

A Physics-Informed IIoT Digital Twin Framework for Automotive Cooling Predictive Maintenance and Odometer Fraud Detection using LSTM-Attention

Ali Haider Mohammed Hussein^{1*}, Muatamed Abed Hajer²

¹Department of Computer Science, College of Computer Science and Information Technology, University of Sumer, Thi-Qar, Iraq

²Department of Computer Science, College of Computer Science and Information Technology, University of Sumer, Thi-Qar, Iraq

* Email address of the Corresponding Author: alihaidar@stu.edu.iq, ²Email: m.hajer@uos.edu.iq
Phone number: 07830975670

Abstract

The shift to Industry 4.0 in the automotive industry requires an approach which moves away from passive monitoring and towards physics-based predictive maintenance. This paper introduces a Physics-Informed IIoT-Enabled Digital Twin (PI-DT) framework based on an Adaptive Physics-Informed Long Short-Term Memory with Attention (PI-LSTM-Attention) model to forecast the Remaining Useful Life (RUL) of Engine Cooling Systems (ECS). The proposed architecture integrates the thermodynamic Monotonicity constraints into the learning process using Physics-Informed, meeting the 100% target of Monotonicity compliance, 99.64% Failure-Point Prediction Accuracy, and a self-healing rate of 0.00%. The framework is a Dual-Layer Edge-Cloud Continuum for ultra-low latency and complete offline operational resilience. In a Zero Trust cybersecurity paradigm, the accumulated ledgers of thermal damage are continually compared with the digital reading on the odometers to identify tamper and warranty fraud. The proposed PI-LSTM-Attention network is compared with Transformer architectures with lower latency and higher prediction accuracy, which proves its unique optimization for the automotive Edge, especially for time-critical applications.

Keywords: Industrial Internet of Things (IIoT); Digital Twin; Physics-Informed Machine Learning; LSTM-Attention; Remaining Useful Life; Engine Cooling System; Odometer Fraud Detection.

إطار عمل التوأم الرقمي قائم على الفيزياء لإنترنت الأشياء الصناعية للصيانة
التنبؤية لتبريد السيارات والكشف عن التلاعب بعدد المسافة باستخدام
خوارزمية الذاكرة طويلة وقصيرة المدى مع آلية الانتباه
علي حيدر محمد حسين^{1*}، معتمد عبد هجر²

¹قسم علوم الحاسوب، كلية علوم الحاسوب و تكنولوجيا المعلومات، جامعة سومر، ذي قار، العراق

الخلاصة

يتطلب الانتقال إلى الثورة الصناعية الرابعة في صناعة السيارات تحولاً من المراقبة السلبية إلى الصيانة التنبؤية القائمة على الفيزياء. تقدم هذه الورقة إطار عمل توأم رقمي قائم على الفيزياء وممكن بتقنية إنترنت الأشياء، يستخدم نموذج ذاكرة طويلة المدى تكيفي مدعوم بالفيزياء مع آلية انتباه للتنبؤ بالعمر المتبقي لأنظمة تبريد المحرك. يدمج التصميم المقترح قيود الرتبة الديناميكية الحرارية في عملية التعلم القائمة على الفيزياء، محققاً رتبة شبيهة مثالية، ودقة تنبؤ بنقاط الفشل بنسبة 99.64%، ومعدل إصلاح ذاتي 0.00%. يتميز إطار العمل ببنية ثنائية الطبقات، من الحافة إلى السحابة، مما يوفر زمن استجابة منخفضاً للغاية ومرونة تشغيلية كاملة، حتى في وضع عدم الاتصال. في نموذج أمن سببراني قائم على مبدأ انعدام الثقة، تتم مقارنة سجلات الأضرار الحرارية المتراكمة باستمرار مع قراءات عداد المسافة الرقمي لتحديد التلاعب بالضمان والاحتيال. تمت مقارنة شبكة ذاكرة طويلة المدى تكيفية مدعومة بالفيزياء مع آلية الانتباه المقترحة بهيكل بنية التحويل المعتمدة على الانتباه، مما أظهر أوقات استجابة أقل ودقة تنبؤ أعلى، مما يدل على كفاءتها الفريدة في تطبيقات الحافة في قطاع السيارات، وخاصة في التطبيقات الحساسة للوقت.

1. Introduction

1.1 Industrial Context & Motivation

Today's automotive industry is entering one of its most transformative periods, with the clear transition from a purely mechanical system to a much more intelligent, connected and autonomous one. The driving force behind this metamorphosis is basically the Fourth Industrial Revolution (Industry 4.0), characterized by the pervasive use of Cyber-Physical Systems (CPS), decentralized decision making and the spread of the Industrial Internet of Things (IIoT). In this context, vehicles have become complex mobile sensor hubs which are part of the larger Internet of Vehicles (IoV) network [1, 2].

The IIoT is the industrial backbone of the transformation, which demands the highest standards in terms of ultra-low latency, deterministic communications, and mission critical reliability. The IIoT allows for seamless bidirectional exchange of thermodynamic and operational data between the physical and digital vehicle, which could otherwise only be obtained through traditional diagnostics, in real time [3]. Modern smart manufacturing offers the technical facilities needed to support the systematic science of estimating a system's current health state and predicting its Remaining Useful Life (RUL) before failure, known as Prognostics and Health Management (PHM) [4].

1.2 Criticality of Engine Cooling Systems (ECS)

The Engine Cooling System (ECS) is one of the most important thermal management subsystems in an Internal Combustion Engine (ICE) that has approximately 2/3 of fuel combustion energy dissipated as heat [5]. If the temperature is not maintained at thermodynamic equilibrium, then thermal runaway can occur, causing warped cylinder heads or seizing the pistons and resulting in catastrophic damage [6]. For ECS, failures directly affect passenger safety, and the addition of new complex stochastic ECS failure modes by modern electronic thermostats and variable-speed water pumps further makes ECS a safety-critical system [7] .

