FEIC 2026

AI-Enabled IoT Framework for Fleet Condition Monitoring and Resilient Transportation Operations

J. A. Frank · K. A. Abdulsalam · A. A. Ayorinde

Department of Electrical and Electronics Engineering, University of Lagos, Nigeria



Roadmap of the talk

FIVE PARTS

1. Background and related work
2. Four-layer architecture, sensing and connectivity
3. Features, vehicle-grouped splits, ANN and decisions
4. KPIs, evaluation protocol and robustness testing
5. Claims, limitations and path to a pilot

Type of paper

A framework and method contribution, not a field-results paper.

Its validation protocol is designed so future results can be independently verified.

Why fleet monitoring matters

Why it matters

Road-freight fleets carry goods and keep supply chains moving.

Where maintenance resources and digital infrastructure are constrained, a sudden fault is more disruptive and costly.

IOT AND AI: COMPLEMENTARY ROLES

1. **Sensors acquire**
   Operating and condition data from the vehicle

2. **Networks move**
   Selected information to edge or cloud services

3. **Models identify**
   Deviations from expected behaviour

Scope: vehicles and fleet operations

1

Unexpected faults

Sudden vehicle failures disrupt transport services

2

Late decisions

Poor visibility delays maintenance decisions

3

Weak links

Intermittent communications limit data reaching operators

In scope

Vehicles and fleet operations: GNSS, load, temperature and vibration data turned into maintenance decisions.

Out of scope

Structural-health monitoring of roads, bridges, pavements or roadside infrastructure.

What prior reviews show

Arena et al. [1]
Sensors and AI used to estimate equipment condition and remaining useful life.

Giannoulidis & Gounaris [2]
Context-aware, unsupervised anomaly detection for fleets, as operating context shifts.

Sharma et al. [3]
Machine-learning diagnosis, limited by sensor availability and scarce data.

Hossain [4]; Mahale [5]
Multi-sensor fusion, explainable AI, scalability and cross-vehicle validation.

Continuing challenges

Data quality · model transferability · explainability · deployment validation [1]–[4]

How the literature was selected

SEARCH CONCEPTS (TERMS WITHIN A GROUP COMBINED WITH OR)

Vehicles
vehicle · fleet · transport

AND

Maintenance topics
predictive maintenance · condition monitoring · fault diagnosis · anomaly detection

AND

Technology
IoT · telematics · machine learning · AI

SELECTION CRITERIA

- Peer-reviewed papers from 2020 to 2025 prioritised
- Earlier foundational work only where needed
- Fleet data, AI diagnosis, IoT or maintenance metrics
- Manufacturing-only and civil-structure papers excluded

Narrative review
Not a systematic, PRISMA-compliant search.

Accuracy alone is not resilience

The gap

A fleet study must also report event-level warning performance, false alarms, repair duration, failure intervals, availability and intervention cost.

Four objectives

1. Specify a multi-sensor IoT architecture for fleet monitoring
2. Formulate health estimation and maintenance decision objectives
3. Define a leakage-resistant validation procedure
4. Identify the indicators needed to assess resilience

A four-layer architecture

	COMPONENTS		OUTPUT
1 Sensing	GNSS; load or axle measurement; temperature; tri-axial accelerometer	→	Timestamped position, speed, load, temperature, acceleration
2 Edge and communication	Filtering, buffering, timestamp alignment; cellular, NB-IoT or other link	→	Quality-checked features; queued packets during outages
3 Cloud analytics	Feature store; ANN or benchmark model; model monitoring	→	Health state, anomaly score and confidence
4 Decision support	Alert prioritisation; work-order interface; audit log	→	Inspection recommendation, event record, maintenance feedback

Sensing, calibration and the 50 Hz limit

Installation and calibration

- Load cells calibrated against traceable reference loads
- Record date, reference, zero offset, span error, uncertainty, interval
- Temperature probes: location-specific placement and checks
- Accelerometers: rigid mount, fixed axes, anti-alias filtering

< 25 Hz

At 50 Hz sampling, the Nyquist limit confines analysis to low-frequency body, suspension and road-induced vibration.

Bearing-fault diagnosis is not claimed.

Designed for intermittent connectivity

1 Buffer

Gateway queues timestamped observations during outages

→

2 Transmit

Queued data is sent when a link becomes available

→

3 Summarise

Raw vibration reduced to summary features locally when bandwidth is limited

Every record carries

- Vehicle pseudonym
- Trip identifier
- Sensor identifier
- Timestamp
- Unit
- Calibration version
- Data-quality flag

From signals to a fixed feature vector

1. **Align and check**
   Timestamps, duplicates, physical ranges, missing intervals

2. **Handle gaps**
   Interpolate short gaps only where justified; mark long gaps missing

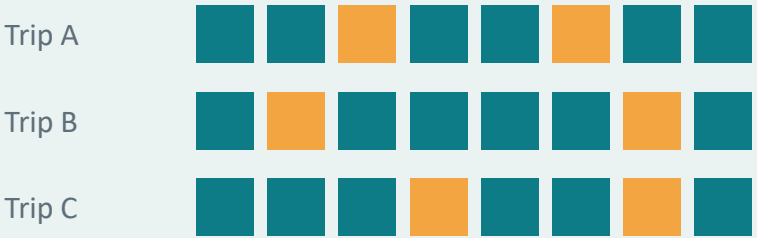
3. **Window**
   Fixed-duration windows; feature vector fixed before training

