

Beyond Prediction by Integrating Physics-of-Failure Analysis with Predictive Maintenance for Failure Prevention: An Empirical Case Study and Integrated Framework

Witantyo^{1*}, Taufik Hidayat², Ahmad Nur Muzaqin¹

¹Mechanical Engineering, Institut Teknologi Sepuluh Nopember (ITS), Surabaya, Indonesia

²Naval Architecture, Institut Teknologi Kalimantan (ITK), Balikpapan, Indonesia

Received: 3 August 2026, Revised: 25 September 2026, Accepted: 30 September 2026

Abstract

Predictive maintenance (PdM) has become the dominant paradigm in industrial maintenance research, driven by IoT, machine learning, and data analytics. This narrative review argues that the field's near-exclusive focus on predicting when failures will occur has overshadowed a more fundamental question: why do failures happen? Examining PdM research (2010–2025) reveals a persistent gap: most predictive models detect degradation without addressing the physical mechanisms that cause failure, so even accurate predictions do not prevent recurrence. A physics-of-failure (PoF) approach—rooted in root cause failure analysis (RCFA) and reliability-centered maintenance (RCM)—offers a complementary, more durable solution when integrated with data-driven methods. A case study of 7,300 recurring spring leaf failures in a heavy truck fleet illustrates the gap: a single overlooked mechanism—shock absorber wear causing increased dynamic stress on springs—perpetuated a cycle that no predictive model could break. An integrated three-layer framework combining data-driven monitoring, mechanism-based root cause analysis, and design feedback is proposed; it is a design proposal, not yet an evaluated operational system. The projected failure reduction and return-on-investment derive from fatigue-life calculations rather than observed field data, because the intervention followed data collection; validating the framework and these projections with before-and-after field data is a critical area of future work.

Keywords: predictive maintenance, physics of failure, root cause analysis, reliability-centered maintenance, truck fleet

1. Introduction

Maintenance is one of the largest operational costs in asset-intensive industries. In the transportation sector alone, maintenance accounts for approximately 59% of total vehicle operating costs [1]. For a typical heavy truck fleet, unplanned downtime from component failures can cost thousands of dollars per hour in lost revenue, emergency repairs, and cascading supply chain disruptions. The economic burden extends beyond direct repair costs to include production delays, logistics disruptions, and safety incidents [2,3].

The evolution of maintenance strategies has followed a well-documented trajectory: reactive maintenance (run-to-failure), where components are repaired only after failure; preventive maintenance (time-based intervals), which reduces unexpected failures but often results in unnecessary replacements; condition-based maintenance (CBM), using sensor data to assess actual equipment condition; and predictive maintenance (PdM), using data-driven models to forecast failures [4]. This evolution

reflects a broader shift from time-based to condition-based and predictive strategies across industrial sectors. The past decade has seen an explosion of PdM research, fueled by affordable IoT sensors, cloud computing, and advances in deep learning [5–7].

However, this review argues that the field has developed a critical blind spot. Most PdM systems are designed to answer “when will this component fail?” but not “why does this component keep failing?” As a result, even when predictions are accurate, the same failures recur—often on the same equipment, from the same root causes. We argue that a physics-of-failure (PoF) approach, grounded in systematic root cause failure analysis (RCFA) and reliability-centered maintenance (RCM), offers a complementary solution. By understanding the physical mechanisms that drive component degradation—fatigue, wear, corrosion, overload, creep—and eliminating their root causes, organizations can achieve not just prediction, but genuine prevention.

This paper makes the following central claim: pre-

*Corresponding author. Email: witantyo@its.ac.id, phone: +628165405389

© 2027. The Authors. Published by DRPM ITS.

dictive maintenance, as currently practiced, primarily addresses the symptom of failure (when it will occur) rather than its cause (why it occurs). Through a critical narrative review of PdM research (2010–2025) and a case study of 7,300 recurring spring leaf failures in a heavy truck fleet, we demonstrate that a single overlooked failure mechanism—shock absorber wear causing increased dynamic stress on springs—perpetuated a cycle that no predictive model could break. We further show that PoF and data-driven PdM are complementary approaches, and propose an integrated three-layer framework that combines the strengths of both. Section 2 describes the review methodology and case-study methods. Section 3 presents results and discussion. Section 4 provides conclusions.

2. Review Methodology and Case-Study Methods

2.1. Review Approach

This paper is a critical narrative review of predictive maintenance research (2010–2025), complemented by a case study of recurring spring leaf failures in a heavy truck fleet. The review was conducted through targeted searches of Scopus, IEEE Xplore, and Google Scholar using the following search terms: “predictive maintenance” AND (“heavy truck” OR “fleet”), “physics of failure” AND “maintenance”, and “root cause analysis” AND “predictive maintenance”. Studies were included if they addressed data-driven PdM methods for vehicle or fleet applications published between 2010 and 2025. The search was run on 15 June 2025 and last updated on 20 June 2025. The initial query returned 312 records (Scopus: 161; IEEE Xplore: 89; Google Scholar: 62); after removal of 47 duplicates, 265 records were screened by title and abstract. Records were excluded if they (i) were not peer-reviewed or were non-English, (ii) addressed stationary plant, manufacturing lines, or components other than vehicles or fleets, (iii) were purely conceptual with no method or dataset, or (iv) did not report at least one quantitative performance metric. This yielded 58 full-text articles, which were assessed against the inclusion criteria and evaluated for study quality on the basis of dataset size, validation methodology, and reproducibility of results. Fourteen studies were retained for detailed analysis, selected for their relevance to the research question and for methodological diversity (classical ML, deep learning, digital-twin, Bayesian, and expert-system approaches). The review is narrative rather than systematic: the search and selection process is documented and reproducible, but no claim of exhaustive coverage is made, and the 14 studies serve as representative examples of persistent gaps rather than a comprehensive census of the field.

2.2. Data Collection and Fleet Description

The primary dataset comprises maintenance records from a fleet of 35 Volvo FMX 440/460 heavy trucks operating in Indonesian mining and logistics applications

over a four-year period (2022–2025). The fleet operates in open-pit mining and haul road conditions with high ambient temperatures (28–38 °C), heavy dust exposure, and unpaved road surfaces. Maintenance records include component replacement events, downtime hours, repair costs, and vehicle identification numbers. Fleet operational data was obtained from the company’s maintenance management system, covering 7,300 spring leaf failure events across all 35 units.