1.3 Problem Statement

There is a basic lack of understanding of how to marry artificial intelligence algorithms, the physical thermodynamic laws, computational limitations, and cyber security frameworks, which currently limits the state of automotive predictive maintenance. There are four major gaps:

Physical Inconsistency Gap; purely data-driven Black-Box models consider ECS as a statistical system and make physically impossible predictions of self heating that violate the 2nd Law of Thermodynamics [8,9]. Heavy architectures like Large Language Models (LLMs) and massive Transformers demand significant computational power that's beyond the reach of resource-constrained vehicles' Edge nodes [10,11]. Automotive datasets are virtually nonexistent; a catastrophic failure of a physical vehicle cooling system is too expensive and too dangerous to run it to failure, hence the Data Scarcity Gap [12,13]. Cyber-Physical Security Gap; Fraud that takes place around the world is related to odometer tampering. Existing security mechanisms only consider this as a digital cryptography challenge and

neglect the physical forensic capabilities from the collected data of thermal degradation [14].

1.4 Research Contributions

Five main scientific contributions are made:

PI-LSTM-Attention Framework: A novel hybrid architecture embedding thermodynamic Monotonicity constraints within the feature engineering pipeline that obtains physical consistency of RUL predictions 100% of the time while avoiding self-healing predictions.

Heterogeneous Transfer Learning Bridge: A domain adaptation methodology that is able to successfully transfer thermodynamic failure knowledge from the NASA C-MAPSS aerospace dataset [15] to the automotive ECS domain.

Dual-Layer Edge-Cloud Continuum Architecture: A deployment strategy that offers comprehensive operational resilience for real-time autonomous safety interventions, without requiring an online connection.

Zero-Trust Physics Based Odometer Verification: The first framework that combines irreversible thermal degradation history as an unalterable physical anchor to detect digital odometer tampering.

Comprehensive Comparative Validation: Systematic quantitative validation over baseline models such as Transformers and LLMs on all key metrics.

2. Literature Review

The industrial revolution 4.0 has introduced a new concept of predictive maintenance. Under the industry 4.0, new concept of predictive maintenance has been introduced.

The problem with conventional condition based maintenance is that it is based on static thresholds and not dynamic degradation monitoring, which leads to systematic false positive rates [16]. Industry 4.0 PdM pipelines with minimal information in the model, as well as models built from aggregated operational data, are less accurate than pipelines built directly from the raw data; and industrial deployment needs models that can be explained and have physical interpretation [17,18].

The concept of digital twins is becoming a reality in the automotive sector. Digital Twins are real-time virtual replicas of physical assets, which are synchronously linked with bidirectional data exchange [19, 20]. In the field of automotive ECS monitoring, existing deep digital twin methods based on Stacked Sparse Autoencoders are purely data-centric, with no restrictions on physical deployments and only applicable for Cloud deployment, without offline Edge operation or security measures [6].

Inconsistencies between the physical and the virtual world. In data-driven models, with no physical constraints, the conservation laws are violated in an inherent way, e.g., the "Self-Healing Phenomenon", which means that spontaneous recovery of RUL is predicted,

contrary to the Second Law of Thermodynamics. Purely data-driven models have systematic failures when data is sparse and distribution of data has shifted, which is an imperative requirement for safety-critical applications in industry [9] that relies on physically interpretable models.

Transformers and Large Language Models (LLMs) have shown great potential for time-series forecasting and fault diagnosis tasks [10, 11]. But these architectures are quadratic in sequence length and need large numbers of parameters, which are not feasible to be realized in Edge Electronic Control Unit (ECU) hardware. Leclerc et al. have formally quantified latency thresholds for safety-critical automotive IoT applications and showed that safety risks will occur directly if the latency is larger than certain millisecond limits [22]. The only viable way to deploy such models into the Edge node according to Chaurasia et al. is through the use of lightweight models that can be computed on the edge node [23]. Physical laws are powerful inductive constraints that make the interpretability, generalization and reliability of prognostic models significantly more powerful [13]. In recent years, great efforts have been made to develop gradient-enhanced Physics-Informed LSTM networks for RUL prediction, which proved that the physical constraints result in the stable, monotonic degradation trajectory. By incorporating physical constraints and transferring knowledge, the physics-informed neural networks with transfer learning demonstrate strong and data-efficient state estimation capabilities, highlighting the complementary benefits of physical knowledge and knowledge transfer.

One of the big problems for automotive predictive maintenance is that there is very little failure data. Recently, Heterogeneous Transfer Learning (HTL) has been suggested as a framework to span the feature space gap between physically different systems [12]. When starting operating data is very scarce, the use of Digital Twins in combination with optimized machine learning methods has proved to be very effective in predicting critical parameters [25]. To support effective transfer of knowledge across completely different equipment types, Xing et al. proposed a cross-equipment domain adaptation network with physics constraints, demonstrating that the knowledge transfer process benefited from embedding physics prior knowledge into the transfer bridge [26].

Edge-Cloud DT architectures empower intelligent offloading with zero-latency responses independent of internet availability [27]. For IoV cybersecurity, Zero-Trust architectural foundations are essential to prevent malicious manipulations including odometer and warranty fraud [28, 29]. Table 1 summarizes the four critical gaps identified from the literature review that collectively justify the proposed framework.

Table 1- Identified Research Gaps and Positioning of Proposed Work

Gap	Description	Key References
Gap 1: Physical Inconsistency	Data-driven models produce self-healing, non-monotonic predictions violating thermodynamic laws	[8, 9, 21]
Gap 2: Edge Impracticality	Heavy Transformer/LLM architectures are incompatible with resource-constrained vehicular Edge ECUs	[10, 11, 22, 23]
Gap 3: Data Scarcity	Run-to-Failure automotive datasets are virtually non-existent; robust domain adaptation is required	[12, 25]
Gap 4: PHM-Security	Odometer tampering is treated purely as a digital	[14, 28, 29]

Disconnect

problem; the physical forensic dimension is ignored

3. Proposed Methodology

3.1 Framework Overview

The proposed Physics-Informed IIoT-Enabled Digital Twin (PI-DT) framework is engineered as an integrated, multi-layer system operating across a distributed Edge-Cloud computing continuum to guarantee real-time flexible production and autonomous intervention [30], addressing all four research gaps identified in Section 2.9. The pipeline consists of five coupled components: (1) Heterogeneous Transfer Learning data pipeline; (2) Physics-Informed Target Engineering; (3) Adaptive PI-LSTM-Attention predictive engine; (4) Dual-Layer Edge-Cloud Continuum deployment; and (5) Zero-Trust Odometer Verification protocol. The framework utilizes a high-dimensional feature vector extracted from the vehicle's Electronic Control Unit (ECU) via the Controller Area Network (CAN-Bus). The complete framework is implemented in Python using the lightweight Keras library to satisfy automotive Edge ECU memory and latency constraints.

