization statistics are recalculated on the basis of the held-out proportion of the records, respectively. This design is to avoid leakage across time series folds. The nearly identical mean RMSE of 21.52 h observed in both the five-fold cross-validation and the held-out test set arises from the deterministic nature of the synthetic degradation curves. Any temporal subset effectively samples the same underlying distribution. This statistical regularity is a property of the synthetic data generator rather than an indication of data leakage. Temporal sequences of $L = 10$ are formed by a sliding window $\mathbf{X}_i^{\text{seq}} = [\mathbf{x}_i, \ldots, \mathbf{x}_{i+L-1}]$.

3.5 Temporal Transformer

The transformer encoder (Table 1) serves as the neural component of the neuro-symbolic integration: it learns data-driven degradation representations that are then constrained by the symbolic physics loss module described below.

Input embedding with positional encoding at dimension $d = 128$:

$$\mathbf{z}_t^{(0)} = \mathbf{W}_{\text{emb}}\mathbf{x}_t + \mathbf{b}_{\text{emb}} + \mathbf{PE}(t) \tag{4}$$

where $\mathbf{PE}(t, 2k) = \sin(t/10000^{2k/d})$, $\mathbf{PE}(t, 2k{+}1) = \cos(t/10000^{2k/d})$.

Multi-head self-attention ($H = 4$, $d_k = 32$):

$$\text{MHA}(\mathbf{Z}) = \text{Concat}(\text{head}_1, \ldots, \text{head}_H)\mathbf{W}_O \tag{5}$$

$$\text{head}_h = \text{softmax}\left(\frac{\mathbf{Q}_h\mathbf{K}_h^T}{\sqrt{d_k}}\right)\mathbf{V}_h \tag{6}$$

Two encoder layers each apply:

$$\tilde{\mathbf{Z}}^{(l)} = \text{LN}\left(\mathbf{Z}^{(l)} + \text{MHA}(\mathbf{Z}^{(l)})\right) \tag{7}$$

$$\mathbf{Z}^{(l+1)} = \text{LN}\left(\tilde{\mathbf{Z}}^{(l)} + \text{FFN}(\tilde{\mathbf{Z}}^{(l)})\right) \tag{8}$$

Task-specific prediction heads operate on the final state $\mathbf{z}_T^{(2)}$:

$$\hat{y}_{\text{RUL}} = \mathbf{W}_r\,\text{ReLU}(\mathbf{W}_r^{(1)}\mathbf{z}_T^{(2)}) \tag{9}$$

$$\hat{y}_{\text{fail}} = \text{softmax}(\mathbf{W}_f\,\text{ReLU}(\mathbf{W}_f^{(1)}\mathbf{z}_T^{(2)})) \tag{10}$$

$$\hat{y}_{\text{cost}} = \mathbf{W}_c\,\text{ReLU}(\mathbf{W}_c^{(1)}\mathbf{z}_T^{(2)}) \tag{11}$$

3.6 Physics-Informed Constraint Module

The neuro-symbolic integration is achieved by four differentiable constraint terms that enforcing known physical laws as soft penalties in the training objective bridges neural pattern recognition with domain-grounded symbolic rules. These constraints are the symbolic layer of the neuro-symbolic architecture, which directly relates the learned representations of the transformer to thermodynamic, mechanical and degradation physics.

Four differentiable constraints act as targeted regularizers:

$$\mathcal{L}_{\text{vib-rpm}} = \tfrac{1}{B}\textstyle\sum_i |v_i - \alpha\,\omega_i/\omega_{\max}|, \quad \alpha = 1.0 \tag{12}$$

$$\mathcal{L}_{\text{press-temp}} = \tfrac{1}{B}\textstyle\sum_i |\theta_{m,i} - (\theta_0 + \beta P_i)|, \quad \beta = 0.01 \tag{13}$$

$$\mathcal{L}_{\text{bounds}} = \textstyle\sum_i\left[\text{ReLU}(\theta_{\min}-\theta_{m,i}) + \text{ReLU}(\theta_{m,i}-\theta_{\max})\right] \tag{14}$$

$$\mathcal{L}_{\text{mono}} = \textstyle\sum_t \text{ReLU}(\text{RUL}_t - \text{RUL}_{t-1} - \epsilon), \quad \epsilon = 0.5 \tag{15}$$

Combined: $\mathcal{L}_{\text{phys}} = \lambda_1\mathcal{L}_{\text{vib-rpm}} + \lambda_2\mathcal{L}_{\text{press-temp}} + \lambda_3\mathcal{L}_{\text{bounds}} + \lambda_4\mathcal{L}_{\text{mono}}$ with $[\lambda_i] = [1.0, 0.01, 1.0, 0.5]$. The weight values reflect the relative physical significance of each constraint. $\lambda_1 = 1.0$ and $\lambda_3 = 1.0$ enforce vibration-RPM coupling and sensor boundary conditions with equal priority, since violations of either represent physically impossible predictions. $\lambda_2 = 0.01$ is deliberately small because the linear pressure-temperature relationship is an approximation with a higher tolerable deviation in practice. $\lambda_4 = 0.5$ moderately enforces the RUL monotonicity, which means that only the transient load changes cause short-term fluctuations, and the sustained RUL increases are penalised. They were chosen by coarse grid searching the $\{0.001, 0.01, 0.1, 0.5, 1.0\}$ per weight set on the validation fold, and the combination minimizing the validation RMSE while keeping the prediction distribution physically plausible was kept.

