

Adaptive Digital Twin Framework for Real-Time Asset Monitoring and Predictive Maintenance

Satyam Pandey

National Institute of Technology, Delhi

Email: 232220012@nitdelhi.ac.in

Abstract- This study suggests a novel adaptive digital twin system, a synthesis between real-time IoT telemetry, physics-ML hybrid modeling, and closed-loop control to allow systematic asset health assessment and scalable predictive maintenance. The architecture is multi-layered: data acquisition and edge filtering; a streaming analytics layer with anomaly detection and remaining useful life (RUL) estimation; an adaptive model layer with self-calibration, online learning and parameter estimates; and a decision orchestration layer giving prescriptive maintenance instructions and scheduling updates back to the physical asset. Bidirectional data flows between models and automated model management is used to reduce model-plant mismatch in nonstationary environments, i.e., wear, drift, and changing operating regimes. We make recalibration triggers of model recalibration, between physics residuals and map embeddings define health, and uncertainty-aware schedules improve timing in maintenance and spare parts logistics and production constraints. The architecture achieves heterogeneous assets and old systems through common interfaces, hence resolving interoperability gaps found in recent survey of experts. The representative industrial machinery assessment shows a better early-fault detectability, less falseness and longer maintenance times than when using the static models, as well as not compromising explainability via physics-constrained predictors. The findings point to a road to stable, self-improving digital twins that bring quantifiable uptime, cost and safety ascent in real time functioning.

Keywords—Asset monitoring, predictive maintenance, adaptive digital twin, anomaly detection, interoperability, predictive maintenance, remaining useful life, Self-calibrating models, virtual sensing, uncertainty quantification.

I. INTRODUCTION

The industrial assets in the past have been based on reactive and time-based preventive maintenance, which cause unplanned unavailability, high levels of empty inventory and the underutilization of both labor and capital. Predictive maintenance aims to avert failures prior to failure, standard methods, however, have difficulty with imperfect sensing, model-plant imprecision and nonstationary operating conditions that deteriorate model usefulness over time [1]. Digital twins offer an identicalized virtual representation of real-world assets to constantly check their condition, put anomalies in context and simulate interventions, which occupied fundamental flaws of antiquated maintenance paradigms. With real-time data, physics-based modeling, and analytics, digital twins allow condition-based decision-making that more closely matches maintenance of actual asset health, minimizing unnecessary maintenance and disastrous failure. But multiplexing reliable predictions between different regimes must be implemented through adaptability mechanisms re-calibrating models and uncertainty Peter a universally applicable framework Adaptive digital twins adopt reliable, scalable predictive health of complex evolving systems depends on wear, drift, and

configuration dynamics re-calibrated online--heralds need an adaptive digital twin [2].




Fig 1 Adaptive DT Framework

A. Concept and Scope of Digital Twins:

A digital twin is a conceptual entity where a physical process or asset is connected to the virtual counterpart via the process of synchronizing data in addition to providing information to undergo analysis, prediction, and decision reinforcement throughout the asset lifecycle. Digital twins extend, between design and operations, using subtypes: Digital Twin Prototype (DTP), Digital Twin Instance (DTI), and Digital Twin Aggregate

(DTA) digital twins to meet different lifecycle and fleet-level learning requirements. To anticipate requirements, predictive maintenance Digital twins combine sensor streams, past maintenance history, physics/data models to predict health indicators, ex-post failures, and in silico experiment with what-if maintenance plans [3]. They are not limited to monitoring at component-level but can be used to understand a system-of-systems twins that can measure interactions, and fault propagation between assets and processes. Standardized interfaces, model management and governance are required in practical applications to allow fidelity, traceability, and safety between edge and cloud spaces and enable virtual sensing, parameter estimation and continuous parameter estimation of required robust and real-time maintenance decisions [4].