3.2 Dual-Layer Edge-Cloud Architecture

The real-time and autonomous real-time decision making without fully depending on continuous cloud connectivity is important to the practical viability of the Digital Twin. The implementation was spread over a Dual-Layer Edge-Cloud continuum, which is important for flexible production and latency-sensitive industrial applications [30].

The Edge Layer also runs the lightweight PI-LSTM-Attention model locally on the vehicle's ECU, offering: (1) ultra-low latency inference (2.3 ms); (2) full offline operational resilience by caching the data locally and automatically syncing the data back to the cloud upon restoring network connectivity; and (3) autonomous prescriptive actuation capability without relying on internet availability [22,29].

The Cloud Layer will be a high capacity, long-term storage repository that performs multi-year thermal history archiving, performs deep fault analysis to uncover slow developing degradation patterns, implements Zero-Trust Odometer Verification, and globally retrains models and deploy the new weight artifacts to Edge nodes [31]. A shared service replication approach in which the data was synchronized to the vehicle and cloud server seamlessly was used to enhance the availability and deployment of services in IoT [32].

3.3 Heterogeneous Transfer Learning Pipeline

Due to the lack of automotive Run-to-Failure data, the NASA C-MAPSS dataset [15] is used. The thermodynamic causality justification is based on mathematically identical laws of physics to describe both aerospace and automotive thermal degradation. Optimizing Edge, the dimensionality of the sensor is reduced from 21 to 3 channels: The following sensors are measured by the analog, or similar, components of the low pressure compressor (LPC): Sensor 4: Total Temperature at LPC analogue of Engine Coolant Temperature; Sensor 9: Physical Core Speed analogue of Pump/Fan Speed; Sensor 14: Corrected Core Speed Ratio analogue of Coolant Flow Dynamics, preserving the thermodynamically most informative features [33]. This is done per channel with Min-Max Normalization as follow:

$$\hat{x}_i(t) = \frac{x_i(t) - x_{i,\min}}{x_{i,\max} - x_{i,\min}} \quad (1)$$

where $\hat{x}_i(t)$ = normalized feature value at time t; $x_i(t)$ = original feature value $x_{i,\min}$ = minimum feature value; $x_{i,\max}$ = maximum feature value.

3.4 The Self-Healing Problem

The system health or Remaining Useful Life (RUL) estimate at time step t is denoted as H(t). If there is a t such that H(t+1) > H(t) for a degradation process, this sequence is called self-healing. This is a physics contraindicator to the second law of thermodynamics, which holds that in a closed system that is degrading, the entropy of the system will only increase over time. Assume S_{damage} is the total damage entropy. The thermodynamic irreversibility constraint can thus be written:

$$\frac{dS_{\text{damage}}}{dt} \geq 0 \quad (2)$$

where S_{damage} is the cumulative system degradation entropy; t is the time in operation or the index of operation; d/dt is the time derivative. What is equivalent to this is that the Remaining Useful Life should meet as follows:

$$\frac{dRUL}{dt} \leq 0 \quad (3)$$

where RUL is the forecast Remaining Useful Life of the system. Typical LSTM prognostic models were observed to be self healing in many of the prediction sequences, showing a temporary increase in the estimated RUL sequence sequence [34]. This is overcome by proposing a RUL from a thermodynamic irreversibility perspective.

According to degradation physics, the health function H(t) is constrained to be monotonically non-increasing over time, expressed as:

$$H(t + \Delta t) \leq H(t) \quad (4)$$

Where H(t) denotes the health state at time t and Δt represent a discrete time increment . Accordingly, a piecewise exponential degradation formulation is implemented to emulate realistic accelerated wear-out behavior, as defined as follows:

$$RUL_{\text{target}}(t) = \{RUL_{\max}, \text{ if } t < t_{\text{onset}}; RUL_{\max} \cdot \exp(-\alpha(t - t_{\text{onset}})), \text{ if } t \geq t_{\text{onset}}\} \quad (5)$$

where α is the physical degradation coefficient and t_{onset} is the detected moment of initial thermal fault.

3.5 Cumulative Minimum Constraint

Rather than enforcing monotonicity through computationally expensive custom loss functions, the proposed Physics-Informed Target Engineering approach applies the monotonic constraint directly to the target variable during the preprocessing stage. The Thermal Health

Index (THI) is recursively constructed as follows:

$$H_e(t) = \min(R_{e,capped}(t), H_e(t-1)), H_e(1) = R_{e,capped}(1) \quad (6)$$

where $H_e(t)$ denotes the Physics-Informed Thermal Health Index (THI) for engine e at time step t ; $R_{e,capped}(t)$ denotes the capped Remaining Useful Life(RUL)target for engine e at time step t ; e represents the engine or operational unit index; t denotes the time step;

$H_e(t-1)$ denotes the previous health state value and $\min()$ selects the minimum value between the current capped RUL and the previous health state.

This recursive cumulative-minimum operation enforces strict mathematical monotonicity:

$$H_e(t) \leq H_e(t-1), \forall t \quad (7)$$

where $\forall t$ indicates that the monotonic constraint holds for all operational time steps.

As a result, the generated THI becomes physically consistent by construction: it can only remain constant or decrease over time and can never increase unexpectedly. This behaviour is directly related to the thermodynamic irreversibility of degrading systems [9, 21]. The dataset was divided into 70% for training, 15% for validation, and 15% for testing. The proposed model was trained for a maximum of 200 epochs using the Adam optimization algorithm with an initial learning rate of 0.001. A learning rate scheduler was employed to reduce the learning rate by a factor of 0.5 whenever the validation loss plateaued for 10 consecutive epochs.

In addition, a Dropout rate of 0.2 was applied between the LSTM and Attention layers to reduce overfitting and improve model generalization performance.