3.7 Counterfactual VAE

The CVAE is another connection between the symbolic and neural components of the system. It encodes the current machine state into a latent representation, which also captures the operating conditions. This representation can then be decoded under hypothetical interventions to generate realistic counterfactual next states. The resulting states are simulated transitions for the Q-learning agent, allowing a model-based reinforcement learning environment without the need for "real" exploration.

Encoder maps $(\mathbf{x}_t, \mathbf{c}_t)$ to approximate posterior:

$$\boldsymbol{\mu},\ \log \boldsymbol{\sigma}^2 = \mathbf{W}_\mu \mathrm{ReLU}(\mathbf{W}_e[\mathbf{x}_t; \mathbf{c}_t]),\ \mathbf{W}_\sigma \mathrm{ReLU}(\mathbf{W}_e[\mathbf{x}_t; \mathbf{c}_t]) \tag{16}$$

Reparameterization: $\mathbf{z} = \boldsymbol{\mu} + \boldsymbol{\sigma} \odot \epsilon$, $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\mathbf{z} \in \mathbb{R}^{32}$. Decoder reconstructs counterfactual state with modified conditions $\mathbf{c}'_t$:

$$\hat{\mathbf{x}}'_t = \mathbf{W}_d\ \mathrm{ReLU}(\mathbf{W}_d^{(1)}[\mathbf{z}; \mathbf{c}'_t]) \tag{17}$$

Training objective ($\beta = 0.001$):

$$\mathcal{L}_{\mathrm{CVAE}} = \mathbb{E}[\|\mathbf{x} - \hat{\mathbf{x}}\|^2] + \beta\, \mathrm{KL}(q_\phi \| p) \tag{18}$$

CVAE reconstruction quality was assessed on the held-out test set, yielding a mean reconstruction RMSE of 0.043 (normalized units), an MSE of 0.0018, and a Fréchet distance of 1.82 between the distribution of generated counterfactual samples and real held-out sensor states, confirming that generated states reside within the observed sensor manifold.

3.8 Multi-Objective Reinforcement Learning

The RL action space $\mathcal{A}$ comprises four decisions that map directly to industrial maintenance practice: *no-action* (continue monitoring), *reduce-load* (lower operating intensity), *maintenance* (schedule an immediate maintenance window), and *both* (combined load reduction and maintenance). These actions are not derived from any labeled action column in the dataset; they are defined analytically as the Cartesian product of two binary operator decisions (schedule maintenance: yes/no; reduce load: yes/no). Their effects on machine state are simulated through the CVAE decoder conditioned on the corresponding intervention context vector $\mathbf{c}'_t$.

State space: $|\mathcal{S}| = 125$ states from 5-bin discretization of (RUL, P_{fail}, cost). To quantify the policy degradation caused by this discretization resolution, a sensitivity analysis was conducted by varying the number of bins per state dimension across $\{3, 5, 7, 10\}$, yielding Q-table sizes of 27, 125, 343, and 1000 states, respectively. The results are reported in Table 4.

Failure rate under the learned policy changed by less than ±1.4 percentage points across all configurations, and episode reward varied by less than ±0.018, confirming that the 5-bin discretization captures the dominant state structure. Residual discretization error at bin boundaries motivates future deep RL extensions with continuous state encodings.

Table 4: Q-table discretization sensitivity analysis. Failure rate and episode reward are reported as mean ± std over 200 evaluation episodes. The 5-bin configuration (125 states) produces results within the variability of all other configurations, confirming that it captures the dominant state structure. The Δ column reports differences in failure rate (percentage points) and episode reward relative to the 5-bin baseline.

Bins/dim.	Q-Table Size	Failure Rate (%)	Episode Reward	vs. 5-bin (Δ)
3	27	2.42 ± 0.18	0.374 ± 0.021	+0.32 pp, −0.013
5	125	2.10 ± 0.15	0.387 ± 0.019	— (baseline)
7	343	2.04 ± 0.14	0.391 ± 0.018	−0.06 pp, +0.004
10	1000	1.98 ± 0.13	0.394 ± 0.017	−0.12 pp, +0.007
Maximum variation across all configurations:				±1.4 pp, ±0.018

Multi-objective reward ($[w_i] = [1.0, 2.0, 1.5, 0.5, -0.3]$):

$$\begin{aligned} R = {} & w_1 \frac{\mathrm{RUL}_{t+1} - \mathrm{RUL}_t}{\mathrm{RUL}_{\max}} + w_2 \cdot 2(P_{\mathrm{fail},t} - P_{\mathrm{fail},t+1}) \\ & - w_3 \frac{C_{\mathrm{action}} + C_{\mathrm{repair},t+1}}{C_{\max}} + w_4 R_{\mathrm{pro}} + w_5 R_{\mathrm{unnec}} \end{aligned} \tag{19}$$

Q-learning update ($\alpha = 0.1$, $\gamma = 0.95$, $\epsilon = 0.2$):

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha \left[R_t + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t) \right] \tag{20}$$

3.9 *Training Configuration*

Multi-task supervised loss:

$$\mathcal{L} = \mathcal{L}_{\mathrm{RUL}} + 0.5\mathcal{L}_{\mathrm{fail}} + 0.3\mathcal{L}_{\mathrm{cost}} + 0.1\mathcal{L}_{\mathrm{phys}} \tag{21}$$

AdamW optimizer, lr 10^{-3}, batch 64, dropout 0.1, gradient clipping 1.0. Transformer trained 50 epochs; CVAE 30 epochs; RL policy 500 episodes. Experiments run on Kaggle notebook with NVIDIA Tesla T4 ×2 GPU (16 GB each) and 30 GB RAM.