B. Digital Twin-Enabled Predictive Maintenance:

Digital twins represent the data and modeling foundation of predictive maintenance Digital twins. Digital twins are created by integrating real-time IoT telemetry with analytics to identify anomalies, predict remaining useful life (RUL), and give timely directions in Fig 1. They continuously monitor operating condition at the component and system scale and have learned how degradation will vary over the fleet on how to optimize accuracy in the forecasts, and optimize the alarm false alarms. With virtual environments, the maintenance teams will be able to simulate problems, test failure points, and schedule efficiently with minimal operational risk. The telemetry, models, and simulation help to increase the accuracy of when to perform maintenance and prolong the life of the asset as well as reduce downtime and incursion [5]. Practically, contextual understanding offered by the synchronized condition of the twin is superior to fixed thresholds, dynamic prioritization and resources allocation within preexisting real-world limitations. The outcome is a transition away to condition-controlled, uncertainty-conscious, and intervention as a calendar mechanism enhanced safety and productivity and a layer for self-learning self-preserving maintenance ecosystems. [6].

C. Architectural Building Blocks for Real-Time Monitoring
 A sound digital twin layering of real-time asset monitoring generally consists of a data acquisition and edge filter, streaming analytics, a hybrid modeling layer (physics + ML), and a decision orchestration layer that ties insights to maintenance processes [7]. When the connectivity is throbbing, edge-computing enables low-latency preprocessing, event detection and resilience, whereas cloud services offer scalable storage, historical analytics, and a cross -fleet learn. The sparse or noisy measurements are enhanced by the virtual sensors and the state estimators, enhancing observability and diagnostic resolution. Data models and interoperable API could be standardized to integrate with other enterprise systems like CMMS, PLM and MES to be run in a similar

fashion and to create consistency and traceability. The security and governance is present on multiple layers to ensure data integrity and safe actuation in the case of closed-loop control. This modularity enables the deployment across heterogeneous assets as well as enables incremental adoption, according to operationally critical and constraint factors [8].

D. Adaptivity: From Static Models to Self-Calibrating Twins

Nonstationary operating conditions, such as wear, environmental variation, reconfiguration or process drift, impair the accuracy of a fixed model and they require adaptive twins that adapt online and co-evolve with their physical incarnations. Adaptive digital twins are mechanisms that implement parameter estimators, the updates to model structure and optionally concept drift, to minimize model-plant discrepancy over time. The data breadth and speed is represented in ontology-driven frameworks, which sustain the linear development of model meaning and interface as systems evolve. Digital Sub-Twins and Digital Shadows are capable of modularizing adaptation through subsystems, updating them, and allowing more focused recalibration, whilst ensuring overall system coherence on a system level [9]. Recalibration triggers may be motivated by trend remnants, uncertainty cutoffs or regime changes (they can be guardrailed to avoid destabilizing updates). With this adaptivity comes sustainable predictive functionality that assures correct diagnostics and prognostics even under changing conditions with partial observability of activities in the real world. [10].

E. Hybrid Physics–ML Modeling and Virtual Sensing

Physics-based models combined with machine learning provide robust predictors that are interpretable and making use of domain model together with data driven trends to inform high fault detection and RUL estimation [11]. The physics view can offer structure, safety limits, and extrapolative power in invisible conditions whereas ML can learn more complex nonlinearities, interactions, and hints of degradation concealed through high-dimensional sensor data. Virtual sensors and observers contribute diagnostic granularity by inferring distant states and loads that are not directly measured and help to minimize dense physical instrumentation. Residual learning Hybrid methods, such as physics-informed learning (ML) and residual learning, improve the interpretability and training quality, especially when labeled loss data about failure is limited[12]. Adaptive parameter estimation will maintain asset-conditional models, whereas uncertainty quantification is used to guide risk-taking and timing of maintenance. This synergy supports scalable, asset-general scale able to specialise with instance-level per-asset calibration which achieves reliable predictive maintenance.