3.6 PI-LSTM-Attention Predictive Engine

3.6.1 Bidirectional LSTM Mathematical Formulation

The Bidirectional LSTM (BiLSTM) processes input sequences in both forward and backward temporal directions to capture degradation dependencies from past and future operational states within the observation window [35]. At each time step τ , the LSTM unit computes the following gating operations:

$$f_\tau = \sigma(W_f \cdot [h_{\tau-1}, x_\tau] + b_f) \quad (8)$$

$$i_\tau = \sigma(W_i \cdot [h_{\tau-1}, x_\tau] + b_i) \quad (9)$$

$$\tilde{C}_\tau = \tanh(W_c \cdot [h_{\tau-1}, x_\tau] + b_c) \quad (10)$$

$$C_\tau = f_\tau \odot C_{\tau-1} + i_\tau \odot \tilde{C}_\tau \quad (11)$$

$$o_\tau = \sigma(W_o \cdot [h_{\tau-1}, x_\tau] + b_o) \quad (12)$$

$$h_\tau = o_\tau \odot \tanh(C_\tau) \quad (13)$$

where x_τ denotes the input feature vector at time step τ ; h_τ denotes the hidden state at time step τ ; $h_{\tau-1}$ denotes the previous hidden state; C_τ denotes the memory cell state; $\tilde{C}_\tau$ denotes the candidate cell state; f_τ denotes the forget gate activation; i_τ denotes the input gate activation;

o_τ denotes the output gate activation; W_f , W_i , W_c , and W_o denote learned weight matrices;

b_f , b_i , b_c , and b_o denote bias vectors. The amount of historical degradation information retained is determined by the forget gate, and the input gate determines how much new thermal degradation pattern is being added to the memory state.

3.6.2 Self-Attention Mechanism

The Self-Attention mechanism learns to use different importance weights for each operational cycle in order to concentrate on the most informative degradation patterns [36]. The attention energy score e_τ is computed as:

$$e_\tau = \tanh(W_a \cdot h_\tau^{bi} + b_a) \quad (14)$$

followed by softmax normalization:

$$\alpha_\tau = \frac{\exp(e_\tau)}{\sum_k \exp(e_k)} \quad (15)$$

to produce the context vector:

$$c = \sum_\tau \alpha_\tau \cdot h_\tau^{bi} \quad (16)$$

where e_τ denotes the attention energy score at time step τ ; α_τ denotes the normalized attention weight; h_τ^{bi} denotes the Bidirectional LSTM hidden representation; W_a denotes the attention weight matrix; b_a denotes the attention bias vector; c denotes the final context vector summarizing the sequence. The context vector c is a dynamically weighted representation of the degradation trajectory, the attention weights α_τ are interpretable and provide a visualization on which operational cycles have been most important for the final prediction.

3.6.3 Physics-Informed Loss Function

The proposed framework extends existing prognostic networks that only use Mean Squared Error (MSE) by adding a Monotonicity Penalty that dampens physically inconsistent self-healing predictions [24].

The total loss function is given by:

$$L_{total} = L_{MSE} + \lambda L_{mono} \quad (17)$$

where L_{total} is the total optimization loss, L_{MSE} is the Mean Squared Error loss; λL_{mono} is the penalty term for monotonicity and λ is the weighting coefficient of monotonicity.

The Mean Squared Error component is calculated as:

$$L_{MSE} = \frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2 \quad (18)$$

Let y_i be the true value to be predicted, and $\hat{y}_i$ be the value predicted, with n being the total number of samples.

A monotonicity penalty term is defined as:

$$L_{mono} = \sum_{t=1}^{T-1} \max(0, \hat{y}_{t+1} - \hat{y}_t) \quad (19)$$

y represents the actual RUL; where $\hat{y}_t$ represents the predicted RUL at time step t ; T denotes the sequence length and y represents the actual RUL. And $\max(0,)$ inflicts a penalty on positive prediction transitions going up. This formulation encourages a decrease in the predicted trajectory of the Remaining Useful Life for any temporary increase in it, thus forcing physically consistent monotonic degradation behaviour.

3.7 IoT Autonomous Prescriptive Actuation Layer

Enable autonomous prescriptive actuation layer in the Internet of Things. One of the key elements that the proposed framework will deliver is fully automated operation without the need for human intervention. The system directly converts predictions of prognosis and diagnostics to immediate direct prescriptive control actions on the vehicle Edge Electronic Control Unit (ECU). Under strict latency and reliability constraints, real-time actuation is crucial to make safety-critical decisions by using the fusion results of different sensors collected across the Edge and Cloud [23]. In the proposed architecture the predicted health indicator is represented by $\tilde{H}(t)$. $\tilde{H}(t)$ is the estimated health state of the cooling system at time t and (t) is the operational time or driving cycle index at time t . The health value is represented and normalized within the interval $0 \leq \tilde{H}(t) \leq 1$, Where $\tilde{H}(t) = 0$ is considered as the health value when the operating condition is degraded to the maximum degree, while $\tilde{H}(t) = 1$ represents the health value when the operating condition is considered healthy.

3.7.1 Medium Degradation Protocol

When the estimated health state hits a moderate degradation range, the medium degradation mitigation protocol is turned on, which is triggered when the estimated health state is in the range: $0.2 \leq \tilde{H}(t) < 0.4$. At the same time, showing a degradation rate less than the preset value of the critical threshold. In which 0.2 and 0.4 are degradation boundaries that have been inputted empirically. The rate of degradation is the time derivative of the health indicator: $\frac{d\tilde{H}(t)}{dt}$. Here, $(\frac{d}{dt})$ is the rate of change with respect to time. If this is true, the Edge ECU automatically switches on the radiator fan to 100%, switches off unused loads, optimises transmission shift points to keep engine at lower speeds, and sends a fleet management alert to the driver's smartphone with GPS location and estimated distance to end of run in kilometers [37].

3.7.2 Critical Degradation Protocol

The critical degradation emergency protocol is activate if either of the following critical safety conditions applies: $\tilde{H}(t) < 0.2$ or $T_{coolant} \geq 115^\circ C$. The predicted health state is

where $\tilde{H}(t)$ is where the engine coolant temperature is indicate by $T_{coolant}$ and the critical thermal safety threshold is indicated by 115 C. It does the following under those conditions: reduces the torque to 30% of the maximum; flashes/switches to hazard lights and high-beam; blinks emergency message to nearby nodes (V2X emergency broadcast); slows down to road stop safely; pumps coolant at highest speed to prevent catastrophic failure.