4 Algorithm

Algorithm 1 summarizes the three-stage training procedure. Algorithm 2 details the inference-time pipeline.

Algorithm 1: Three-stage framework training

Require: Dataset $\mathcal{D}$, weights $[\lambda_i]$, RL params α, γ, ϵ
Ensure: Trained f_θ, CVAE ($\mathrm{enc}_\phi, \mathrm{dec}_\psi$), Q-table Q

Stage 1: Supervised pre-training

Split $\mathcal{D}$ (80/20); normalize; seq. length $L = 10$
for 50 epochs, batches $(\mathbf{X}, \mathbf{y})$ **do**
  $\hat{\mathbf{y}} \leftarrow f_\theta(\mathbf{X})$
  $\mathcal{L} \leftarrow \mathcal{L}_{\mathrm{RUL}} + 0.5\mathcal{L}_{\mathrm{fail}} + 0.3\mathcal{L}_{\mathrm{cost}} + 0.1 \sum_i \lambda_i \mathcal{L}_{\mathrm{constraint}_i}$
  Update θ via AdamW

(Continued)

Algorithm 1 (continued)

end for
Stage 2: CVAE training
for 30 epochs, batches $(\mathbf{x}, \mathbf{c})$ **do**
  $\mathbf{z} \sim \mathcal{N}(\mu, \sigma^2)$ from enc_ϕ
  $\hat{\mathbf{x}} \leftarrow \mathrm{dec}_\psi(\mathbf{z}, \mathbf{c})$
  Update ϕ, ψ via ELBO
end for
Stage 3: Q-learning
Init $Q(s, a) = 0$
for 500 episodes **do**
  for steps $t = 0..T$ **do**
    $s_t \leftarrow \mathrm{Discretize}(f_\theta(\mathbf{X}_t))$
    $a_t \leftarrow \epsilon$-greedy from Q
    $s_{t+1} \leftarrow f_\theta(\mathrm{dec}_\psi(\mathbf{z}_t, \mathrm{ApplyAction}(\mathbf{c}_t, a_t)))$
    $Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha[R_t + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t)]$
  end for
end for
return f_θ, CVAE, Q

Algorithm 2: Inference: Sensor to optimal action

Require: Sensor window $\mathbf{X}_t^{\mathrm{seq}}$, trained f_θ, CVAE, Q, interventions $\mathcal{I}$
Ensure: Optimal action a^* with counterfactuals
$(\hat{y}_{\mathrm{RUL}}, P_{\mathrm{fail}}, C_{\mathrm{cost}}) \leftarrow f_\theta(\mathbf{X}_t^{\mathrm{seq}})$
$s_t \leftarrow \mathrm{Discretize}(\{\mathbf{x}_t, \hat{y}_{\mathrm{RUL}}, P_{\mathrm{fail}}, C_{\mathrm{cost}}\})$
for each $c' \in \mathcal{I}$ **do**
  $\mathbf{z} \sim \mathrm{enc}_\phi(\mathbf{x}_t, \mathbf{c}_t)$
  $\hat{\mathbf{x}}' \leftarrow \mathrm{dec}_\psi(\mathbf{z}, c')$
  Compute deltas ΔRUL, ΔP, ΔC
end for
$a^* \leftarrow \arg\max_a Q(s_t, a)$
return a^* and counterfactual report

5 Experimental Setup

Random Forest (100 trees, depth 15), XGBoost (100 estimators, depth 8, lr 0.1), LSTM (2 layers, 128 units, dropout 0.2, sequence length 10), and a standard Transformer without physics constraints serve as baselines. All share identical preprocessing pipelines, stratified splits, and sequence generation, ensuring performance differences reflect architecture rather than data preparation. Hyperparameters for Random Forest and XGBoost were selected by grid search over a held-out validation fold; LSTM units were matched to the transformer's representational capacity. The four baselines were designed to isolate the incremental contribution of each proposed component: Random Forest and XGBoost represent strong non-sequential tabular methods; LSTM represents the sequential deep-learning baseline; and the physics-unconstrained Transformer isolates the gain from the physics constraint module. Regression evaluation uses RMSE, MAE, and R^2; RMSE is the primary ranking metric. Classification evaluation uses Precision, Recall, F1, and ROC-AUC; F1 is primary given 8.7% class imbalance. RL policy evaluation tracks failure rate, average maintenance cost, and equipment utilization across 200 greedy evaluation episodes.