The analysis focuses on spring leaf failures as the primary failure mode. One unit alone (DT FMX 440-2021) recorded 342 spring failures over the four-year period, corresponding to approximately 1.64 failures per week (342 failures ÷ 208 weeks). A “failure event” is defined as a documented spring leaf fracture requiring replacement, verified against maintenance work orders and parts requisition records. Duplicate entries (e.g., same event recorded in both work order and inventory systems) were removed during data cleaning. Each failure resulted in 2–4 hours of downtime and direct repair costs of approximately Rp 850,000–3,350,000 per event. Data collection followed a structured protocol: failure date, vehicle identification, component position, operating hours at failure, and repair actions taken. Failure rates were normalized by operating hours where available; the fleet accumulated approximately 1.2 million operating hours over the four-year period, yielding a fleet-wide spring failure rate of approximately 6.1 failures per 1,000 operating hours.

2.3. Root Cause Failure Analysis (RCFA) Methodology

Root cause failure analysis was conducted using a multi-method approach. The 5-Whys technique was applied iteratively to each recurring failure event to trace the causal chain from the observed failure mode back to its latent root cause. Ishikawa (fishbone) diagrams were constructed for the spring leaf failure mode, organizing potential causes into the standard categories: man, machine, material, method, measurement, and environment. Fault tree analysis (FTA) was used to evaluate the logical relationships between failure events and their contributing factors. The FTA was conducted qualitatively to identify the most probable root cause pathways; quantitative probability estimation was not performed due to the absence of component-level reliability data for the specific operating conditions [9,10].

The RCFA process was complemented by field observations, maintenance staff interviews, and review of operating procedures. This mixed-methods approach ensured that both technical and organizational root causes were captured, following the methodology described by Nowlan and Heap [10] in the foundational RCM framework.

2.4. Failure Analysis Approach

Physical failure analysis was conducted on failed spring leaf specimens to identify the failure mechanism. The methodology included: (1) visual inspection of frac-

ture surfaces at magnifications of $50\times$ to $500\times$ to identify fatigue beach marks and striations; (2) material composition analysis using optical emission spectroscopy (OES) to verify compliance with SUP9A spring steel specifications (0.56–0.64% C, 1.30–1.60% Si, 0.70–1.00% Mn); (3) hardness testing (HRC) at the fracture origin and bulk material to detect softening or hardening zones; and (4) dimensional measurement of worn shock absorber piston rods and seals to quantify the degree of wear. The 20,000 km inspection/replacement interval used later in the prevention strategy is an author-derived value, not a verbatim quotation from a single manufacturer document. It is informed by (i) OEM and industry maintenance practice: the Volvo Trucks preventive-maintenance program includes dedicated front- and rear-shock-absorber inspections in its scheduled annual service [20]; and current industry guidance for Asia-Pacific heavy fleets recommends regular scheduled shock-absorber inspection, treating visible oil leakage — beyond normal seal misting — as grounds for replacement, with damping performance affected once roughly 10% of the oil volume is lost [21]. It also reflects (ii) the fleet's own field records, in which seal-related oil loss of 40–60% of nominal capacity and piston-rod wear beyond the 0.10 mm service limit were already observed at approximately 20,000 km. The primary literature basis comprises references [20] and [21]; the 20,000 km value itself is derived from the fleet's wear data.

The results of the failure analysis are presented in Section 3.4. Visual inspection confirmed classic fatigue fracture surfaces with beach marks and striations on all examined specimens. OES analysis confirmed that the spring material met SUP9A specifications. Hardness testing showed no significant deviation between fracture origin and bulk material (HRC 38–42 vs. HRC 39–43), ruling out material degradation as a primary cause. Shock absorber measurements revealed that units with high spring failure rates had piston rod wear of 0.15–0.30 mm (beyond the 0.10 mm service limit) and seal wear resulting in hydraulic oil loss of 40–60% of nominal capacity, confirming shock absorber degradation as the critical contributing factor. These values lie within the typical reported hardness range for quenched-and-tempered SUP9A/SAE 5160-class leaf-spring steel (approximately HRC 38–45). This band is a typical reported range rather than a limit prescribed by JIS G4801, which specifies composition and manufacture requirements but no hardness value; it is consistent with published measurements on quenched-and-tempered SUP9/9A-class spring steel (mean HRC 39; individual readings HRC 37–41) [22] and with the observed values, supporting the conclusion that material degradation or a heat-treatment defect was not the cause of premature failure.

2.5. P-F Interval Analysis

The potential-failure (P-F) interval concept from RCM theory [10] was used to quantify the time between

when a failure condition becomes detectable and when it becomes a functional failure. For the spring leaf case, the theoretical P-F interval was estimated by analyzing the time from initial fatigue crack initiation (detectable via vibration monitoring) to final fracture (functional failure). This analysis provides the theoretical basis for both predictive and preventive intervention strategies, as formalized in Equation 5. The P-F interval is presented as a conceptual estimate based on published crack growth rates for SUP9A spring steel; experimental validation under the specific operating conditions is recommended as future work.

2.6. Statistical Analysis

Descriptive statistical analysis was performed on fleet maintenance data to identify failure patterns and correlations. The key metric analyzed was the spring-to-shock replacement ratio, defined as the number of spring leaf replacements divided by the number of shock absorber replacements on the same unit over the same time period. Units with high spring failure rates (ratio $> 10:1$) were compared against units with normal failure rates (ratio 2–3:1) to identify the systemic factor driving the disparity. Correlation analysis between shock absorber replacement frequency and spring failure frequency was performed using Pearson's correlation coefficient ($r = -0.87$, $p < 0.01$, $n = 35$ units, 95% CI $[-0.93, -0.75]$), indicating a strong negative correlation: units with more frequent shock absorber replacements experienced fewer spring failures.

3. Results and Discussion

3.1. Review of Predictive Maintenance Approaches

The majority of PdM research falls under data-driven approaches, where models learn patterns directly from operational data without explicit physical modeling. Early studies employed traditional ML algorithms for failure classification and anomaly detection. Falkovén [2] investigated association rule mining, classification trees, and autoregressive modeling on Volvo truck data in Sweden, concluding that overall data quality was insufficient for reliable PdM implementation. This highlighted a recurring theme: data quality is the foundation of any data-driven approach. Data-driven detection has also been applied to specific degradation mechanisms; for example, Effendi et al. [8] used a convolutional neural network for automated corrosion detection on steel structures, illustrating a mechanism-specific detection capability.

