TECHNICAL WHITEPAPER

Non-intrusive predictive monitoring for railway rolling stock

System architecture, MEMS sensor technology, and residual diagnostic engine based on extended Kalman filter for early anomaly detection in bogies, doors, and bearings.

Document	IN-SIGHT-WP-001 · Public version 1.0
Date	June 2026
Organization	Ingérop Spain — Transport Division (T3)
Program	IN³ Season IV — Ingérop Incubates Your Innovations
Audience	Fleet engineers, maintenance managers, technical directors

1. The Problem: The Digitalization Gap

Approximately 40–50% of the European passenger railway rolling stock fleet is in the mid-life range of 15 to 30 years: fully operational vehicles whose OEM warranty has expired and which mostly lack a modern TCMS system capable of generating diagnostic data. For these assets, there is no native telemetry: the operator works blindly, alternating calendar-based preventive maintenance—which replaces healthy components—with corrective maintenance after failure, whose downtime cost can reach €15,000 per hour of lost service.

IN-SIGHT closes this gap with a radical design premise: **total non-intrusiveness.** The system captures the vehicle's vital signs—vibration, acoustic signature, temperature, position, and dynamics—using external sensor pods, without connecting to any vehicle data bus (MVB, CANbus, vehicular Ethernet), without accessing the TCMS, and without modifying any functional safety certification.

> **Principle of Non-Intrusiveness**
> The IN-SIGHT kit does not connect to any vehicle system. It is installed as an independent onboard equipment, with its own power supply and dedicated cellular communication. This eliminates the main regulatory and commercial barrier of traditional monitoring solutions, making the system manufacturer-agnostic: it works on any fleet, from any OEM, of any age.

2. Three-layer architecture

The system is structured into three decoupled functional layers, communicating via standard protocols. Each layer can evolve independently without affecting the others.

Layer	Components	Function	Technology
1 — Edge	Pod A (bogie, ×2) · Pod B (doors, ×4) · IoT Edge Gateway	Multi-physical capture, local feature extraction (FFT, RMS, envelope), buffer for coverage loss	MEMS COTS + Raspberry Pi CM5 + Azure IoT Edge
2 — Transport	4G cellular module + eSIM · WiFi backup in depot	Secure uplink of low-density features (no raw signal)	MQTT over TLS 1.3
3 — Cloud	Ingestion · time series · alert engine · dashboard	Residual diagnostic, alert management, automatic reports, administration portal	Azure IoT Hub · Data Explorer · Functions · Power BI

Edge processing is deliberate: the high-frequency vibration signal is locally reduced to a compact feature vector (spectral band energies, kurtosis, crest factor, acoustic envelope), which reduces the transmitted volume by orders of magnitude and makes operation feasible with intermittent cellular connectivity — the local buffer covers weeks without coverage, with automatic retroactive synchronization.

3. Industrial-grade MEMS sensor technology

The entire capture chain uses standard commercial off-the-shelf (COTS) industrial-grade components, eliminating single supplier dependencies and obsolescence risks. The selection by subsystem:

Magnitude	Sensor	Key specification	Subsystem
High-frequency vibration	ADXL357B (triaxial MEMS)	Low noise, bandwidth suitable for bearing signatures	Bogie / grease box
Acoustic signature	ICS-43434 (digital MEMS microphone)	I2S output, immune to EMI in analog chain	Bogie
Contactless temperature	MLX90614 (IR)	Contactless surface measurement of the axle	Grease box
Position and dynamics	u-blox NEO-M9N (GNSS)	Operational context: speed, event location	Vehicle
Door pressure / position	MLH016 + ADXL355 + SS49E	Cycle profiling: time-pressure curve of each cycle	Doors (×4)

In Pod B, each door integrates an ESP32-S3 microcontroller with a W5500 Ethernet interface that digitizes the signal at the source, avoiding long runs of analog cable susceptible to electromagnetic interference in the railway environment. The onboard backbone is industrial Ethernet with M12 connectors according to EN 50155.

4. The diagnostic engine: EKF and statistical residual

4.1 The golden run as a model, not as a recording

The core of the system is the continuous comparison between what the vehicle should be doing if it were healthy and what the sensors actually measure. The health reference—the golden run—is not a historical recording: it is a parametric dynamic model that, given the instantaneous operational conditions (speed, load, ambient temperature, traction regime), predicts what the sensors should read. This solves the structural flaw of systems based on historical averages: a healthy bearing at 60 km/h with 12 t/axle produces a very different signature from the same bearing at 220 km/h with 16 t/axle, and directly comparing them generates systematic false positives.

4.2 Extended Kalman Filter

The model-measurement confrontation is formalized through an extended Kalman filter (EKF), the canonical state estimation tool in nonlinear dynamic systems, validated for over six decades in safety-critical industries. At each instant, the EKF produces the **innovation**: the difference between the health model prediction and the actual multi-modal observation. Under the healthy asset hypothesis, the normalized innovation squared (NIS) follows a known chi-square distribution.

> **Mathematically Controlled False Alarm Rate**
> The fact that the NIS follows a chi-square distribution under the health hypothesis allows setting detection thresholds with a false alarm probability chosen a priori (typically 1e-3 or 1e-4), instead of calibrating them empirically by trial and error. To avoid triggers from benign transients, a sequential probability ratio test (SPRT) accumulates evidence over time before raising a maintenance alert.

4.3 Multi-physical Fusion by Design

Unlike solutions that apply a detector per modality and combine results with heuristic rules (AND/OR, voting), the EKF fuses vibration, acoustic, and temperature at the innovation level, capturing correlations between modalities. A slight but consistent deviation across the three modalities—invisible to each isolated detector—produces a clearly detectable joint NIS. This is the formal basis for the system's reduction of false positives.

4.4 Second stage: machine learning on residuals

The EKF does not decide maintenance: it provides statistically characterized raw data. A second stage of machine learning (gradient boosting with quantile regression) operates on features derived from the innovation sequence—trends, residual spectrum, accumulated degradation, operational context—to classify failure modes and estimate remaining useful life (RUL) with P10/P50/P90 confidence intervals. The combination of a rigorous physical model and modern ML is what differentiates IN-SIGHT from purely data-driven solutions that behave as black boxes: each alert is explainable and auditable, a design requirement aligned with the European AI Act.

5. Validation plan

The validation of the residual diagnostic engine follows a three-level progression with measurable criteria: (1) public reference datasets — Case Western Reserve University Bearing Dataset for detection, FEMTO-ST/PRONOSTIA for RUL prognosis — with comparison against published benchmarks; (2) proprietary test bench with synchronized multi-modal acquisition on healthy and induced-fault components, measuring the false positive rate of the fused system versus mono-modal detectors; and (3) instrumentation of a real vehicle in commercial service, with an acquisition campaign on healthy assets and statistical consistency validation of the NIS under real operational conditions.

For further details on the system architecture, calibration protocol, or an evaluation of your fleet, please contact the team at **in3-insight.cloud**.