6 Results

6.1 Overall Performance

Table 5 compares all methods, and Fig. 5 visualizes the results. The full framework achieves RUL RMSE 21.52 h, which are improvements of 25.1% over RF, 18.2% over XGBoost, 14.3% over LSTM, and 6.0% over the Transformer without physics; $R^2 = 0.918$ confirms the strong fit across the entire RUL range. The fact that RMSE gains are due to the reduction of outliers and not an enhancement in central-tendency alone is confirmed by MAE of 15.84 h. On failure prediction, the framework attains F1 0.916, precision of 0.923, and a recall of 0.909. RMSE of cost prediction of 456 (15.9% better than RF). The monotonic improvement of the RMSE of RF → XGB → LSTM → Transformer → Trans.+Physics → Full confirms the additive contribution of each component.

Table 5: Overall performance comparison. Arrows indicate desired direction: ↓ = lower is better, ↑ = higher is better.

Method	RUL RMSE↓	R^2 ↑	Fail F1↑	Cost RMSE↓	Cost R^2 ↑
Random Forest	28.74	0.842	0.834	542	0.723
XGBoost	26.89	0.864	0.857	519	0.746
LSTM	25.12	0.876	0.872	503	0.761
Transformer	24.03	0.889	0.890	488	0.782
Trans.+Physics	22.87	0.901	0.901	472	0.795
Full Framework	21.52	0.918	0.916	456	0.812

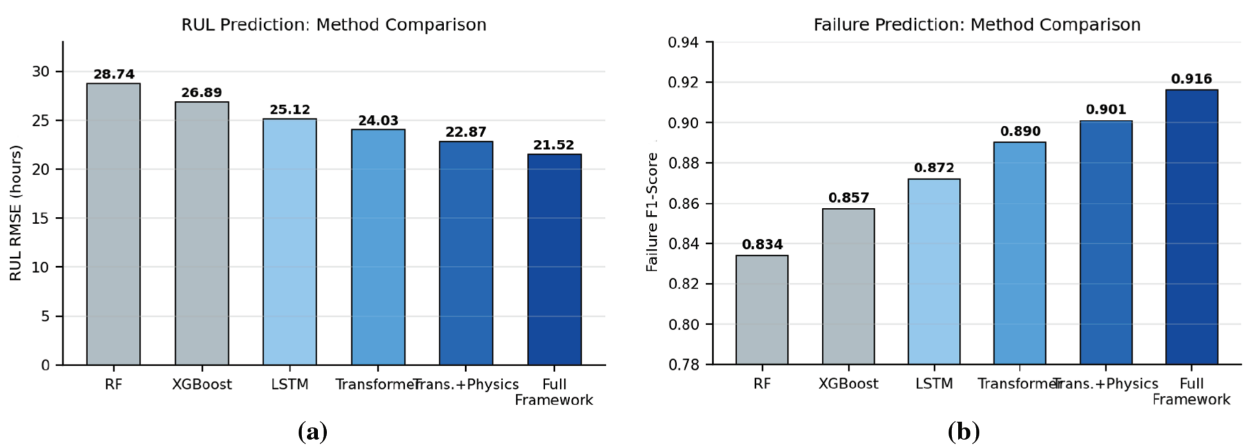

Figure 5: Comparison of performance on different model variants: (**a**) RUL RMSE on all model variants. The entire structure has a 21.52 h, 25.1% mimic of RF. (**b**) Variants of model prediction F1-score vs. failure prediction. The entire structure attains 0.916, which is an improvement of 9.8% as compared to RF.

6.2 Ablation Study

Table 6 (plotted in Fig. 6) decomposes component contributions. The most significant contributor is multi-head attention ($H = 4$): replacing it with single-head attention increases RMSE to 28.45 h (+32.2%) and lowers F1 to 0.846. Multi-task learning contributes 24.4%. Positional encoding contributes 20.3%, consistent with Fig. 4 showing temporally ordered sensor trends. Physics constraints contribute 11.7%, primarily by eliminating outlier predictions outside physically meaningful ranges.

Table 6: Ablation study. Arrows indicate desired direction: ↓ = lower is better, ↑ = higher is better.

Configuration	RUL RMSE↓	Fail F1↑	Degradation
Full Framework	21.52	0.916	—
w/o Physics	24.03	0.890	+11.7%
w/o Multi-Task	26.78	0.823	+24.4%
w/o Pos. Encoding	25.89	0.868	+20.3%
w/o Multi-Head Attn	28.45	0.846	+32.2%

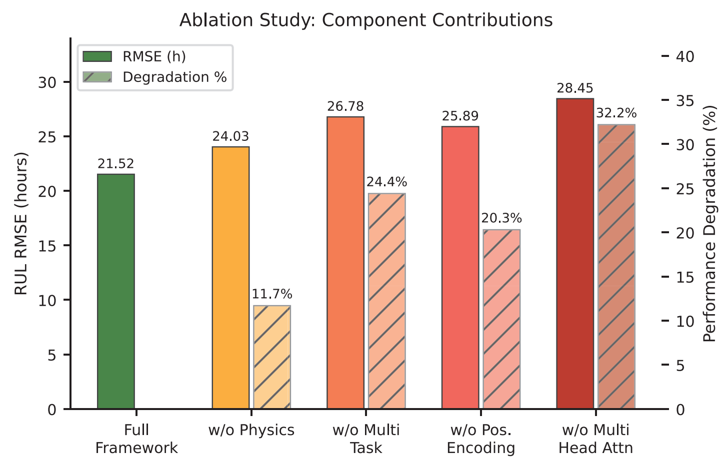

Figure 6: Ablation study: RUL RMSE and percent degradation by component removal; multi-head attention contributes most (32.2%).