Tree-based ensemble methods have been widely adopted. Biteus and Lindgren [3] trained a random forest model on Scania operational data to estimate vehicle breakdown probability, integrating the output with a constraint-based maintenance planner and route optimizer. Barbosa [11] compared XGBoost, LightGBM, and Random Forest for anomaly detection in Scania heavy-duty trucks, finding that tree-based models showed competitive performance but were ultimately outperformed by LSTM in cost-sensitive evaluation.

Table 1. Summary of predictive maintenance approaches.

Approach	Key Capability	Limitation
Classical ML (RF, XGB, LGBM)	Classification, feature importance	Limited temporal modeling
Deep Learning (LSTM, AE)	Temporal pattern capture	Data-hungry, black-box
Ensemble (SSED)	High accuracy, dimensionality reduction	Complex training
Digital Twins	Comprehensive simulation	Complex implementation
Bayesian	Uncertainty quantification	Prior specification required
Expert Systems	Domain knowledge integration	Not adaptive
IoT Frameworks	Real-time data collection	Infrastructure-dependent

Table 2. Critical gaps in predictive maintenance research.

Gap	Description	Consequence
Prediction–prevention	Models predict failure but not its cause	Same failures recur
AD over RCA	Most research stops at anomaly detection	No actionable diagnostics
Generalization	Models fail on new fleets/conditions	Fleet-specific retraining needed
Data quality	Insufficient for reliable prediction	High false positive/negative rates
Recurrence blind spot	No metric for failure rate reduction	Models optimized for wrong objective

Deep learning methods [12] have gained prominence for capturing complex, non-linear relationships. Khoshkangini et al. [1] proposed a Snapshot-Stacked Ensemble Deep Neural Network (SSED) for vehicle breakdown prediction, combining genetic algorithm-based feature selection with cyclic cosine annealing. Barbosa [11] evaluated LSTM, Autoencoder, and LSTM-AE models for anomaly detection on the Scania APS dataset, achieving the best cost-performance trade-off (prediction cost of 3.44, F2-score of 0.64 on training data), though generalization remained challenging with F2-scores dropping to approximately 0.20 on validation and test sets.

Beyond purely data-driven methods [13], digital twin approaches [14] create virtual replicas of physical assets that integrate real-time data with simulation models. Tang et al. [14] developed a vehicle deterioration digital twin framework for commercial trucks and trailers, using K-means clustering and Latent Dirichlet Allocation to identify failure modes from millions of inspection records. Digital twins offer advantages over purely data-driven methods: they can simulate what-if scenarios, incorporate physical models, and provide continuous lifecycle information. However, they are complex to implement and still depend on the quality of underlying failure mode identification.

Bayesian methods [15] handle the censored nature of workshop data (components replaced before failure) better than deterministic approaches but require careful prior specification and are computationally intensive for high-dimensional problems. Expert systems [16] bridge the gap between data-driven methods and domain expertise but are limited to known failure patterns and cannot discover novel mechanisms. IoT-based fleet management

systems [17, 18] provide the enabling infrastructure for all data-driven approaches, but without reliable data collection, no predictive model can function. Table 1 summarizes the key approaches, their capabilities, and their limitations.

3.2. Critical Gaps in the Predictive Maintenance Paradigm

The most significant limitation of current PdM research is the disconnect between prediction and prevention. A model that accurately predicts “this spring will fail in 500 operating hours” provides valuable information, but it does not answer the question “why does this spring keep failing every 2,000 hours when its design life is 50,000 hours?” Without understanding the root cause, the maintenance team replaces the spring, the same underlying condition persists, and the spring fails again at the same shortened interval. Prediction without prevention is just better-informed reaction.

A critical review of the analyzed literature reveals a disproportionate focus on anomaly detection (AD) at the expense of root cause analysis (RCA); the review is narrative rather than systematic, as documented in Section 2.1. Barbosa [11] explicitly identified this gap, noting that most studies stop at “detecting abnormal behavior without providing information about its causes.” Of the 14 studies examined in this review, only two [11, 14] include any form of root cause analysis, and neither provides a complete physical mechanism framework.

Multiple studies report significant degradation in model performance when applied to new data. Barbosa [11] documented F2-scores dropping from 0.64 on training data to approximately 0.20 on validation and test sets. This generalization failure stems from fleet heteroge-

Table 3. Failure mechanism taxonomy for truck components.

Mechanism	Description	Indicators	Affected Components
Fatigue	Cyclic stress below yield causes crack growth	Beach marks, striations	Springs, shafts, bearings, welds
Overload	Single stress exceeds ultimate strength	Ductile/brittle fracture	Structural parts, bolts
Wear	Progressive material removal	Dimensional change, scoring	Brakes, bearings, seals
Corrosion	Chemical/electrochemical degradation	Rust, pitting, material loss	Exposed metal, exhaust
Creep	Time-dependent deformation under stress	Dimensional change	Engine, exhaust
Fretting	Micro-motion wear under vibration	Surface pitting, debris	Bolted joints, splines
Thermal cycling	Repeated expansion/contraction	Cracking at interfaces	Exhaust, electrical

Table 4. Prevention strategies by failure mechanism.

Mechanism	Strategy 1	Strategy 2	Strategy 3
Fatigue	Reduce stress amplitude	Increase fatigue strength	Eliminate stress raisers
Overload	Add overload protection	Increase design margin	Operator training
Wear	Improve lubrication	Harder surface coating	Filtration
Corrosion	Protective coating	Material selection	Environmental control
Creep	Lower temperature	Creep-resistant alloy	Reduce sustained stress
Fretting	Increase clamping force	Lubricate surfaces	Reduce vibration

neity (different vehicle makes, models, and operating environments), temporal drift (component behavior changes over time), and data distribution shift (training data may not represent future failure modes). A physics-of-failure approach is inherently more transferable: the fatigue mechanism of a spring leaf is the same regardless of which truck it is installed on, and the prevention strategy (reduce stress amplitude) applies universally.

Perhaps the most telling gap is the absence of recurrence metrics in PdM evaluation. Studies report accuracy, precision, recall, F1-score, and cost metrics, but none measure whether the same component fails again on the same equipment after a “successful” prediction and intervention. A model that achieves 95% accuracy but does not reduce the long-term failure rate is, in a fundamental sense, failing at its primary objective. Table 2 summarizes the critical gaps identified in the current PdM paradigm.