Candidate features

- Mean
- Standard deviation
- RMS
- Peak-to-peak
- Spectral energy
- Speed statistics
- Route deviation
- Load change
- Temperature range
- Dwell duration

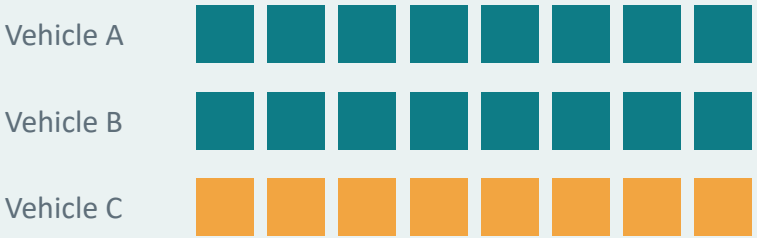
Split by vehicle, not by window

Random split of windows



Correlated windows land on both sides, so test scores can be inflated.

Grouped split by vehicle



Held-out vehicles give an honest estimate of generalisation.

Training Test

Fit scaling, imputation and balancing on training data only

Health-state model and ANN classifier

Latent health state

$$x_{k+1} = f(x_k, u_k, w_k)$$

$$z_k = h(x_k, v_k)$$

x health state, z sensors, u context (speed, payload, route), w process uncertainty, v noise. f and h must be identified per subsystem.

ANN classifier

$$\hat{p}_k = \mathrm{softmax}(W_2\, \sigma(W_1 \varphi_k + b_1) + b_2)$$

$$J_CE = -(1/N)\, \Sigma\, y_n c \log \hat{p}_n c$$

ϕ is the d-dimensional feature vector and σ a nonlinearity such as ReLU. Class weighting or oversampling stays inside the training partition; hyperparameters are tuned on grouped validation data.

Linking health estimates to decisions

CANDIDATE ACTIONS

- **Continue operation**
- **Inspect**
- **Schedule maintenance**
- **Remove from service**

Minimise expected cost over horizon H

- **C_F** Failure cost × failure probability
- **C_I** Inspection cost
- **C_M** Maintenance cost
- **C_D** Downtime cost

Subject to safety and service constraints. Numerical optimisation needs fleet-specific cost and reliability data.

Engineering KPIs for resilience

KPI	Definition
MTBF (Mean Time Between Failure)	Total operating time ÷ observed failures
MTTR (Mean Time to repair/Recovery)	Active repair time ÷ completed repairs (parts waiting reported separately)
Availability	Uptime ÷ (uptime + downtime)
Event recall	Detected eligible events ÷ confirmed eligible events
False alarms per vehicle-day	Unconfirmed alerts ÷ monitored vehicle-days
Warning lead time	Failure time minus first qualifying alert (median, IQR, event count)
Maintenance cost	Inspection + repair + downtime costs, with stated counterfactual

A reproducible evaluation protocol

1 Labels

- Dated work orders, inspections, trouble codes
- Fixed prediction horizon for pre-failure
- Two independent reviewers; report agreement

2 Baselines

- Documented rule-based threshold model
- One-class SVM; report kernel, ν, γ
- Tune on training or validation data only

3 Reporting

- Confusion matrix; per-class P, R, F1
- Macro and weighted averages with CIs
- Repeated seeds: mean ± SD

Effective sample size = independent events and vehicles, not overlapping windows

Robustness and communication testing

Packet-loss masks

State whether losses are independent or bursty, the percentage missing, affected channels and the random seed.

Sensor-noise perturbations

State the distribution and the amplitude relative to each sensor's calibrated range.

Applied after the train/test split. Report for every condition:

- **Event recall**
- **False alarms per vehicle-day**
- **Macro-F1**
- **Confidence intervals**

What the paper does and does not claim

Contribution

- Implementable system architecture
- Engineering model and decision objective
- Leakage-resistant evaluation protocol

Withdrawn claims

- Earlier classification metrics
- Baseline comparisons and inference latency
- Downtime, lead-time and carbon-saving figures

Why: the reported figures could not be reconciled, no auditable dataset, confusion matrix or logs were available, and random window splits may inflate results.

Limitations and the path to a pilot

Not yet established

- Field performance
- Generalisation across fleets
- Maintenance savings and emissions cuts
- Measured MTBF / MTTR gains

A future pilot must provide

- Data dictionary and event register
- Calibration records and software version
- Split identifiers and model configuration
- Repeated-run outputs, confusion matrices
- Timestamped alert and repair records
- Ethics approval or exemption for human data

Claim benefits only after a controlled before-and-after or matched-control evaluation

Key takeaways

1. A four-layer architecture from sensing to maintenance decisions
2. A leakage-resistant, event-based protocol with reliability KPIs
3. Next: a documented, vehicle-grouped field evaluation before quantitative claims

Thank you. Questions welcome.

jfrank@unilag.edu.ng · University of Lagos