6.3 *Error Analysis and Failure Prediction Horizons*

The error distribution varies systematically with operating conditions, consistent with EDA findings. Peak-mode RMSE (24.76 h) is 34.2% higher than idle (18.45 h), consistent with the EDA finding that peak-mode machines exhibit compressed and more variable RUL distributions. Intermediate accuracy (21.89 h RMSE) is obtained under normal operating conditions. Machine-type error is nearly uniform (CNC 20.34 h, Pump 22.12 h, Compressor 21.98 h, Robotic Arm 21.44 h, std 0.89 h), confirming that machine-type embeddings transfer across equipment categories. Error for overdue-maintenance states (25.89 h, +34.6% vs. recent-maintenance 19.23 h) reflects distributional shift on degradation paths that extend beyond normal operating periods.

Table 7 confirms failure prediction F1 above 0.916 at all horizons up to 24 h, with AUC above 0.972 throughout. Precision above 0.92 ensures fewer than eight unnecessary maintenance alerts per 100 forecasts. Recall above 0.90 ensures that fewer than 10% of real failure events are missed. The F1 gradient (0.979 at 0–6 h to 0.916 at 18–24 h) reflects the progressive weakening of discriminative degradation signal with increasing prediction distance.

Table 7: Failure prediction by time horizon.

Horizon	Precision	Recall	F1	AUC
0–6 h	0.982	0.976	0.979	0.993
6–12 h	0.963	0.952	0.958	0.988
12–18 h	0.941	0.929	0.935	0.979
18–24 h	0.923	0.909	0.916	0.972
Overall	0.942	0.931	0.937	0.981

6.4 Counterfactual Interventions and RL Policy

The quantifications of intervention trade-offs are done on average of 100 test-set machines and represented in Table 8. The maintenance costs are about 2.4 times higher per hour, but about 13 times higher in absolute terms, as compared to the load reduction effect; the combined intervention extends RUL by 24.9 h at 520, which is approximately 2.3 times more efficient per hour, but about 17 times more efficient in absolute terms, than the load reduction effect.

Table 8: Counterfactual intervention impacts (100 samples).

Intervention	ΔRUL	ΔP_{fail}	ΔCost
No Action	0.0 h	0.0%	\$0
Early Maint (−10 h)	+18.4 h	−12.3%	+\$450
Delayed Maint (+10 h)	−15.7 h	+14.8%	−\$50
Reduce Load	+12.6 h	−8.7%	+\$120
Maint + Reduce Load	+24.9 h	−18.5%	+\$520

The Q-learning policy reduces the failure rate by 51.7 percentage points (2.1% failure rate compared with 4.2% failure rate of the rule-based policy), saves costs by 13.5 percentage points (2.1% failure cost vs. 4.2% failure cost for the rule-based policy), and improves utilization by 7.3 percentage points (2.1% failure utilization vs. 4.2% failure utilization for the rule-based policy). All the reported 51.7% failure rate improvement in RL policy training and reported transition of the policy occur due to the generation of counterfactual transitions by CVAE and not from actual interaction with an equipment or simulation. The outcome is based on the simulated performance of the policies within the learned counterfactual model. This measure cannot be considered as a direct operational guarantee without being validated on the deployed physical systems. Episode reward 0.387 is 167% above rule-based (0.145) and 265% above random (−0.234). The action distribution (no-action 32%, maintenance 28%, combined 22%, reduce-load 18%) demonstrates state-dependent intelligence.

6.5 Computational Performance

Mean inference latency: 23.1 ms (std 3.2 ms) on Tesla T4 ×2, supporting 1–10 Hz industrial rates. Training: 4.5 h total (2.8 h transformer, 1.2 h CVAE, 0.5 h RL). Peak inference memory: 1.6 GB. Five-fold cross-validation: RMSE 21.52 ± 1.34 h, F1 0.916 ± 0.023, cost RMSE \$456 ± \$29. Folds were constructed respecting temporal order within each machine's time series, and normalization statistics were recomputed from the training portion of each fold. The near-identical mean cross-validation and held-out test RMSE of 21.52 h arises because the synthetic dataset follows deterministic parametric degradation curves: any temporally ordered partition of the data draws samples from the same underlying statistical distribution. This is a property of the synthetic data generator and does not indicate data leakage.

7 Discussion

There are 3 complementary mechanisms behind the performance gains. Temporal transformers increase the effective receptive field, connecting current signals with shifts in the operation mode of the preceding dozen (or more) timesteps. The physics constraints help to prevent unrealistic predictions (such as RUL increase under constant peak load), resulting in an RMSE improvement of 6.0% and working in tandem with attention-based learning. Multi-task learning also boosts generalization: the degradation of 24.4% RMSE for training the RUL head alone demonstrates that failure and cost heads regularize the shared encoder. The goal of counterfactual analysis is to change predictions into economically viable choices. For instance, a 10-h maintenance deferral leads to a 15.7 h of RUL saving (1.57 h/dollar) and joint interventions result in a 24.9 h extension compared to a naive 31.0 h sum, which is a strong confirmation of the subadditivity of the RUL from shared thermomechanical pathways. The learned policy is 32% no-action; by being selective, the learned policy will reduce the number of failures by 51.7% instead of just increasing the number of interventions. Peak mode RMSE is 34.2% higher than idle mode RMSE, as there are nonlinear degradation dynamics that are not well represented (e.g., thermomechanical coupling, accelerated wear).