3.3. The Physics-of-Failure Alternative

Physics-of-failure (PoF) is an engineering approach that uses knowledge of failure mechanisms and stress analysis to predict and prevent component failure. Unlike data-driven methods that learn correlations from historical data, PoF models the actual physical processes that cause degradation and failure. The fundamental premise is that every failure has a physical cause, and understanding this cause is the most reliable path to prevention. The approach is

grounded in materials science, solid mechanics, thermodynamics, and tribology.

All component failures can be classified into a finite set of physical mechanisms, as shown in Table 3. Each mechanism has specific, well-understood prevention strategies, as summarized in Table 4.

A complete failure analysis traces the chain from latent root cause to functional failure, as illustrated in Figure 1. Each level of this chain represents an intervention point. Data-driven PdM typically only detects the final stages (failure mode or functional failure). PoF analysis identifies intervention points at every level, with the most effective interventions being those closest to the latent root cause.

The following equations provide the theoretical foundation for the physics-of-failure approach applied in this study.

$$da/dN = C(\Delta K)^m \quad (1)$$

where da/dN is the crack growth rate per loading cycle (mm/cycle), C is the Paris constant (material-dependent; with da/dN in mm/cycle and ΔK in $\text{MPa}\sqrt{\text{m}}$, C typically ranges from 10^{-11} to 10^{-9} in the corresponding units of $(\text{mm/cycle})/(\text{MPa}\sqrt{\text{m}})^m$ for spring steels), ΔK is the stress intensity factor range ($\text{MPa}\sqrt{\text{m}}$), and m is the Paris exponent (typically 2–4 for spring steels).

$$\zeta = c/(2m\omega_n) \quad (2)$$




Figure 1. Failure mechanism chain from latent root cause to functional failure.

where ζ is the damping ratio (dimensionless), c is the damping coefficient (N·s/m), m is the sprung mass (kg), and ω_n is the natural frequency of the spring-mass system (rad/s). A shock absorber with $\zeta = 0.30$ at nominal condition provides effective damping [23]; when degraded to $\zeta < 0.05$, the system approaches resonance and spring stress amplifies dramatically.

$$K_a = \sigma_d / \sigma_0 = 1 / (2\zeta) \quad (3)$$

where K_a is the dynamic stress amplification factor (dimensionless), σ_d is the dynamic stress amplitude (MPa), and σ_0 is the static stress (MPa). For $\zeta = 0.05$ (degraded shock absorber), $K_a = 10$, meaning the spring experiences ten times the static stress amplitude from road impacts. Equation 3 is derived, however, for a single-degree-of-freedom system at resonance ($\omega = \omega_n$); under random road excitation the actual amplification depends on the frequency ratio (ω/ω_n), the road input spectrum, and the multi-degree-of-freedom suspension dynamics, and can fall below or above the resonant value. $K_a = 10$ should therefore be read as an idealized resonance-based reference estimate rather than a directly measured or universally applicable upper bound for the actual vehicle suspension; the typical operating increase is estimated at 3–5× from the observed spring fracture rates.

$$N_f = A \cdot (\sigma_a)^{-b} \quad (4)$$

where N_f is the number of cycles to failure, σ_a is the applied stress amplitude (MPa), A is the fatigue strength coefficient (in units of cycles·MPa^{*b*}, so that N_f is in cycles when σ_a is in MPa), and b is the Basquin exponent in this formulation (typically 3–6 for spring steels; the conventional Basquin form $\sigma_a = \sigma_{f'}(2N_f)^b$ corresponds to $|b| \approx 0.17$ – 0.33). For SUP9A spring steel, $A \approx 4.1 \times 10^{13}$ and $b \approx 3.6$ are illustrative Basquin coefficients: A is calibrated to reproduce the fleet-observed normal-condition life ($N_f \approx 5 \times 10^4$ cycles at $\sigma_a \approx 300$ MPa), and $b = 3.6$ lies within the 3–6 range typical of published S–N data for quenched-and-tempered spring steels such as SUP-9 [22]. Under normal conditions ($\sigma_a \approx 300$ MPa), $N_f \approx 5 \times 10^4$ cycles. Under degraded conditions with 3–5× stress amplification ($\sigma_a \approx 900$ – 1500 MPa), N_f falls to roughly 150–1,000 cycles, i.e. to less than 2% of the normal life, a reduction of more than 98%. At the upper bound (5×

amplification, $\sigma_a \approx 1,500$ MPa), the stress approaches the ultimate tensile strength of SUP9A ($\sim 1,300$ – $1,500$ MPa — values typical of quenched-and-tempered spring steels; see [22] for measured hardness and endurance-limit data for SUP-9), and the failure mode transitions from fatigue to quasi-static overload; the fatigue-life equations are therefore approximate in this regime. The stress ratio $R = \sigma_{min}/\sigma_{max}$ was approximately 0.1–0.2 for the spring under normal loading (referring to the oscillating stress component after subtracting the static preload; the absolute R ratio including preload would be negative); no mean stress correction (e.g., Goodman or Gerber) was applied, as the analysis focuses on comparative life reduction rather than absolute life prediction. The fatigue limit of SUP9A spring steel, however, is typically 500–600 MPa (see [22] for measured endurance-limit and hardness data of SUP-9), and the 300 MPa normal stress amplitude lies below it; the observed 50,000-cycle life therefore reflects the specific operating conditions — occasional overload events and surface defects from mining debris — rather than the material's intrinsic high-cycle fatigue behavior. The comparative life-reduction analysis remains valid regardless of the absolute life value. The fatigue-life calculations in this section are therefore illustrative estimates, intended to quantify the contrast between normal and degraded stress states rather than to serve as experimentally validated life predictions for this component; a design-grade analysis would additionally require material-specific S–N data, mean-stress correction (e.g., Goodman), and measured load spectra.

$$PFI = t_F - t_P \quad (5)$$

where PFI is the P–F interval (hours), t_F is the time of functional failure (hours), and t_P is the time when the potential failure condition becomes detectable (hours). For the spring leaf case, the interval is estimated from three explicit assumptions. (i) Initial detectable crack size $a_0 \approx 1$ mm, the smallest crack assumed reliably detectable by the fleet's vibration-based monitoring and visual inspection. (ii) Critical crack size $a_c \approx 4$ mm, identified from beach marks on failed springs, with an uncertainty of roughly 3–5 mm. (iii) Average crack growth rate $da/dt \approx 0.001$ mm/hour, derived from the Paris law (Equation 1) with $C \approx 1 \times 10^{-10}$ mm/cycle and $m \approx 3$ (mid-range values typical of spring steels), and $\Delta K \approx 15$ MPa√m — consis-

Table 5. Data-driven PdM vs. physics-of-failure approach.