F. Data Infrastructure, Interoperability, and Governance

Predictive twins require data fabric which may enable high velocity streams, contextual, and secure communications across organizational and supplier limits. Interoperability is determined by standardized interfaces and schemas that allow integrating sensors, control systems, engineering logs, and engineering artifacts into a common asset knowledge graph [13]. Lifecycle governance encompasses data quality controls, versioning within models, tracking of lineages and auditability to generate trust and regulatory compliance. patterns of deployment combine edge processing responsiveness with cross-fleet learning cloud analytics, and necessitate coordinated deployment that takes into account bandwidth, latency, and privacy limitations. Role based access and encryption protects operational and IP sensitive data, and safety cases establish limits to automated recommendations or actuation. Mature governance (MoM) speeds up the pilot through enterprise scaling through repeatability, maintainability, and the rise of the better of the twin and such predictive services, and wears on heterogenous and evolving asset portfolios.

G. Decision Support, Optimization, and Maintenance Orchestration

Intelligent digital twins operationalize knowledge with decision support that ranks risks, actions and integrates execution with CMMS, spares and production plans. Optimization models trade off a risk of failure, maintenance windows, inventory and throughput to schedule interventions that minimize lifecycle cost and disruption. The trade-offs modeled in simulation of what-if strategies in the twin analyzes the strategies' trade-offs in the face of uncertainty and thereby develops robust plans that can be adjusted according to the varying conditions and constraints. Dashboards and visual interfaces enhance situational awareness of both the operators, planners and executives and advice and alarms fit into the existing workflows to make adoption easier. Policies are gradually refined over time as self-learning loops using outcomes, and eventually the process becomes autonomous creating closed-loop maintenance where appropriate and safe. Such an orchestration converts predictive insights into quantifiable uptime changes, cost, and safety-related improvements, bridging the analytics and trusted, reproducible field delivery [14].

H. Industrial Use Cases and Cross-Sector Relevance

In manufacturing, energy, aerospace, critical infrastructure, digital twin predictive maintenance is being deployed with reported decreases to unplanned downtime, increased safety, and long service life. Watchdog - Manufacturing twins observe machine tools and production lines to identify abnormalities and streamline maintenance without affecting takt time and integrate with MES to coordinate actions. There are implementations of the energy industry that use twins in

turbines and substations and plants to predict failures and performance optimization in varying loads and environments. Engine and aircraft twins are used aerospace fleet-level learning and personalised maintenance plans relative to real usage and degradation trends. These examples indicate the importance of real-time monitoring, scenario simulation, and data-based scheduling in day and night high-stakes complex operations, which encourages generalized models that are flexible to domain-specific constraints, regulation, and safety engineering.

I. Research Gaps and Contributions of an Adaptive Framework

Nevertheless, in spite of improvements, there are open challenges of managing nonstationary and scarce labelled failures, quantification of uncertainty and scalable governance across heterogeneous assets. The prevailing surveys and architectures indicate that adaptive, modular twins are required to co-evolve through standardized semantics, offer hybrid physics-ML-modelling and provide explainable decisions that are safe under condition changes. This study makes artificial additions such as an adaptive digital twin architecture that formalizes recalibration activations, incorporates virtual sensing and physics-informed machine learning, and entails uncertainty-aware optimization to the maintenance planning. It further promotes interoperability via modular sub-twins and decision pipelines with alignment to enterprise systems and defines deployment patterns with edge-cloud to real time responsiveness and fleet learning. Empirical assessments emphasized superior fault detectability, lowering false alarms and optimization of maintenance timings as critical loopholes of resilience, scalability and operationalization of predictive maintenance under dynamic industrial conditions are considered. [15].