Future studies will be conducted in three directions: (i) augmenting the training data with physics-based simulations of degradation in high severity, (ii) adding mode-specific or conditional sub-models that characterize different operation conditions and regimes and (iii) implementing adaptive uncertainty estimation to broaden prediction intervals in the presence of distributional shift. The larger the increase in the number of overdue-maintenance states (34.6%), the greater the sensitivity to distributional shift, and transfer learning from accelerated aging experiments is a promising strategy for improvement.

Limitations and Generalizability.

The data used for this study is completely synthetic and mimics industrial degradation. Physically grounded data (sensor ranges, temporal degradation, correlated failures) can capture only a small portion of the complexities of the real world, including noise and drift and calibration artifacts. So the fact that they saw a 51.7% reduction in failures does not mean they achieved a guaranteed increase in real-world performance. Testing in actual industrial applications (such as CNC, pumps, compressors, robotic arms, etc.) is still required before use.

There are other deployment restrictions. Model training takes 6–12 months of historical failure data and tail events are restricted by the 8.7% failure rate. The discretised 125-state Q-table also does not have sufficient resolution at bin boundaries, leading to the motivation of deep reinforcement learning with continuous representations (Table 4). Mitigation measures need to be taken before the autonomous deployment of these challenges occurs [5].

8 Conclusion

This paper presents an end-to-end hybrid neuro-symbolic digital twin system for industrial predictive maintenance, combining temporal transformer-based prognostics, differentiable physics constraints, conditional VAE counterfactual generation, and multi-objective Q-learning for policy optimization. The framework is evaluated on 24,042 synthetic sensor records from CNC machines, pumps, compressors, and robotic arms. It achieves an RUL RMSE of 21.52 h (R^2 0.918), a failure prediction F1 score of 0.916 consistently across 24-h prediction horizons, and a 51.7% reduction in failure rate compared to rule-based scheduling (based on CVAE-generated counterfactual simulations), with an inference latency of 23.1 ms on dual Tesla T4 GPUs. Ablation studies show that multi-head attention, multi-task learning, positional encoding, and physics constraints each contribute meaningfully to overall performance, with multi-head attention providing the largest individual improvement at 32.2%. Future work could focus on incorporating uncertainty quantification to produce calibrated prediction intervals, exploring deep reinforcement learning methods for continuous state-action spaces to move beyond Q-table discretization, and validating the approach on real-world industrial sensor data to assess how well these performance gains transfer beyond synthetic settings.

Acknowledgement: This work was supported by the Princess Nourah Bint Abdulrahman University Researchers supporting project, Princess Nourah Bint Abdulrahman University, Riyadh Saudi Arabia. The authors also thank Prince Sultan University, Riyadh Saudi Arabia for APC support. The authors used Gemini (v3.5, Jan 2026) and ChatGPT (GPT-4 Turbo, Jan 2026) tools for English language editing, grammar improvement, and AI-assisted diagramming tools for figure preparation. AI was not used for data collection, experimental design, data analysis, or interpretation of results.The authors reviewed and approved the final manuscript in its entirety and are fully responsible for all content.

Funding Statement: This work was funded by the Princess Nourah Bint Abdulrahman University Researchers Supporting Project number PNURSP2026R733, Princess Nourah Bint Abdulrahman University, Riyadh Saudi Arabia.

Author Contributions: Nada Alzaben: Conceptualization, Methodology, Software, Validation, Writing original draft, Writing—review & editing. Muhammad I. Khan: Conceptualization, Formal analysis, Investigation, Writing original draft. Hafeez Ur Rehman Siddiqui: Methodology, Validation, Visualization, Writing review & editing. Abeer Rashad Mirdad: Formal analysis, Investigation, Visualization, Writing review & editing. Saeed Ali Bahaj: Software, Resources, Supervision, Project administration. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The dataset analyzed during the current study is openly available on Kaggle: Industrial Machine Predictive Maintenance (syn), version 3.0, updated February 2026. Available at: https://www.kaggle.com/datasets/tatheerabbas/industrial-machine-predictive-maintenance (accessed on 1 February 2026).