Aspect	Data-Driven PdM	Physics-of-Failure
What it predicts	When failure will occur	Why failure occurs
Data required	Large volumes of sensor data	Failure context, material properties
Generalization	Poor (fleet-specific)	Excellent (mechanisms are universal)
Prevents recurrence	No	Yes (by design)
Interpretability	Low (black-box models)	High (physical causality)
Implementation cost	High	Moderate
Maturity	Rapidly evolving	Well-established
Best for	Detecting degradation onset	Eliminating root causes




Figure 2. P-F curve showing the spring failure cycle: shock absorber wear phase, stress transfer and crack initiation, and fatigue crack propagation. The time axis is qualitative (dimensionless, not to scale); the numerical P-F interval corresponding to t_P and t_F is estimated in Section 3.3 (Equation 5).

tent with $\Delta K = Y\Delta\sigma\sqrt{\pi a}$ for $Y \approx 1.1$ (edge-crack geometry), a degraded dynamic stress range $\Delta\sigma \approx 150\text{--}250$ MPa, and $a = 1\text{--}4$ mm — combined with an effective road-induced loading frequency $f \approx 1$ Hz, giving $da/dN \approx 3.4 \times 10^{-7}$ mm/cycle and $da/dt = f \cdot da/dN \approx 0.001$ mm/hour. With these values, $PFI = (a_c - a_0)/(da/dt) \approx 3,000\text{--}4,000$ hours, i.e. approximately 4,000 hours at the upper end; propagating the parameter uncertainties ($a_0 \approx 0.5\text{--}1.5$ mm, $a_c \approx 3.5\text{--}5$ mm, $da/dt \pm 20\%$) widens this to roughly 1,700–5,600 hours, consistent with the 2,000–5,000-hour order-of-magnitude range used throughout this paper. The Paris law predicts accelerating crack growth (da/dN increases with crack length), so the constant-rate estimate is an order-of-magnitude approximation and should be treated as an upper bound, to be

validated with experimental crack growth data under the specific operating conditions. Figure 2 illustrates the qualitative progression from shock wear to crack initiation to final fracture on a dimensionless (qualitative) time axis; the time axis is not to scale.

The comparison between data-driven PdM and physics-of-failure approaches is summarized in Table 5. We emphasize that PoF and data-driven PdM are complementary, not competing. The ideal approach uses: (1) data-driven monitoring to detect when a component begins to deviate from normal behavior; (2) physics-of-failure analysis to determine why the deviation is occurring; (3) root cause elimination to prevent recurrence; and (4) design feedback to improve future component reliability.




Figure 3. Spring leaf versus shock absorber replacement patterns across six fleet units: (a) absolute counts, (b) replacement ratio.

3.4. Case Study: Recurring Spring Leaf Failures in a Heavy Truck Fleet

A fleet of 35 Volvo FMX 440/460 heavy trucks operating in Indonesian mining and logistics applications experienced 7,300 spring leaf failures over a four-year period (2022–2025); the fleet description and data-collection protocol are given in Section 2.2. One unit alone (DT FMX 440-2021) recorded 342 spring failures over four years, corresponding to approximately 1.64 failures per week. Each failure resulted in 2–4 hours of downtime and significant indirect costs from production delays.

The standard maintenance response was reactive: replace the broken spring leaf and return the vehicle to service. No systematic investigation was conducted to determine why springs were failing at such high rates. The same unit, the same position (F3), and the same failure mode (broken spring leaf) recurred with predictable regularity.

Analysis of maintenance records revealed a critical pattern: the spring-to-shock replacement ratio on the worst-affected unit was 28.5:1 (342 spring replacements vs. 12 shock absorber replacements over the same period). For comparison, units with lower spring failure rates showed ratios of 2–3:1, suggesting that shock absorbers were being replaced more frequently relative to springs on healthier units. Pearson correlation analysis confirmed this relationship: $r = -0.87$ ($p < 0.01$, $n = 35$, 95% CI $[-0.93, -0.75]$), indicating a strong negative correlation between shock absorber replacement frequency and spring failure rate. Figure 3 illustrates this pattern across six representative fleet units. This unit-level correlation is observational and may be influenced by operational

confounders such as driver behavior, severe overloading beyond normal damping limits, and variable road-surface quality; these factors were not instrumented continuously over the four-year period and therefore could not be statistically controlled. Three features nevertheless strengthen the causal interpretation without proving it: the effect is unit-level and persistent rather than fleet-wide; the failure position (F3) and failure mode are constant on the worst-affected units; and the mechanism is independently corroborated by measured shock-absorber degradation (piston-rod wear 0.15–0.30 mm; 40–60% oil loss) and by the fatigue stress-amplification analysis (Equations 3–4). We therefore present the correlation as strong supporting evidence within a convergent-evidence argument, not as proof of causation.

The failure mechanism chain was reconstructed through systematic RCFA, as illustrated in Figure 1 and the P–F curve in Figure 2. The physical mechanism proceeds through six steps: (1) Normal wear of the shock absorber seal causes progressive hydraulic oil loss over approximately 20,000 km of operation. (2) At approximately 50% oil loss, the shock absorber's damping coefficient drops to less than 20% (Equation 2) of its design value. (3) Without functional damping, the spring leaf must absorb the full road impact energy, increasing dynamic stress by a factor of 3–5 (Equation 3). (4) The spring, designed for a fatigue life of approximately 50,000 cycles under normal conditions, now experiences stress exceeding its fatigue limit, initiating a crack at a surface defect or stress raiser (Equation 4). (5) The crack propagates under continued cyclic loading (Equation 1) until final fracture occurs. (6) The spring breaks, the vehicle is taken out of service, the




Figure 4. Comparison of prediction-only vs. physics-of-failure maintenance approaches.

spring is replaced—but the worn shock absorber remains, and the cycle repeats.

The quantitative analysis confirms the mechanism. Using the stress amplification factor from Equation 3, with a fully degraded shock absorber ($\zeta = 0.05$), the idealized resonance-based reference estimate is $K_a = 10$. Under typical operating conditions, the actual stress increase is estimated at $3\text{--}5\times$ based on the observed reduction in spring fatigue life. From Equation 4, a $3\text{--}5\times$ stress amplification reduces the fatigue life from approximately 50,000 cycles to 1,000 cycles or fewer—a reduction of more than 98%. The theoretical P-F interval from Equation 5 provides an estimated intervention window of approximately 4,000 hours, consistent with the 2,000–5,000-hour order-of-magnitude range derived in Section 3.3; Figure 2 illustrates the interval qualitatively.