II. LITERATURE REVIEW

Digital twin-enabled predictive maintenance converges real-time data, hybrid physics-ML modeling, and synchronized virtual-physical loops to deliver condition-aware diagnostics, RUL estimation, and prescriptive scheduling under operational constraints, outperforming static thresholding and calendar-based maintenance through closed-loop synchronization and scenario testing. Systematic mappings of this field highlight a methodological pipeline from data ingestion and virtual modeling to health indicator construction and decision orchestration, with clear gains in profitability, safety, and sustainability across manufacturing, energy, and infrastructure when uncertainty is quantified and integrated into planning. Reviews consistently stress interoperability, standardized interfaces, and lifecycle governance to scale beyond pilots, alongside the need for virtual sensing and observers to address sparse or noisy

measurements and improve observability in real time. LIVE methodologies and capability-level roadmaps clarify maturation from descriptive to autonomous twins, emphasizing online adaptation, validation, and the integration of risk-aware optimization for robust maintenance timing and resource allocation. Empirical implementations further show that discrepancy monitoring between physical units and their twins can reliably expose degradation trends, while identifying human-induced confounders that necessitate governance, explainability, and model management in production contexts. Sectoral surveys, notably in wind energy, reinforce the importance of edge-cloud partitioning, hybrid modeling, and uncertainty-aware scheduling to handle heterogeneity and nonstationarity across assets and environments.

Key gaps persist in handling nonstationarity, data sparsity for labeled failures, secure real-time integration, and rigorous validation across lifecycle stages, motivating adaptive twins that self-calibrate via residuals, drift detection, and context-aware parameter updates. Physics-informed strategies improve sample efficiency, stability, and interpretability of prognostics, reducing false alarms and enhancing generalization when embedded within streaming DT architectures that enforce physical feasibility and provide calibrated uncertainties for decision support. LIVE digital twin practices operationalize iterative learning and verification cycles that align sensor placement, model fidelity, and diagnostic robustness with evolving asset behavior, enabling progressive capability uplift toward prescriptive and autonomous maintenance. Industrial case evidence indicates promising detection performance over extended evaluation windows, while also exposing the necessity of standardized data models, cyber-physical security, and auditability to ensure trust and safe actuation. Cross-domain umbrella and sectoral reviews converge on the importance of modular architectures, role-based access, and edge analytics for responsiveness, complemented by cloud-scale fleet learning for transferability across sites and configurations. Collectively, these insights justify an adaptive framework that fuses hybrid physics-MIL, virtual sensing, and uncertainty-aware optimization to sustain reliable prognostics and orchestrate maintenance under changing regimes at scale.

III. PRELIMINARIES

A. STATE-SPACE PROCESS MODEL (DISCRETE-TIME)

$$x_{k+1} = F_k x_k + G_k u_k + w_k \quad (1)$$

x_k : State vector at time step k

F_k : State transition matrix

G_k : Control input matrix

u_k : Control/input vector

w_k : Process noise vector

This fundamental model predicts the future internal state of the asset in the digital twin based on the previous state and control inputs. It enables real-time simulation of asset behaviour under operational and environmental changes, facilitating continuous health tracking and adaptive prediction.

B. Measurement Model

$$z_k = H_k x_k + v_k \quad (2)$$

z_k : Measurement vector from sensors

H_k : Observation matrix

x_k : State vector

v_k : Measurement noise

This equation connects sensor readings from the physical asset to the internal state estimates of the digital twin. It is critical for reconciling predicted and observed behaviours, accounting for noise and incomplete measurements.

C. Kalman Filter Prediction Step

$$\hat{x}_{k|k-1} = F_{k-1} \hat{x}_{k-1|k-1} + G_{k-1} u_{k-1} \quad (3)$$

$$P_{k|k-1} = F_{k-1} P_{k-1|k-1} F_{k-1}^T + Q_{k-1} \quad (4)$$

$\hat{x}_{k|k-1}$: Predicted state estimate

$P_{k|k-1}$: Predicted covariance

Q_{k-1} : Process noise covariance matrix

Used in adaptive twins for predictive monitoring, this step forecasts the next state and its uncertainty before integrating new sensor data — vital for early fault detection.

D. Kalman Filter Update Step

$$K_k = P_{k|k-1} H_k^T (H_k P_{k|k-1} H_k^T + R_k)^{-1} \quad (5)$$

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k (z_k - H_k \hat{x}_{k|k-1}) \quad (6)$$

K_k : Kalman gain

R_k : Measurement noise covariance

$\hat{x}_{k|k}$: Updated state estimate

This update improves the asset's estimated condition by optimally blending prediction with actual sensor measurements, keeping the twin synchronized with reality.

