AI-Driven Anomaly Detection for Industrial Motor Fault Diagnosis

A Proof-of-Concept developed by AIRS ML on data shared by ABB.

Background

- Three-phase industrial motors are the workhorses of high-value manufacturing, and unexpected faults cause severe downtime, safety incidents and financial loss.
- Conventional supervised machine-learning approaches to fault diagnosis depend on large, labelled datasets of failure modes, but in real industrial settings this data is:
  - **Rare & expensive** — faults are infrequent and costly to recreate.
  - **Reliant on labels** — every fault class must be observed and annotated.
  - **Not scalable** — models do not transfer across motors, plants or operating regimes.
- A scalable solution must instead learn from *healthy* operation alone and flag deviations as potential faults.

Objective

- Develop a deep-learning model that detects motor faults using **only normal operating data**, eliminating the need for labelled fault datasets.
- Validate the approach on a real industrial-scale 3-phase motor.
- Demonstrate clear separation between normal operation and previously-unseen fault types.

Fig 1. ABB M3BP 400LB 4B3 — 3-phase industrial motor.

Industrial Setup

Motor	**M3BP 400LB 4B3**
Voltage / Frequency	**690 V · 50 Hz**
Power / Current	**630 kW · 629 A**
Speed / Poles	**1491 rpm · 4**
Signals captured	**6 (3×I, 3×V)**
Samples / class	**20 each: Normal, Fault A, Fault B**

Proposed Approach

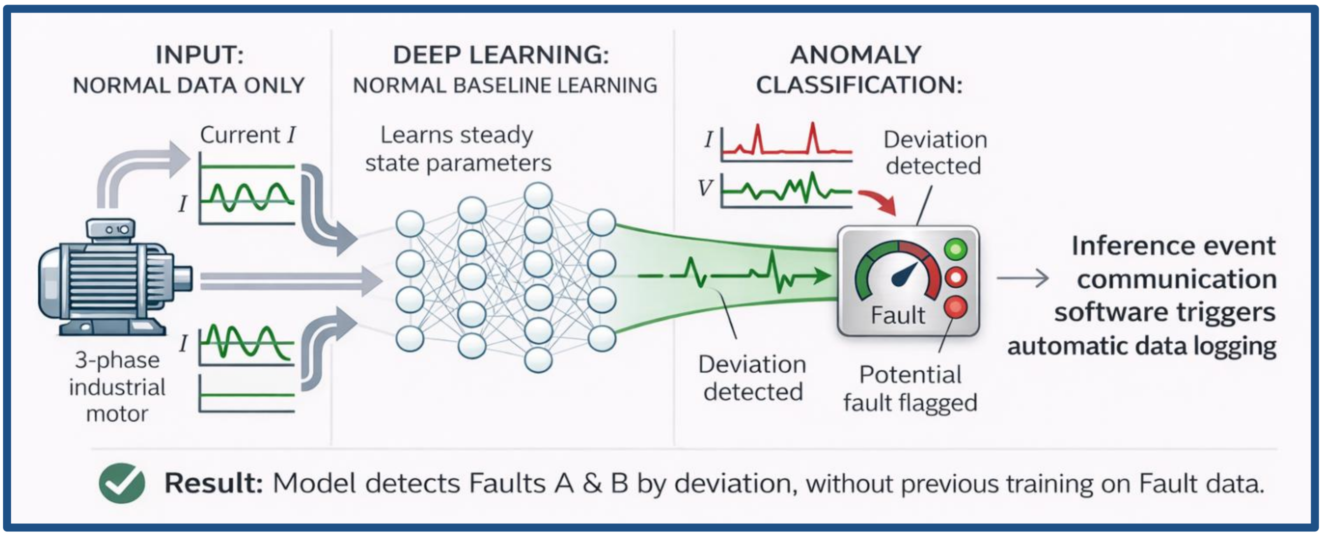

Fig 2. Pipeline — only healthy data feeds the model. At inference, deviations from the learnt baseline are scored as potential faults and trigger automatic data logging.

Machine Learning

- An **unsupervised deep-learning model** was trained exclusively on the *Normal* class, learning a compact representation of the motor's steady-state current and voltage signatures.
- At inference, the per-window reconstruction error rises sharply for any operating regime the model has not seen before, providing a continuous health score that requires no fault labels.
- Sample-level mean reconstruction error gives a single per-sample metric used to rank and classify health state, forming the basis for AUC analysis.
- The model output is consumed by an **inference event communication layer** which automatically logs flagged windows for downstream review.

Results

Trained exclusively on normal data, the model achieved

100%

diagnostic accuracy — detecting both fault types without ever seeing a single fault during training.

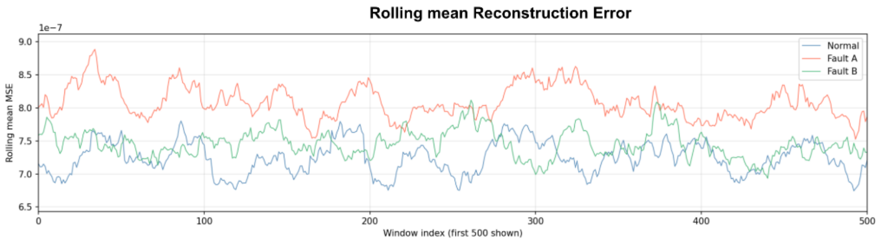

Fig 3. Our model's rolling mean reconstruction error helps increase specificity

- Normal vs Fault A → **clear separation**
- Normal vs Fault B → **clear separation**
- Robust on a **very small dataset** (20 samples / class)

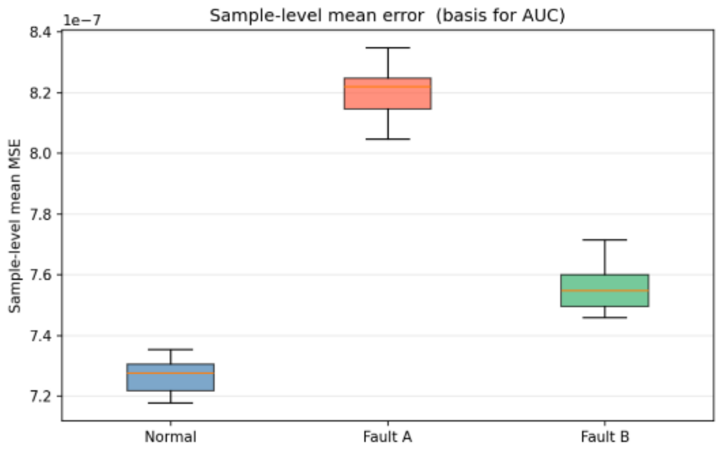

Fig 4. Sample-level mean reconstruction error — non-overlapping distributions across all three classes.

Conclusion

- The PoC demonstrates that **label-free anomaly detection** is viable on real industrial 3-phase motor data: faults are detected reliably without ever being observed during training.
- Removing the labelled-fault-data bottleneck makes the approach **scalable** across motor types, plants and operating conditions — a critical enabler for fleet-wide condition monitoring.
- Combined with AIRS ML's edge-AI deployment stack, this method supports **real-time predictive maintenance** with low latency and strong data-privacy guarantees.

Interested in a similar PoC for your equipment? Get in touch with the AIRS ML team at **prateek.tripathi@airsml.co.uk** to discuss anomaly-detection and predictive-maintenance solutions for your industrial assets.