A predictive model trained on this fleet's data would likely achieve reasonable accuracy in predicting individual spring failures. However, because the model does not address the root cause (degraded shock absorber), the spring will continue to fail at the same elevated rate. In contrast, the physics-of-failure analysis identifies a simple, effective prevention strategy: replace the shock absorber at its recommended interval (or when damping drops below a threshold), thereby restoring the spring to normal stress conditions and eliminating the fatigue driver. Figure 4 illustrates the comparison between prediction-only and physics-of-failure maintenance approaches.

The prevention strategy derived from the PoF analysis has three components. Immediate action: implement shock absorber inspection and replacement at 20,000 km intervals or when damping drops below $\zeta = 0.15$. Short-term action: develop and implement a paired replacement SOP requiring shock absorber inspection whenever a spring leaf is replaced. Long-term action: redesign the suspension system to include damping monitoring and integrate this data into the fleet's condition monitoring system. The 20,000 km interval is an author-derived recommendation based on the fleet's observed wear data (Section 2.4) and supported by OEM and industry scheduled-inspection

practice [20,21]; it is not a direct manufacturer-prescribed replacement interval.

The 85–95% reduction in spring failure rate attributed to the PoF intervention is a projection, not a measured outcome: the intervention was implemented after the data-collection period, so no before-and-after field data exist for the fleet. The projection follows from the fatigue-life analysis (Equations 3–4). Restoring normal stress levels raises fatigue life from roughly 1,000 cycles back to roughly 50,000 — a reduction of more than 98% at the level of an individual spring subject to the degraded stress state. The translation to a fleet-wide failure-rate reduction is bounded as follows. Within the same fleet and duty cycle, units with maintained shock absorbers (spring-to-shock ratio 2–3:1) failed roughly an order of magnitude less frequently than the worst-affected unit (28.5:1), indicating that shock-absorber condition dominates spring failure risk on this fleet. Eliminating the degraded-shock mechanism should therefore restore the affected units toward the low observed baseline. The projection is nevertheless discounted from the >98% per-unit fatigue-life recovery to 85–95% at fleet level because (i) a residual failure population arises from mechanisms the intervention does not address (overload, road debris, manufacturing defects), and (ii) adherence to the scheduled replacement SOP will be imperfect in operations. Confirming these figures with post-intervention data is a critical area of future work.

3.5. Integrated Framework

We propose a three-layer framework that integrates data-driven monitoring with physics-of-failure analysis and systematic root cause elimination, as illustrated in Figure 5. Layer 1 (Condition Monitoring) uses IoT sensors and anomaly detection algorithms [5, 6] to identify when a component begins to deviate from normal behavior. This layer leverages the strengths of data-driven PdM: continuous monitoring, early detection, and pattern recognition across large fleets.

Layer 2 (Root Cause Analysis) applies RCFA and PoF




Figure 5. Three-layer integrated framework: condition monitoring → root cause analysis → design feedback.

methods to determine why the deviation is occurring. When an anomaly is detected, the RCFA process traces the causal chain from failure mode back to physical, human, and latent root causes. This layer uses the 5-Whys, fishbone diagrams, and fault tree analysis described in Section 2.3.

Layer 3 (Design Feedback) implements corrective actions and feeds lessons learned back into design specifications, maintenance procedures, and training programs. This is where prevention becomes permanent: a design change or procedure update eliminates the root cause for all similar equipment, not just the specific unit that failed.

3.5.1. Novelty and Contribution of the Proposed Framework

We acknowledge that the individual elements of the framework—condition monitoring, root cause analysis, and design feedback—are individually well established in maintenance and reliability engineering, root cause failure analysis (RCFA), and reliability-centered maintenance (RCM). The novelty of this work is therefore not the invention of these elements, but their integration into a single closed loop with an explicit, mechanism-based hand-off between them, and the demonstration of that loop on a large, real failure dataset.

Specifically, the contribution is threefold. First, the framework defines an explicit interface at which data-driven anomaly detection (Layer 1) triggers a physics-of-failure root cause analysis (Layer 2) rather than stopping at detection—a hand-off that is largely absent from the re-

viewed PdM literature (Section 3.2). Second, it closes the loop from root cause to prevention and to design (Layer 3), so that a corrective action is propagated to all similar equipment instead of only the unit that failed. Third, it introduces recurrence as the evaluation criterion: the success of maintenance is measured by whether the same failure recurs, not only by prediction accuracy. Within the literature examined in this study, we did not identify a framework that explicitly combines these three features and demonstrates their application using a multi-thousand-event industrial dataset.

A note on implementation status: the complete three-layer framework has not yet been implemented and evaluated as an operational system. The case-study evidence covers the mechanism-based root cause analysis (Layer 2) and the prevention decision derived from it (timely shock absorber replacement); continuous data-driven condition monitoring (Layer 1) and fleet-wide design feedback (Layer 3) remain proposed elements. Consequently, all quantitative benefits associated with the framework as a whole—including the failure reduction and ROI estimates in Sections 3.4 and 3.6—are projections (Sections 3.4 and 3.7), and evaluating the framework as an integrated loop requires post-intervention, before-and-after field data.

3.6. Economic Analysis

The economics of prevention versus prediction are compared in Table 7. Prediction costs include sensors (\$50–500 per monitoring point), data infrastructure (\$10,000–50,000 per fleet for cloud services and analyti-

Table 6. Comparison of prediction-only vs. physics-of-failure outcomes for the case study.

Metric	Prediction-Only	Physics-of-Failure
Spring failures (4 years)	7,300 (recurring)	~365–1,095 (projected 85–95% reduction)
Cost (4 years)	Rp 6.2–24.5 billion (ongoing)	Rp 0.7–2.0 billion (one-time + scheduled shock replacement)
Root cause addressed	No	Yes
Recurrence	Same failures repeat	Failures eliminated at source
ROI	N/A (cost center)	5:1 to 10.5:1 (projected)

Table 7. Economic comparison of prediction-only vs. physics-of-failure approaches.