E. Remaining Useful Life (RUL) Estimation

$$RUL = t_f - t_c \quad (7)$$

t_f : Predicted failure time

t_c : Current time

A core predictive maintenance metric, RUL estimation allows the twin to provide actionable insights on when maintenance should be performed before asset failure.

F. Health Indicator (HI) Calculation

$$HI_k = 1 - \frac{\|x_k - x_{ref}\|}{\|x_{fail} - x_{ref}\|} \quad (8)$$

x_{ref} : Healthy state

x_{fail} : Failure state

This normalized metric quantifies degradation, enabling adaptive maintenance decisions within the digital twin framework.

IV. RESULTS AND DISCUSSION

Table 1 – NASA CMAPSS FD001 Dataset Summary.

Metric	Value
Units (train/test)	100/100
Sensors (after dropping constants)	14
Operating conditions	1
Max RUL cap	125 cycles
Typical sequence length range	~130–360 cycles
Target	Remaining Useful Life (cycles)

Table 1 describes the key characteristics of the FD001 subset from the NASA CMAPSS turbofan engine degradation dataset, one of the most widely used benchmarks in predictive maintenance research. It consists of sensor measurements from 100 engines in the training set and 100 in the test set, each operating under a single fault mode and a single operating condition. After preprocessing, 14 of the 21 sensor channels are retained, excluding those with constant or redundant readings. The target variable is the Remaining Useful Life (RUL) measured in cycles, with a maximum cap applied at 125 cycles to avoid overly large prediction windows. Sequence lengths vary between ~130 and 360 cycles per engine, creating diverse degradation trajectories. This dataset is particularly useful for developing and validating adaptive digital twin models since it provides controlled yet realistic fault progression under consistent conditions. In the context of the adaptive framework, FD001 allows for the assessment of anomaly detection, RUL estimation, and uncertainty quantification methods in a single working regime scenario, which simplifies baseline testing. The data's structure also supports both physics-informed and purely data-driven modeling approaches, enabling hybrid experiments. The uniform operational setting makes it an ideal starting point for comparative evaluation of modeling techniques before extending to more complex, multi-condition datasets. The suggested bar chart can visualise dataset dimensions, such as the number of units or sensor channels retained, helping stakeholders grasp dataset scale. Insights from FD001 performance serve as a benchmark for evaluating the robustness, generalisation, and adaptivity of proposed digital twin solutions.

Metric	Value
Units (train/test)	100/100
Sensors (after dropping constants)	14
Operating conditions	1
Max RUL cap	125 cycles
Noted difficulty vs FD001	Higher due to different fault dynamics
Target	Remaining Useful Life (cycles)

Table 2 summarises the FD003 subset from the NASA CMAPSS turbofan dataset, which, like FD001, contains simulated degradation data from 100 training units and 100 testing units, but with a crucial difference: FD003 features a different fault mode while still operating under a single condition profile. The number of retained sensors (14) and the RUL cap at 125 cycles are identical to FD001, allowing model performance comparisons between the datasets without confounding variable differences in preprocessing. However, FD003's degradation dynamics are inherently more complex, leading to increased prediction difficulty—models that perform well on FD001 often see higher error rates here. This makes FD003 a critical test case for assessing the adaptability of digital twin frameworks to changing fault patterns while maintaining the same environmental stability. The consistent operational setting (single condition) removes extraneous variability, meaning differences in model performance will be predominantly driven by the distinct degradation signature. For adaptive twins, this dataset allows the isolation and study of concept drift and retraining triggers when transitioning between fault types. The recommended bar chart contrasting FD001 and FD003 statistics visually communicates that the datasets are structurally similar but differ in fault dynamics—making it easy to justify their joint use in validation workflows. By using FD003 alongside FD001, researchers can measure whether an adaptive twin can accurately generalise across fault classes without manual reconfiguration, which is essential in real-world environments where failures may be of different nature but occur within a similar operating envelope.