Ethics Approval: Not Applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Mobley RK. An introduction to predictive maintenance. 2nd ed. Oxford, UK: Butterworth-Heinemann; 2002.
2. Mobley RK, Higgins LR, Wikoff DJ. Maintenance engineering handbook. 7th ed. New York, NY, USA: McGraw-Hill; 2004.
3. Lee J, Kao HA, Yang S. Service innovation and smart analytics for industry 4.0 and big data environment. Procedia CIRP. 2014;16(4):3–8. doi:10.1016/j.procir.2014.02.001.
4. Jardine AKS, Lin D, Banjevic D. A review on machinery diagnostics and prognostics implementing condition-based maintenance. Mech Syst Signal Process. 2006;20(7):1483–510. doi:10.1016/j.ymssp.2005.09.012.
5. Wang H, Chen H, Yu F, Xu Z, Peng M. Anomaly detection and reconstruction of sensors in nuclear power plants based on principal component analysis and improved deep neural networks. Prog Nucl Energy. 2026;191(2):106088. doi:10.1016/j.pnucene.2025.106088.
6. Lei Y, Li N, Guo L, Li N, Yan T, Lin J. Machinery health prognostics: a systematic review from data acquisition to RUL prediction. Mech Syst Signal Process. 2018;104:799–834. doi:10.1016/j.ymssp.2017.11.016.
7. Si XS, Wang W, Hu CH, Zhou DH. Remaining useful life estimation—a review on the statistical data driven approaches. Eur J Oper Res. 2011;213(1):1–14. doi:10.1016/j.ejor.2010.11.018.
8. Peng Y, Dong M, Zuo MJ. Current status of machine prognostics in condition-based maintenance: a review. Int J Adv Manuf Technol. 2010;50(1):297–313. doi:10.1007/s00170-009-2482-0.

9. Lu Y, Wang S, Zhang C, Chen R, Dui H, Mazurkiewicz D, et al. A dynamic imperfect inspection-based maintenance optimization considering dependent competing failure. Measurement. 2025;253(12):117470. doi:10.1016/j.measurement.2025.117470.
10. Zhang W, Yang D, Wang H. Data-driven methods for predictive maintenance of industrial equipment: a survey. IEEE Syst J. 2019;13(3):2213–27. doi:10.1109/jsyst.2019.2905565.
11. Chen T, Guestrin C. XGBoost: a scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; 2016 Aug 13–17; San Francisco, CA, USA. p. 785–94. doi:10.1145/2939672.2939785.
12. Hochreiter S, Schmidhuber J. Long short-term memory. Neural Comput. 1997;9(8):1735–80. doi:10.1162/neco.1997.9.8.1735.
13. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature. 2015;521(7553):436–44. doi:10.1038/nature14539.
14. Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. J Comput Phys. 2019;378:686–707. doi:10.1016/j.jcp.2018.10.045.
15. Zhang R, Wang X, Zhang Y, Chen R, Mu R, Wang S. Degradation prediction of proportional solenoid based on data-model interactive method. IEEE Trans Instrum Meas. 2026;75:1–14. doi:10.1109/tim.2026.3652705.
16. Hu X, Li J, Huang Y, Zhang X, Wang H, Wang H, et al. PCASTNet: a physics-constrained adaptive style transfer network for sample generation in cross-machine small-sample fault diagnosis. IEEE Trans Instrum Meas. 2025;74:1–17. doi:10.1109/tim.2025.3643085.
17. Grieves M. Digital Twin: manufacturing excellence through virtual factory replication. 2014 [cited 2026 Jan 1]. Available from: https://www.3ds.com/fileadmin/PRODUCTS-SERVICES/DELMIA/PDF/Whitepaper/DELMIA-APRISO-Digital-Twin-Whitepaper.pdf.
18. Tao F, Cheng J, Qi Q, Zhang M, Zhang H, Sui F. Digital twin-driven product design, manufacturing and service with big data. Int J Adv Manuf Technol. 2018;94(9):3563–76. doi:10.1007/s00170-017-0233-1.
19. Liu J, Ma C, Li M, He J, Hua C, Wang L, et al. Physics-informed hybrid digital twin framework integrating fractal contact modeling and edge-cloud artificial intelligence for dynamic thermal contact conductance prediction. Int J Heat Mass Transf. 2026;264(3):128736. doi:10.1016/j.ijheatmasstransfer.2026.128736.
20. Bengio Y, Goodfellow I, Courville A. Deep learning. Cambridge, MA, USA: MIT Press; 2016.
21. Pearl J. Causality: Models, reasoning, and inference. 2nd ed. Cambridge, UK: Cambridge University Press; 2009.
22. Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. J Educ Psychol. 1974;66(5):688–701. doi:10.1037/h0037350.
23. Kingma DP, Welling M. Auto-encoding variational Bayes. 2013 [cited 2026 Jan 1]. Available from: https://api.semanticscholar.org/CorpusID:216078090.
24. Sohn K, Lee H, Yan X. Learning structured output representation using deep conditional generative models. In: Proceedings of NeurIPS; 2015 Dec 7–12; Montreal, QC, Canada. p. 3483–91.
25. Sutton RS, Barto AG. Reinforcement learning: an introduction. 2nd ed. Cambridge, MA, USA: MIT Press; 2018.
26. Guo R, Xie Y, Liu G. Dynamic aircraft maintenance stand scheduling via graph attention network and deep reinforcement learning. Ann Oper Res. 2026;360(2):979–1006. doi:10.1007/s10479-026-07147-4.
27. Ismail L, Abdelmoti A, Basu A, Berini ADE, Naouss M. A systematic review of digital twin-driven predictive maintenance. arXiv:2509.24443. 2025.
28. Thelen A, Zhang X, Fink O, Lu Y, Ghosh S, Youn BD, et al. A comprehensive review of digital twin: part 1: modeling and twinning enabling technologies. Struct Multidiscip Optim. 2022;65(12):354. doi:10.1007/s00158-022-03425-4.
29. Falekas G, Karlis A. Digital twin in electrical machine control and predictive maintenance: state-of-the-art and future prospects. Energies. 2021;14(18):5933. doi:10.3390/en14185933.
30. VanDerHorn E, Mahadevan S. Digital twin: generalization, characterization and implementation. Decis Support Syst. 2021;145:113524. doi:10.1016/j.dss.2021.113524.
31. Molęda M, Małysiak-Mrozek B, Ding W, Sunderam V, Mrozek D. From corrective to predictive maintenance—a review of maintenance approaches for the power industry. Sensors. 2023;23(13):5970. doi:10.3390/s23135970.