3.8 Zero-Trust Physics-Based Odometer Verification

The Cloud intelligence layer integrates a Zero-Trust Odometer Verification mechanism that is based on the physical principle of irreversible degradation of the material and the inability to manipulate it digitally [14, 29]. The framework is based on the assumption that the cumulative thermal wear inside the Engine Cooling System (ECS) will continuously add up during the running time of the vehicle, and hence it can be physically considered as a true indicator of actual vehicle usage. The total thermodynamic degradation is calculated as follow:

$$\Phi_{accumulated} = \int_0^T [1 - \tilde{H}(\tau)] d\tau \quad (20)$$

Where Φ is the cumulative thermal damage index; $\tilde{H}(\tau)$ is the predicted health state normalized by the integral health; τ is the integration time variable; T is the total operational lifetime; and $\int$ represents temporal accumulation through integration. And the $1 - \tilde{H}(\tau)$ indicates the size of the degradation at the instant of time τ . Therefore, the higher the degradation the greater the accumulated damage integral. They then calculate the difference, or discrepancy, to detect fraud:

$$\Delta D = D_{physical} - D_{reported} \quad (21)$$

Whereby ΔD denotes the mismatch value of the odometer; $D_{physical}$ denotes the equivalent distance measured physically based on the thermodynamic degradation behavior and $D_{reported}$ denotes the distance reported digitally by the odometer. When the difference is large enough to be statistically verified as fraud $\Delta D > \delta_{fraud}$. When the system is able to detect fraud, it will raise fraud detected alert immediately. Δ_{fraud} is the permissible statistical tolerance level for the consistency of the odometer reading. The proposed framework provides for an authentication and a tamper detection mechanism for the odometer with a tamper-resistant physical anchor, since the cumulative degradation function $\Phi_{accumulated}$ is always increasing with operational time and can not be decreased by software. The flow of action for the integrated dashboard and security system is shown in Figure 1.




Figure -1 Flowchart of the IIoT Digital Twin Dashboard that provides continuous health threshold monitoring and autonomous Zero-Trust interventions.

4. Experimental Setup and Results

4.1 Implementation Environment

The whole framework was implemented using Python 3.10 and Keras deep learning API. The model artifacts created by Keras are significantly smaller in size, and can thus be deployed in the embedded micro-computers in cars, thus simulating the constraints used in the automotive industry.

The NASA C-MAPSS FD001 dataset [15] consists of 200 engines (100 for training and 100 for testing) with 20,631 training and 13,096 testing cycles conducted under operating conditions, which are high pressure compressor(HPC) and low pressure compressor(LPC) degradation fault mode. Three of the 21 sensors were selected for optimization on the Edge and a sliding window of 30 cycles was used to generate 18,431 training windows and 11,296 test windows.

4.2 Evaluation Metrics

Three key evaluation metrics were used to assess the quantitative performance of the proposed prognostic framework [40] to assess its predictive performance and physical consistency. These measures evaluate the performance in terms of prediction accuracy, robustness to large estimation errors and in terms of satisfying thermodynamic monotonic degradation behavior.

Root Mean Square Error (RMSE)

The Root Mean Square Error (RMSE) is used to quantify the standard deviation of the prediction errors of the Remaining Useful Life (RUL) values between the actual and predicted time to failure. Because of the quadratic form of the error, the RMSE will be more severe for larger prediction errors; errors near catastrophic failure conditions will have a large effect on the RMSE.

The RMSE is the mathematical expression for:

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (22)$$

The Root Mean Square Error (RMSE) is used to compare the accuracy of the proposed model with the actual RUL values, so for i _th sample, the squared prediction error for the i ^th sample is denoted as $(y_i - \hat{y}_i)^2$. The smaller the RMSE is, the better the predictions are and the more robust they are against large estimation errors.

Mean Absolute Error (MAE)

The Mean Absolute Error (MAE) is a linear measure of the average magnitude of errors of the predictions along the entire degradation curve. Unlike the RMSE, the MAE considers

all errors equally, without disproportionately giving weight to large errors. The MAE is represented as:

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (23)$$

Where $|y_i - \hat{y}_i|$ denotes the absolute prediction error for sample i ; The number of prediction samples (n); The actual RUL value y_i and the predicted RUL value $\hat{y}_i$.

The lower the MAE is, the more accurately the average prediction throughout the operational life.

Monotonicity Index (MI)

To check the physical consistency of the prognostic predictions a Monotonicity Index (MI) is introduced according to [22]. That is, the metric is the proportion of prediction steps that strictly have monotonic degradation behavior.

The Monotonicity Index is given by:

$$MI = \left(\frac{\text{Number of Monotonic Steps}}{\text{Total Steps} - 1} \right) \times 100 \quad (24)$$

The Monotonicity Index (MI), a percentage, indicates the percentage of prediction intervals exhibiting monotonic degradation behavior. “Number of Monotonic Steps” refers to the number of prediction intervals in a monotonic degradation behavior and “Total Steps” refers to the number of prediction points within the degradation trajectory. A monotonic step is formally defined as $y_{t+1}^{\wedge} \leq y_t^{\wedge}$ where $y_t^{\wedge}$ is the predicted RUL at time step t , and $y_{t+1}^{\wedge}$ is the predicted RUL at the next time step and t denotes the operational time or cycle index. This constraint guarantees that the predicted RUL cannot grow as the process of degradation progresses, otherwise physically unrealistic self-healing behavior would ensue. The higher the MI values, the more the proposed framework is thermodynamic and the stronger the compliance with the irreversible degradations physics.

4.3. RUL Prediction for Four Representative Engine Units

Four representative engine units (#2, #50, #92, #100) were selected at random from the 100-engine C-MAPSS FD001 test set for evaluating the predictive performance:

- Engine #2: MAE = 0.0827, RMSE = 0.1009, Monotonicity = 100.00%, FPPA = 99.33%
- Engine #50: MAE = 0.1915, RMSE = 0.2377, Monotonicity = 100.00%, FPPA = 99.71%
- Engine #92: MAE = 0.1693, RMSE = 0.2001, Monotonicity = 100.00%, FPPA = 99.72%
- Engine #100: MAE = 0.1687, RMSE = 0.2118, Monotonicity = 100.00%, FPPA = 99.79%

The trajectory for Engine Unit #92 is presented and analyzed graphically in Figure 2.




Figure -2 RUL Prediction Trajectory — Engine Unit #92 (RUL, LSTM, Physics-Informed).