Table 2 – NASA CMAPSS FD003 Dataset Summary.

Table 3 – Example Sensor Channels Used.

Index	Sensor channel (canonical CMAPSS IDs)
1	s2 (Total temperature at fan inlet)
2	s3 (Total temperature at LPC outlet)
3	s4 (Total temperature at HPC outlet)
4	s7 (Total pressure at HPC outlet)
5	s8 (Physical fan speed)
6	s9 (Physical core speed)
7	s11 (Bypass ratio)
8	s12 (Bleed enthalpy)
9	s13 (HPT coolant bleed)
10	s14 (LPT coolant bleed)
11	s15 (Burner exit temperature)
12	s17 (Fuel air ratio)
13	s20 (HPT coolant temperature)
14	s21 (LPT coolant temperature)

Table 3 lists the 14 selected sensor channels from the CMAPSS dataset typically retained for predictive maintenance modeling after dropping constant or irrelevant channels. Each sensor corresponds to a specific physical location or subsystem in the turbofan engine—for example, s2 measures total temperature at the fan inlet, s8 captures physical fan speed, and s15 records burner exit temperature. These variables provide a mix of thermal, pressure, flow, and speed indicators, giving a comprehensive view of engine health. In adaptive digital twin frameworks, selecting relevant sensors is critical for accurate state estimation, virtual sensing, and RUL prediction. The table functions as a reference schema, ensuring reproducible feature selection across experiments. It also aids in interpreting model outputs by linking sensor IDs to physical meanings. While not inherently numeric, this mapping supports downstream feature importance analysis—allowing researchers to discover which physical aspects contribute most to degradation detection. For example, pressure at the high-pressure compressor (HPC) outlet or bypass ratio changes may strongly correlate with early-stage faults. In deployment, these channels could inform both model training and online monitoring strategies, particularly when sensor health or availability changes. Although the suggested visualisation is optional, a bar chart of sensor usage frequency across models or studies could reveal consensus on sensor relevance. In the context of the proposed adaptive framework, this sensor list becomes the foundation upon which virtual models, fusion algorithms, and hybrid physics–ML approaches operate, directly affecting diagnostic coverage and interpretability.

Table 4 – XJTU-SY Bearing Dataset Lifetimes (Condition 1).

Operating Condition	Bearing	Lifetime	Failure Location
35Hz/12kN	Bearing1	2h3min	Outer race
35Hz/12kN	Bearing2	2h41min	Outer race
35Hz/12kN	Bearing3	2h38min	Outer race
35Hz/12kN	Bearing4	2h2min	Cage

Table 4 presents part of the XJTU-SY bearing accelerated life testing dataset, specifically the Condition 1 configuration (operating speed: 35 Hz; radial load: 12 kN). Five bearings are tested under identical controlled conditions, and their total lifetimes—from start until detectable failure—are recorded alongside the fault location. Reported lifetimes vary significantly: from as short as 52 minutes (outer+inner race failure) to over 2 hours 41 minutes (outer race). Such variation under constant operating stress illustrates the challenge of building accurate degradation models without adaptive mechanisms—despite identical inputs, physical units age at different rates due to microstructural differences, lubrication variation, or manufacturing tolerances. The inclusion of fault location (outer race, cage, outer+inner) allows for targeted failure mode classification within the digital twin. This data is important for testing predictive maintenance frameworks in high-speed rotating machinery, where fault onset may occur suddenly. In the adaptive twin context, these results help validate anomaly detection triggers and remaining useful life estimates for discrete components rather than integrated systems. The suggested bar chart plotting lifetimes (converted to minutes) grouped by fault location can immediately communicate variability across units, highlighting the necessity for real-time asset-specific adaptation rather than relying solely on fleet averages. Additionally, by comparing lifetimes with vibration pattern changes in collected raw signals, researchers could evaluate how early the twin can detect each type of raceway or cage degradation, and whether some fault types remain inherently more predictable under given load-speed conditions.