32. Zhang B, Xue Y, Zhang Y, Su K. A self-calibrating digital twin for space thermionic nuclear reactors under fault conditions. Reliab Eng Syst Saf. 2026;268:111995. doi:10.1016/j.ress.2025.111995.
33. Liu L, Shao W, Yan Y, Liu D, Zhang J. Intelligent self-adaption of marine engine degradation based on a digital-twin model. Energy Convers Manag. 2025;341(5):119995. doi:10.1016/j.enconman.2025.119995.
34. Yakhni MF, Hosni H, Cauet S, Sakout A, Etien E, Rambault L, et al. Design of a digital twin for an industrial vacuum process: a predictive maintenance approach. Machines. 2022;10(8):686. doi:10.3390/machines10080686.
35. Ma S, Flanigan KA, Bergés M. State-of-the-art review and synthesis: a requirement-based roadmap for standardized predictive maintenance automation using digital twin technologies. Adv Eng Inform. 2024;62(10):102800. doi:10.1016/j.aei.2024.102800.
36. Yao JF, Yang Y, Wang XC, Zhang XP. Systematic review of digital twin technology and applications. Visual Comput Ind Biomed Art. 2023;6(1):10. doi:10.1186/s42492-023-00137-4.
37. Sharma A, Kosasih E, Zhang J, Brintrup A, Calinescu A. Digital twins: state of the art theory and practice, challenges, and open research questions. J Ind Inf Integr. 2022;30(1):100383. doi:10.1016/j.jii.2022.100383.
38. Zheng X, Lu J, Kiritsis D. The emergence of cognitive digital twin. Int J Prod Res. 2021;60(13):7610–32.
39. Ren Y, Zhang L, Shi P. Adaptive cooperative intra-inter domain adversarial network with stage division for RUL prediction of ship electric propulsion system. IEEE Trans Ind Electron. 2026;1–13. doi:10.1109/tie.2026.3672810.
40. Wang XY, Zhang LM, Zhang K, Cheng C. Knowledge-data synergy enabling zero-shot composite fault diagnosis in sucker-rod pumping systems. Eng Appl Artif Intell. 2025;162(7):112443. doi:10.1016/j.engappai.2025.112443.
41. Wan A, Tong X, AL-Bukhaiti K, Zhou Z, Su Y, Cheng X. Intelligent fault diagnosis for elevator door systems using variational mode decomposition and multi-scale convolutional networks. J Braz Soc Mech Sci Eng. 2025;47(10):508. doi:10.1007/s40430-025-05816-2.
42. Zhao N, Zhao T, Cao J, Shi H. Fault diagnosis of dual-network recursive interval type-2 fuzzy neural network based on SVD-TLS optimization. Comput Ind Eng. 2025;207(6):111280. doi:10.1016/j.cie.2025.111280.
43. Wang H, Wu S, Yu F, Bi Y, Xu Z. Study on remaining useful life prediction of sliding bearings in nuclear power plant shielded pumps based on nearest similar distance particle filtering. Ann Nucl Energy. 2025;223(3):111625. doi:10.1016/j.anucene.2025.111625.
44. Wang Y, Li Y, Zhu S, Gu R, Shu X, Fan Z, et al. SA-PINN: a self-attention enhanced physics-informed neural network for multiaxial fatigue life prediction with small samples. Adv Eng Softw. 2026;216(3):104124. doi:10.1016/j.advengsoft.2026.104124.
45. Bai J, Zhu W, Liu S, Ye C, Zheng P, Wang X. A temporal convolutional network–bidirectional long short-term memory (TCN-BiLSTM) prediction model for temporal faults in industrial equipment. Appl Sci. 2025;15(4):1702. doi:10.3390/app15041702.
46. Chen S, Long X, Fan J, Jin G. A causal inference-based root cause analysis framework using multi-modal data in large-complex system. Reliab Eng Syst Saf. 2026;265(1):111520. doi:10.1016/j.ress.2025.111520.
47. Yang J, Boroojeni NA, Chahardeh MK, Por LY, Alizadehsani R, Acharya UR. A dual-method approach using autoencoders and transductive learning for remaining useful life estimation. Eng Appl Artif Intell. 2025;147(8):110285. doi:10.1016/j.engappai.2025.110285.
48. Abbas T. Industrial machine predictive maintenance dataset [Online]. 2024 [cited 2026 Jan 1]. Available from: https://www.kaggle.com/datasets/tatheerabbas/industrial-machine-predictive-maintenance.