Cost Category	Prediction-Only	Physics-of-Failure	Notes
Sensor hardware	\$50–500/point	Minimal (visual inspection)	PoF needs fewer sensors
Data infrastructure	\$10–50K/fleet	Minimal	Cloud + analytics vs. spreadsheets
Model development	\$20–100K	\$7–20K (one-time RCFA + SOP)	ML models need ongoing updates
Maintenance cost (4 yr)	Rp 6.2–24.5B (recurring)	Rp 0.7–2.0B (one-time + scheduled shock replacement)	Spring replacement only
ROI	N/A (cost center)	5:1 to 10.5:1	Projected; conservative-to-central assumptions (Section 3.6)

cs), and model development (\$20,000–100,000 for custom ML models) [19]. For the case study fleet of 35 units with 7,300 spring failures over four years, the total spring replacement cost alone exceeds Rp 6.2 billion (approximately \$387,500 at an exchange rate of Rp 16,000/USD as of June 2026). These costs include parts and labor but exclude downtime costs, which would substantially increase the total. These figures are projections derived from the case-study fleet's failure counts and standard parts and labor rates; they are not audited post-intervention financial results.

Prevention costs include the RCFA investigation (\$5,000–15,000 one-time), shock absorber replacement at the proposed 20,000 km interval (\$50–150 per unit per interval), and SOP development and training (\$2,000–5,000 one-time). The estimated ROI for the case study is 5:1 to 10.5:1 over a four-year period, based on the projected reduction in spring replacement costs; the explicit cost model behind this range is as follows. NASA's RCM program, by comparison, documented a return of 1.75–2.20 over four years, with cost avoidance of \$26.8–33.6 million [9]—notably lower than the case-study estimate, because the exceptionally high failure frequency of mining operations amplifies the savings from root cause elimination. These figures suggest that root cause elimination can be more economically efficient than prediction alone, though the actual ROI depends on fleet size, failure rates, and implementation costs. The stated ROI is therefore a projected value rather than a validated one; empirical

confirmation requires before-and-after cost data (Section 3.7).

Baseline: 7,300 failures over four years at Rp 850,000–3,350,000 per event (parts and labour; Section 2.2) gives Rp 6.2–24.5 billion ($\approx$ US\$0.39–1.53 million at Rp 16,000/USD; mid-range Rp 2.1 million per event). Intervention: one-time RCFA (\$5,000–15,000) and SOP development and training (\$2,000–5,000), plus scheduled shock absorber replacement every 20,000 km — assuming $\approx$ 100,000 km per truck per year (five replacement intervals per truck per year) for the 35-unit fleet at \$50–150 per shock absorber per interval, i.e. \$35,000–105,000 over four years — for a total intervention cost of $\approx$ \$42,000–125,000 (Rp 0.67–2.0 billion) over four years. Savings: avoiding 85–95% of failures prevents 6,205–6,935 events, worth Rp 5.3–23.2 billion (US\$329,000–1,452,000), before any downtime savings (14,600–29,200 vehicle-hours over four years), which are excluded here and therefore make the estimate conservative. The resulting ROI = savings $\div$ intervention cost spans $\approx$ 2.6:1 (low-end savings, high-end cost) to $\approx$ 34.6:1 (high-end savings, low-end cost), with a central estimate of $\approx$ 10:1 (mean savings Rp 13.8 billion, mean cost Rp 1.3 billion). The 5:1–10.5:1 headline range reported in Tables 6–7 corresponds to conservative-to-central paired assumptions (e.g., Rp 5.3 billion savings against Rp 1.05 billion cost gives 5:1; Rp 13.8 billion against Rp 1.3 billion gives $\approx$ 10.5:1) and is retained as the prudential headline estimate.

3.7. Limitations

Several limitations should be acknowledged. The case study data is drawn from a single fleet operating in specific conditions (Indonesian mining), which may limit generalizability. The P-F interval analysis relies on estimated parameters that should be validated with additional experimental data. The integrated framework requires organizational commitment to root cause investigation, which may be challenging in resource-constrained environments. In addition, the narrative review analyzed 14 studies, which, while sufficient to identify persistent gaps, limits the generalizability of conclusions about the broader PdM research landscape. Finally, this review focuses on mechanical failure mechanisms; software failures, human errors in operation, and organizational factors require separate treatment. The 85–95% failure reduction and the 5:1–10.5:1 ROI are projections — derived from fatigue-life calculations and cost assumptions, respectively — not measured outcomes. The single most important item of future work is to collect post-intervention, before-and-after field data and report the observed change in failure rate, so that these projections can be confirmed, revised, or refuted.

4. Conclusions

This review has identified a fundamental gap in the predictive maintenance paradigm: most PdM approaches address the symptom of failure (when it will occur) rather than its cause (why it occurs). Through a critical narrative review of PdM research (2010–2025) and a case study of 7,300 recurring spring leaf failures in a heavy truck fleet, we have demonstrated that:

1. Current PdM approaches achieve prediction accuracy but do not reduce failure recurrence rates. A model that correctly predicts each failure event is still, in a fundamental sense, failing if the same component fails again under the same conditions.
2. Physics-of-failure analysis, grounded in root cause failure analysis and reliability-centered maintenance, provides a complementary and more durable solution by identifying and eliminating the physical, human, and latent root causes of failure.
3. The case study demonstrates that a single overlooked failure mechanism—shock absorber wear causing increased dynamic stress on springs—perpetuated a cycle of 7,300 failures over four years that no predictive model could break. Addressing the root cause (timely shock absorber replacement) is projected to reduce the spring failure rate by an estimated 85–95%, based on fatigue life analysis. Validation with post-intervention field data is recommended as future work.
4. An integrated three-layer framework—condition monitoring, root cause analysis, and design feedback—combines the strengths of both approaches: early detection from data-driven methods and permanent elimination from physics-of-failure analysis. The framework it-

self is proposed rather than field-evaluated end-to-end: the case study demonstrates its mechanism-analysis core (root cause analysis and the resulting prevention logic), while the integrated monitoring-analysis-design loop requires validation with post-intervention field data (Section 3.7).

The path forward is not to abandon data-driven methods, but to integrate them with mechanism-based analysis in a three-layer framework: condition monitoring → root cause analysis → design feedback. This integrated approach addresses both the immediate need (detect and respond to degradation) and the long-term goal (eliminate the causes of failure). The ultimate objective of maintenance should not be to predict failure, but to make failure prediction unnecessary by preventing its causes.

Acknowledgments

The authors gratefully acknowledge the fleet management team at PT Bukit Asam for providing maintenance data and access to fleet records. We also thank the Institut Teknologi Sepuluh Nopember (ITS) and Institut Teknologi Kalimantan (ITK) for institutional support.