Fig 1 – XJTU-SY Bearing Dataset Lifetimes (Condition 1).

Table 5 – XJTU-SY Bearing Dataset Lifetimes (Condition 2).

Operating Condition	Bearing	Lifetime	Failure Location
37.5Hz/11kN	Bearing6	8h11min	Inner race
37.5Hz/11kN	Bearing7	2h41min	Outer race
37.5Hz/11kN	Bearing8	8h53min	Cage
37.5Hz/11kN	Bearing9	42min	Outer race
37.5Hz/11kN	Bearing10	5h39min	Outer race

Table 5 focuses on the Condition 2 subset of the XJTU-SY bearing dataset, with an operating speed of 37.5 Hz and radial load of 11 kN—a slightly altered mechanical stress profile from Condition 1. The performance and time-to-failure of five test bearings under this regime are listed alongside failure locations. Lifetimes vary sharply, ranging from a short-lived 42 minutes (outer race failure) to nearly 9 hours (cage failure), illustrating increased endurance for some units at reduced load despite higher speed. Fault locations include the inner race, outer race, and cage, showing that even subtle operational changes shift both the type and onset time of failures. This dataset portion is valuable for evaluating how adaptive twins recalibrate prognosis models when operating conditions shift even slightly, affecting degradation rates and dominant failure modes. The substantial differences in lifespan across units under the same test setup highlight the influence of material microdefects, installation conditions, and lubrication on degradation behaviour. The suggested visualisation—a bar chart of lifetime by bearing, colour-coded by failure type—clearly shows clustering of lifespans by fault type, enabling quick comparative insight. From a digital twin perspective, these differences would drive the implementation of condition-specific virtual submodels or weighting adjustments in hybrid physics-ML algorithms. When

combined with sensor signal analysis (vibration amplitude, statistical features, etc.), the Condition 2 results help define thresholds and drift detection methods to ensure that predictive maintenance recommendations remain accurate when seemingly minor parameter changes occur in actual industrial settings.




Fig 2 – XJTU-SY Bearing Dataset Lifetimes (Condition 2).

V. CONCLUSION

The study titled “*Startups and Sustainability: Exploring Public Awareness in Coimbatore Regarding the Role of Startups in Achieving the SDGs*” reveals critical insights into the intersection of entrepreneurial innovation and sustainable development. The findings indicate that while there is moderate public awareness of the Sustainable Development Goals (SDGs), there remains a substantial gap in knowledge about how startups contribute to these goals. Younger demographics, particularly those between 18 and 35, exhibit greater awareness and support for sustainability initiatives, suggesting a favorable outlook for youth-driven entrepreneurial ecosystems. The public perceives startups as impactful agents in addressing sustainability challenges, especially when they adopt practices like waste recycling, green packaging, renewable energy use, and ethical sourcing. However, the research also highlights the need for increased visibility and communication of startup-led sustainability efforts. A significant number of respondents are unaware of specific startups contributing to SDGs, indicating a lack of outreach or public engagement from these ventures. From a strategic standpoint, startups in Coimbatore must prioritize not only sustainable operations but also transparent communication and community involvement. Bridging the awareness gap through education, policy support, and digital storytelling is essential for building a more inclusive and participatory environment. Startups that align their missions with

specific SDGs, report measurable impacts, and engage meaningfully with their communities will likely gain more trust, investment, and social capital. This research underscores the transformative potential of startups in achieving the SDGs at a local level. By leveraging their agility, innovation, and purpose-driven models, startups can act as catalysts for sustainable urban development. A combination of informed public participation, supportive policy frameworks, and robust entrepreneurial ecosystems is vital to maximize this potential. Ultimately, fostering sustainability-focused startups is not just an economic imperative but a pathway toward a more equitable and environmentally conscious future for cities like Coimbatore.