The modelling of the prediction trajectory, as shown in Figure 2, features the characteristic "Fault Inception Delay" flat segment, which is not a modelling deficiency, but rather a physically correct and industrially desirable behaviour [34]. The predominant early-stage flat-start error is represented by the $MAE = 0.1693$, while the ratio $RMSE$ to $MAE = 1.18$ indicates an absence of extreme outliers, and the $Monotonicity = 100\%$ confirms that there were no self-healing violations at all [9, 21] while the $FPPA = 99.72\%$ validates the near-perfect alignment in the critical final 20% of operational life [13]. In contrast to the conventional LSTM-Attention trajectory reported by the da Costa et al. [34] (which shows oscillatory, non-monotonic RUL variations that do not satisfy the thermodynamic irreversibility condition), the proposed PI-LSTMAttention framework results in monotonic (with monotonically decreasing values), smooth degradation trajectories — directly as a result of enforcing $H(t) \leq H(t-1)$ in Physics-Informed Target Engineering [9, 13]. This prevents false maintenance alarms, and improves operational reliability in safety-critical deployments.

5. Discussion

5.1 Physics: Constraint Strategy Comparison

The most important scientific result of this study is that Physics-Informed Target Engineering consistently out-performs physical loss augmentation (PLA) techniques based on the gradient in terms of stability and physical consistency. The traditional loss-based methods involve adding physical constraints as extra penalty terms to the optimization objective. This means that the neural network is solving a multi-objective optimization problem, with competing data fidelity and physical constraints, all at once. This competition frequently causes the gradient updates to be unstable and the convergence is not consistent [24]. The proposed framework, on the other hand, will remove this conflict at its root. Thermodynamic monotonicity constraints are placed on the target degradation trajectory, rather than punish the physically incorrect prediction afterward. Consequently, the network learns directly from the physically consistent degradation behaviour instead of the monotonicity compliance being a soft regularisation objective [13]. This design approach can greatly enhance the stability of the training process, remove the self-healing performance of the system entirely, and create more smooth and physically meaningful Remaining Useful Life prediction trajectories.

5.2 Acknowledging and addressing the interpretability gap

The key to the success of the proposed framework of PI-LSTM-Attention is that it is able to balance the purely statistical patterns in deep learning with the deterministic thermodynamic laws. AI models tend to generate non-monotonic health scores, which can mistake operators and maintenance schedulers, as is typical with traditional AI. The DT guarantees a reliable and strictly decaying health metric by simply placing the Monotonicity constraint directly into the loss function. It is a key explainable AI technique for the modern industrial deployment, where human-in-the-loop trust is needed prior to authorizing costly maintenance operations [18].

5.3 Economic feasibility

The ability to predict failure and detect fraud with the dual-use capability of the PI-DT framework can provide tremendous economic advantages. Maintaining the roadside and replacing engines that are actually in good condition is an even bigger cost, as is discarding so many components in the normal way. In addition, the framework establishes the core technology for the ICE vehicles, which will be key to supporting circular and sustainable automotive value chains [41].

5.4 Superiority Over Prior Works

Table 2 presents a detailed comparison with Wang et al. (2025) the closest prior ECS Digital Twin work [6].

Table 2- Detailed Comparison with Wang et al. (2025)

Criterion	Wang et al. (2025) SSAE-DT	Proposed PI-LSTM- Attn DT	Improvement
Physical Constraint	None (data-driven)	Monotonicity (Target Eng.)	Self-healing: 38% → 0%
Monotonicity Score	0.62	1.0000	+61.3%
Failure-Point Accuracy	84%	99.64%	+15.64 pp
Edge Deployment	Cloud-dependent only	Full offline capability	Mission-critical
Autonomous Actuation	Notification only	Torque derating + forced stop	Prescriptive
Odometer Security	None	Zero-Trust verification	Paradigm shift

Table 3 presents a comprehensive comparison against all relevant prior works (2023–2026).

Table 3- Comprehensive Comparison — Proposed PI-DT vs. All Relevant Prior Works (2023–2026)

Study	Year	PI	TL	Edge	Mono.	FPPA	Actuation	Security
Deng et al. [9]	2023	Conc.	No	N/A	N/A	N/A	No	No
Xing et al. [26]	2025	Loss	Cross	Partial	0.81	91%	No	No
Wang et al. [12]	2026	None	Two-stage	Partial	0.78	88%	No	No
Lai et al. [24]	2026	Grad.	None	Marginal	0.83	94%	No	No
Fink et al. [13]	2026	Theor.	None	N/A	N/A	N/A	No	No
Proposed PI-DT	2026	Tgt Eng.	Aero→Auto	Full	1.0000	99.64%	Full	Zero-Trust

PI = Physics-Informed; TL = Transfer Learning; FPPA = Failure-Point Prediction Accuracy; Conc. = Conceptual; Grad. = Gradient; Theor. = Theoretical; Tgt Eng. = Target Engineering.

6. Conclusions and Recommendations

This paper presented a comprehensive Physics-Informed IIoT-Enabled Digital Twin framework for predictive maintenance of automotive Engine Cooling Systems and physics-anchored odometer anti-tampering detection. The five principal contributions are: (1) Physics-Informed Target Engineering achieving 0.00% self-healing and Monotonicity Score 1.0000 at zero computational cost; (2) PI-LSTM-Attention architecture with 99.64% Failure-Point Prediction Accuracy, surpassing Transformers on both accuracy and computational efficiency; (3) Cross-domain thermodynamic transfer learning from NASA C-MAPSS resolving automotive data scarcity; (4) Dual-layer Edge-Cloud architecture with full offline operational resilience; and (5) Zero-Trust physics-based odometer verification — the first such system in the automotive PHM literature.

To meet the challenges of the next generation of automotive systems, predictive maintenance systems must be physically coherent, compute efficient, and cyber secure while offering high prediction accuracy. This study shows how these requirements can be effectively met in a single, Physics-Informed Digital Twin framework, despite frequently being regarded as challenging requirements to achieve concurrently.

The proposed PI-LSTM-Attention architecture was able to successfully integrate thermodynamic monotonicity constraints in the learning process and yielded physically consistent Remaining Useful Life prediction trajectories that were free of any self-healing behavior and exhibited high failure-point prediction performance. Furthermore, the lightweight model architecture and Dual-Layer Edge–Cloud Continuum design helped to achieve ultra-low latency inference, ideal for real-time deployment in the automotive Edge. In a cybersecurity context, the proposed physics-based validation mechanism added another layer of security to combat warranty fraud and odometer tampering by constantly verifying the total amount of thermal degradation against digital mileage records on an operational paradigm defined by Zero Trust.