Data Availability

The fleet maintenance data used in this case study is proprietary and confidential. Aggregated data and analysis results are presented in the manuscript. Interested researchers may contact the corresponding author for further information, subject to confidentiality agreements. To support reproducibility while protecting proprietary information, Supplementary File S1 (S1_synthetic_sample.csv) provides a synthetic, anonymized dataset of 35 fleet-unit records (four-year spring and shock absorber replacement counts per unit, with unit-level wear indicators), matching the $n = 35$ units of the correlation analysis, and Supplementary File S2 (S2_analysis.py) provides the Python script that reproduces the unit-level Pearson correlation (with 95% confidence interval), the spring-to-shock ratio summary, and the fatigue-life comparison of Equation 4. Both files are uploaded with this submission and are consistent with this statement; Supplementary File S2 operates on Supplementary File S1, so that the statistical and fatigue analyses can be reproduced without disclosing confidential fleet data.

Conflict of Interest

The authors declare no conflict of interest.

Author Contributions

Witantyo: Conceptualization, methodology, investigation, writing – review and editing, supervision. Taufik Hidayat: Writing – original draft, formal analysis, visualization, validation. Ahmad Nur Muzaqin: Data curation, investigation.

References

- [1] R. Khoshkangini, M. Tajgardan, J. Lundström, M. Rabbani, and D. Tegnered, "A snapshot-stacked ensemble and optimization approach for vehicle breakdown prediction," *Sensors*, vol. 23, p. 5621, 2023.
- [2] A. Falkovén, "On data-driven predictive maintenance of heavy vehicles: A case study on Swedish trucks," M.S. thesis, Chalmers University of Technology, Gothenburg, Sweden, 2017.
- [3] J. Biteus and T. Lindgren, "Planning flexible maintenance for heavy trucks using machine learning models, constraint programming, and route optimization," *SAE Int. J. Mater. Manuf.*, vol. 10, no. 3, pp. 306–315, 2017.
- [4] R. K. Mobley, *An Introduction to Predictive Maintenance*. Burlington, MA: Butterworth-Heinemann, 2nd ed., 2002.
- [5] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Computing Surveys*, vol. 41, no. 3, pp. 1–58, 2009.
- [6] O. Serradilla *et al.*, "Deep learning models for predictive maintenance: A survey, comparison, challenges and prospects," *Applied Intelligence*, vol. 52, pp. 10934–10964, 2022.
- [7] S. Gawde, S. Patil, S. Kumar, P. Kamat, K. Kotecha, and S. Alfarhood, "Explainable predictive maintenance of rotating machines using LIME, SHAP, PDP, ICE," *IEEE Access*, vol. 12, pp. 29345–29361, 2024.
- [8] M. K. Effendi, B. Atmaja, A. Wahjudi, and D. B. Purwanto, "Automated corrosion detection on steel structures using convolutional neural network," *JMES: The Int. J. Mechanical Engineering and Sciences*, vol. 7, no. 1, pp. 36–44, 2023.
- [9] NASA, "Reliability-centered maintenance guide for facilities and collateral equipment." National Aeronautics and Space Administration, Washington, D.C., NASA/TM-2008-215280, 2008.
- [10] F. S. Nowlan and H. F. Heap, "Reliability-centered maintenance." Department of Defense, Washington, D.C., Report AD-A066579, 1978.
- [11] I. A. Barbosa, "Root cause analysis of automotive system failures: A predictive maintenance approach," M.S. dissertation, University of Porto, Porto, Portugal, 2025.
- [12] F. Nurdin, M. K. Effendi, and D. Mohakul, "Machine learning-based prediction of diesel engine health using operational parameters: Comparison of SVM and BPNN models," *JMES: The Int. J. Mechanical Engineering and Sciences*, vol. 9, no. 2, pp. 69–80, 2025.
- [13] P. Bethaz *et al.*, "Empowering commercial vehicles through data-driven methodologies," *Electronics*, vol. 10, p. 2381, 2021.
- [14] P. Tang *et al.*, "Deterioration digital twins of commercial trucks and trailers for targeted inspection and maintenance." USDOT University Transportation Centers Program, 2022.
- [15] S. Park and J. Lee, "Bayesian-based component lifetime prediction model using workshop and telematics data," *Proc. PHM Society Asia-Pacific Conf.*, vol. 4, no. 1, 2023.
- [16] T. Lindgren and J. Biteus, "Expert guided adaptive maintenance," *Proc. European Conf. Prognostics and Health Management Society*, vol. 2, no. 1, 2014.
- [17] K. Lee *et al.*, "A study on IoT-based fleet maintenance management," *Int. J. Control and Automation*, vol. 10, no. 4, pp. 287–296, 2017.
- [18] A. Giannoulidis *et al.*, "Predictive maintenance in a fleet management system: The Navarchos case," in *Lecture Notes in Business Information Processing*, vol. 520, Springer, 2024.
- [19] S. M. Bayzid, Y. Mohamed, and M. Al-Hussein, "Prediction of maintenance cost for road construction equipment: A case study," *Canadian J. Civil Engineering*, vol. 43, no. 5, pp. 480–492, 2016.
- [20] Volvo Trucks North America, "Service Manual Trucks, Group 177–501: Preventive Maintenance, Annual Service VN, VHD." Publication PV776-TSP158779, Volvo Trucks North America, Inc., Greensboro, NC, USA, 2002. (OEM scheduled front- and rear-shock-absorber checks in the annual service.).
- [21] Hendrickson Asia Pacific, "Shock Absorber Inspection." Technical Bulletin 97117-208, Rev. E, Hendrickson Asia Pacific Pty Ltd, Dandenong, Vic., Australia, 2024. (Regular scheduled inspection of heavy-vehicle shock absorbers; misting vs. leaking criteria; damping affected at roughly 10% oil-volume loss.).
- [22] J. Daudpoto, D. Kumar, S. A. Shaikh, T. Hussain, and Z. A. Memon, "Fatigue failure of SUP-9 spring steel," *Materials Research Express*, vol. 8, p. 066509, 2021. (Quenched-and-tempered SUP-9 spring steel: hardness mean HRC 39, individual readings HRC 37–41.).
- [23] T. Kurniawan and H. L. Guntur, "Comparative study of hydro-magneto-electric regenerative shock absorber (HMERSA) with two outputs hydraulic generator installed series and parallels," *JMES: The Int. J. Mechanical Engineering and Sciences*, vol. 5, no. 2, pp. 45–54, 2021.