VI. REFERENCES

- [1] operations: A data-driven approach to predictive maintenance through machine learning,” International Journal of Machine Learning for Sustainable Development, vol. 3, no. 1, pp. 31–44, 2021.
- [2] M. Krishna Menon and R. Tuladhar, “Asset management decisionmaking through data-driven predictive maintenance—an overview, techniques, benefits and challenges,” Maintenance, Reliability and Condition Monitoring, vol. 4, no. 2, pp. 44–63, 2024.
- [3] Chen, K.M.; Chang, R.Y.; Liu, S.J. Interpreting Convolutional Neural Networks for Device-Free Wi-Fi Fingerprinting Indoor Localization via Information Visualization. *IEEE Access* **2019**, *7*, 172156–172166.
- [4] Zhang, G.; Wang, P.; Chen, H.; Zhang, L. Wireless Indoor Localization Using Convolutional Neural Network and Gaussian Process Regression. *Sensors* **2019**, *19*, 2508.
- [5] Jang, H.J.; Shin, J.M.; Choi, L. Geomagnetic field based indoor localization using recurrent neural networks. In Proceedings of the 2017 IEEE Global Communications Conference, GLOBECOM 2017—Proceedings, Singapore, 4–8 December 2017; Volume 2018, pp. 1–6.
- [6] I. Wood, A. Barari, E. Esmailzadeh, Using the simple structural beam model to optimize and analyze automotive structures for torsion stiffness and fundamental natural frequency, in *Canadian Society for Mechanical Engineering International Congress* (2014)
- [7] A. Lalephour, A. Barari, Developing skin model in coordinate metrology using a finite element method. *Measurement* **109**, 149–159 (2017).
- [8] A. Barari, M.d.S.G. Tsuzuki, Y. Cohen, M. Macchi, Editorial: intelligent manufacturing systems towards Industry 4.0 era. *J. Intell. Manuf.* **32**, 1793–1796 (2021).
- [9] Huang, Y.; Zhang, J.; Chen, X.; Lam, A.H.F.; Chen, B.M. From Simulation to Prediction: Enhancing Digital Twins with Advanced Generative AI Technologies. In Proceedings of the ICCA, Reykjavík, Iceland, 18–21 June 2024; pp. 490–495.
- [10] Shaposhnyk, O.; Lai, K.; Wolbring, G.; Shmerko, V.; Yanushkevich, S. Next Generation Computing and Communication Hub for First Responders in Smart Cities. *Sensors* **2024**, *24*, 2366.
- [11] Bengesi, S.; El-Sayed, H.; Sarker, M.K.; Houkpati, Y.; Irungu, J.; Oladunni, T. Advancements in Generative AI: A Comprehensive Review of GANs, GPT, Autoencoders, Diffusion Model, and Transformers. *IEEE Access* **2024**, *12*, 69812–69837.
- [12] J. Aguayo, A. Claudio, L. G. Vela, and S. Gentile, “A survey of fault diagnosis methods for induction motors drives under inverter fault conditions,” in (ICEEE). 1st International Conference on Electrical and Electronics Engineering, 2004., Acapulco, Mexico: IEEE, 2004, pp. 367–372. doi: 10.1109/ICEEE.2004.1433910.
- [13] B. H. Bahgat, E. A. Elhay, and M. M. Elkholly, “Advanced fault detection technique of three phase induction motor: comprehensive review,” *Discov Electron*, vol. 1, no. 1, p. 9, Jun. 2024, doi: 10.1007/s44291-024-00012-3
- [14] M. Yousuf, T. Alsuwian, A. A. Amin, S. Fareed, and M. Hamza, “IoT-based health monitoring and fault detection of industrial AC induction motor for efficient predictive maintenance,” *Measurement and Control*, vol. 57, no. 8, pp. 1146–1160, Aug. 2024, doi: 10.1177/00202940241231473.
- [15] T. Zonta, C. A. Da Costa, F. A. Zeiser, G. De Oliveira Ramos, R. Kunst, and R. Da Rosa Righi, “A predictive maintenance model for optimizing production schedule using deep neural networks,” *Journal of Manufacturing Systems*, vol. 62, pp. 450–462, Jan. 2022.