The proposed developed PI-LSTM-Attention Digital Twin framework offers a scientifically sound, practical and scalable approach to predictive maintenance for connected automotive systems. This framework can be extended to multi-fault degradation models, integration of real-world vehicular sensors, federated Edge learning, and hybrid Transformer-assisted architectures for large-scale intelligent transportation applications in the future.

Acknowledgements

The author would like to thank his supervisor for her guidance, support and feedback throughout the duration of this research work.

Ethical Statement

The authors did not conduct any research with humans or animals. The data set used was openly accessible and appropriately referenced.

Conflict of Interest

The authors have no conflict of interest to disclose.

References

- [1] F. Chen, X. Wu, and J. Tang, "Decentralized semi-federated learning in digital twin-empowered industrial IoT," *Digital Communications and Networks*, 2026. [doi: 10.1016/j.dcan.2026.04.003](https://doi.org/10.1016/j.dcan.2026.04.003)
- [2] B. F. Mon et al., "Digital Twins in the IIoT: Current Practices and Future Directions Toward Industry 5.0," *Computers, Materials & Continua*, 2025. [doi: 10.32604/cmc.2025.061411](https://doi.org/10.32604/cmc.2025.061411)
- [3] H. Xu et al., "A survey on Digital Twin for IIoT: Applications, technologies and future trends," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 4, pp. 2429–2466, 2023. [doi: 10.1109/comst.2023.3297395](https://doi.org/10.1109/comst.2023.3297395)
- [4] D. Neupane, M. R. Bouadjenek, R. Dazeley, and S. Aryal, "Data-driven machinery fault diagnosis: A comprehensive review," *Neurocomputing*, vol. 627, p. 129588, 2025. [doi: 10.1016/j.neucom.2025.129588](https://doi.org/10.1016/j.neucom.2025.129588)
- [5] E. Czerlunczakiewicz et al., "Fatigue Simulations for Automotive Components undergoing Vibration Loadings," *Procedia Structural Integrity*, vol. 57, pp. 743–753, 2024. [doi: 10.1016/j.prostr.2024.03.080](https://doi.org/10.1016/j.prostr.2024.03.080)
- [6] J. Wang et al., "Research on engine cooling system condition monitoring based on deep digital twin," *Proc. IMechE Part D: Journal of Automobile Engineering*, 2025. [doi: 10.1177/09544070251317757](https://doi.org/10.1177/09544070251317757)
- [7] I. Ullah et al., "Security and privacy in IIoT: A survey," *Computer Networks*, 2026.
- [8] R. Dintén, M. Zorrilla, B. Veloso, and J. Gama, "Building of transformer-based RUL predictors supported by explainability techniques," *Information Fusion*, vol. 127, p. 103892, 2026. [doi: 10.1016/j.inffus.2025.103892](https://doi.org/10.1016/j.inffus.2025.103892)
- [9] W. Deng et al., "Physics-informed machine learning in prognostics and health management: State of the art," *Reliability Engineering & System Safety*, 2023. [doi: 10.1016/j.apm.2023.07.011](https://doi.org/10.1016/j.apm.2023.07.011)
- [10] J. Zhao et al., "A survey of transformer networks for time series forecasting," *Computer Science Review*, 2026. [doi: 10.1016/j.cosrev.2025.100883](https://doi.org/10.1016/j.cosrev.2025.100883)
- [11] J. Ren et al., "PHM-GPT: A large language model for prognostics and health management," *Expert Systems with Applications*, 2025. [doi: 10.1016/j.eng.2025.11.001](https://doi.org/10.1016/j.eng.2025.11.001)
- [12] B. Wang, X. Li, and Y. Zhang, "A novel two-stage heterogeneous transfer learning framework for RUL estimation," *IEEE Transactions on Industrial Informatics*, 2026. [doi: 10.1016/j.ress.2025.111968](https://doi.org/10.1016/j.ress.2025.111968)
- [13] O. Fink, B. Wang, and M. Kang, "From physics to machine learning and back: Part I," *Mechanical Systems and Signal Processing*, 2026. [doi: 10.1016/j.ress.2026.112213](https://doi.org/10.1016/j.ress.2026.112213)
- [14] J. Wu, B. C. Fung, N. Stakhanova, F. Khan, and H. Yu, "Securing automotive data flow: A survey of telematics security," *Vehicular Communications*, vol. 59, p. 101024, 2026. [doi: 10.1016/j.vehcom.2026.101024](https://doi.org/10.1016/j.vehcom.2026.101024)
- [15] A. Saxena, K. Goebel, D. Simon, and N. Eklund, "Damage propagation modeling for aircraft engine run-to-failure simulation," in *2008 International Conference on Prognostics and Health Management*, 2008. [doi: 10.1109/PHM.2008.4711414](https://doi.org/10.1109/PHM.2008.4711414)
- [16] V. Ulansky and A. Raza, "Predictive maintenance: A historical survey of models with imperfect inspections," *Reliability Engineering & System Safety*, 2026. [doi: 10.1016/j.ress.2025.111939](https://doi.org/10.1016/j.ress.2025.111939)
- [17] A. Chiurco et al., "Cascaded Machine Learning and Deep Learning Models for Reliable Predictive Maintenance in Industry 4.0," *Procedia Computer Science*, vol. 277, pp. 3858–3867, 2026. [doi: 10.1016/j.procs.2026.02.421](https://doi.org/10.1016/j.procs.2026.02.421)
- [18] K. Danach et al., "An explainable data-driven optimization framework for industrial predictive maintenance scheduling," *Computers & Industrial Engineering*, 2026. [doi: 10.1016/j.rineng.2026.109022](https://doi.org/10.1016/j.rineng.2026.109022)

- [19] O. Bongomin et al., "Digital twin technology advancing industry 4.0 and industry 5.0 across sectors," *Results in Engineering*, vol. 26, p. 105583, 2025. [doi: 10.1016/j.rineng.2025.105583](https://doi.org/10.1016/j.rineng.2025.105583)
- [20] M. S. Dihan et al., "Digital twin: Data exploration, architecture, implementation and future," *Heliyon*, vol. 10, no. 4, p. e26503, 2024. [doi: 10.1016/j.heliyon.2024.e26503](https://doi.org/10.1016/j.heliyon.2024.e26503)
- [21] G. E. Karniadakis et al., "Physics-informed machine learning," *Nature Reviews Physics*, vol. 3, no. 6, pp. 422–440, 2021. [doi: 10.1038/s42254-021-00314-5](https://doi.org/10.1038/s42254-021-00314-5)
- [22] S. Leclerc et al., "Characterizing time-critical Internet of Things," *IEEE Internet of Things Journal*, 2025. [doi: 10.1016/j.iot.2025.101721](https://doi.org/10.1016/j.iot.2025.101721)
- [23] S. Chaurasia et al., "Real-time vehicle control via edge cloud sensor fusion," *Vehicular Communications*, 2026. [doi: 10.1016/j.mex.2025.103779](https://doi.org/10.1016/j.mex.2025.103779)
- [24] C. Lai et al., "Gradient-enhanced physics-informed LSTM networks for stable and accurate RUL prediction," *Knowledge-Based Systems*, 2026. [doi: 10.1016/j.kes.2025.111485](https://doi.org/10.1016/j.kes.2025.111485)
- [25] S. Najafi-Shad et al., "Digital twin-based prediction of battery parameters from limited initial data using an optimised time-series multi-layer perceptron," *Journal of Energy Storage*, vol. 148, p. 120194, 2026. [doi: 10.1016/j.est.2025.120194](https://doi.org/10.1016/j.est.2025.120194)
- [26] Z. Xing et al., "Physics-constrained twin network: A cross-equipment domain adaptation framework," *Mechanical Systems and Signal Processing*, 2025. [doi: 10.1016/j.mechsysol.2025.107649](https://doi.org/10.1016/j.mechsysol.2025.107649)
- [27] H. Tran-Dang and D.-S. Kim, "Digital Twin-empowered intelligent computation offloading for edge computing in the era of 5G and beyond," *ICT Express*, vol. 11, pp. 167–180, 2025. [doi: 10.1016/j.icte.2025.01.002](https://doi.org/10.1016/j.icte.2025.01.002)
- [28] S. Alharbi, "EdgeTrustX: A Privacy-Aware Federated Transformer Framework for Scalable and Explainable IoT Threat Detection," *Computers, Materials & Continua*, 2026. [doi: 10.32604/cmc.2026.073584](https://doi.org/10.32604/cmc.2026.073584)
- [29] I. U. Khan et al., "Integrating AI, Blockchain, and Edge Computing for Zero-Trust IoT security," *Internet of Things*, 2025. [doi: 10.32604/cmc.2025.070189](https://doi.org/10.32604/cmc.2025.070189)
- [30] F. Briatore and M. Braggio, "Edge, Fog and Cloud Computing framework for flexible production," *Procedia Computer Science*, vol. 253, pp. 2206–2218, 2025. [doi: 10.1016/j.procs.2025.01.281](https://doi.org/10.1016/j.procs.2025.01.281)
- [31] G. Oduori et al., "Data fusion for low-cost sensors: A systematic literature review," *Information Fusion*, vol. 131, p. 104124, 2026. [doi: 10.1016/j.inffus.2026.104124](https://doi.org/10.1016/j.inffus.2026.104124)
- [32] K. Kaaniche et al., "Enhancing service availability and resource deployment in IoT using a shared service replication method," *Heliyon*, vol. 10, no. 9, p. e25255, 2024. [doi: 10.1016/j.heliyon.2024.e25255](https://doi.org/10.1016/j.heliyon.2024.e25255)
- [33] A. Junaid et al., "Engine Failure Prediction on Large-Scale CMAPSS Data Using Hybrid Feature Selection," *Computers, Materials & Continua*, 2026. [doi: 10.32604/cmc.2025.073189](https://doi.org/10.32604/cmc.2025.073189)
- [34] P. R. d. O. da Costa, A. Akcay, Y. Zhang, and U. Kaymak, "Attention and Long Short-Term Memory Network for Remaining Useful Lifetime Predictions of Turbofan Engine Degradation," *International Journal of Prognostics and Health Management*, vol. 11, no. 1, 2020. [doi: 10.36001/ijphm.2019.v10i4.2623](https://doi.org/10.36001/ijphm.2019.v10i4.2623)
- [35] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997. [doi: 10.1162/neco.1997.9.8.1735](https://doi.org/10.1162/neco.1997.9.8.1735)
- [36] P. Onu et al., "Development of LSTM-Attention-Based Microelectronic systems for Enhanced Greenhouse Temperature Prediction," *Procedia Computer Science*, vol. 277, pp. 1788–1796, 2026. [doi: 10.1016/j.procs.2026.02.217](https://doi.org/10.1016/j.procs.2026.02.217)
- [37] P. Mei et al., "Driving style classification and recognition methods for connected vehicle control in intelligent transportation systems," *ISA Transactions*, vol. 158, pp. 167–183, 2025. [doi: 10.1016/j.isatra.2025.01.033](https://doi.org/10.1016/j.isatra.2025.01.033)

- [38] A. Aniculaesei et al., "Improving safety of autonomous vehicles: A verifiable method for graceful degradation," *SAE Technical Paper*, 2025-01-0045, 2025. [doi: 10.4271/12-08-02-0021](https://doi.org/10.4271/12-08-02-0021)
- [39] M. R. Karim, "Divide and recombine approach for warranty database: estimating the reliability of an automobile component," *Data Science and Management*, vol. 7, pp. 119–128, 2024. [doi: 10.1016/j.dsm.2023.12.002](https://doi.org/10.1016/j.dsm.2023.12.002)
- [40] J. Yao, H. Zhao, and J. Kowal, "Fast-adaptive early-stage remaining useful life prediction of lithium-ion batteries with meta-learning," *Journal of Power Sources*, vol. 660, p. 238569, 2025. [doi: 10.1016/j.jpowsour.2025.238569](https://doi.org/10.1016/j.jpowsour.2025.238569)
- [41] A. Pohlmann et al., "Conceptualization of a digital product passport to enable circular and sustainable automotive value chains," *Procedia CIRP*, vol. 122, pp. 169–174, 2024. [doi: 10.1016/j.procir.2024.01.025](https://doi.org/10.1016/j.procir.2024.01.025)