Applied Structured Corrective Diagnosis of a Legacy Autonomous Mobile Robot: A Comparison of Three Prioritisation Strategies

Zen Hassoun[1] and Hazem Eissa[1]

[1] *School of Computer Science and Engineering, University of Sunderland, United Kingdom*
bi54zr@student.sunderland.ac.uk
hazem.eissa@sunderland.ac.uk

ABSTRACT

Corrective diagnosis of legacy robotic assets is routinely performed in operating contexts where manufacturer documentation is incomplete, telemetry is unavailable, and the failure mode is non-progressive. These operations fall outside the typical scope of both data-driven and knowledge-based PHM techniques. This paper presents a structured framework for diagnosing legacy AMR startup failures under limited-information conditions, together with a comparison of three diagnostic prioritisation strategies applied to the same practical case: (i) a qualitative power-tracing logic that combines Fault Tree Analysis (FTA), Failure Mode and Effects Analysis (FMEA), and a decision flowchart, which was the method applied to the asset; (ii) an analytical FMEA with Risk Priority Number (RPN) scoring; and (iii) an analytical decision-cost ordering that ranks investigation steps by expected total cost. Methods (ii) and (iii) are then applied retrospectively to the same case using the operator's experience gained during the practical investigation to assign scores and parameters. The case is a no-start fault on a legacy MiR100 autonomous mobile robot in a university laboratory. The practical investigation, following Method (i), traced the fault to a blown fuse in the robot-side input path; replacement and a brief charging recovery restored normal operation, confirmed via the built-in hardware-health diagnostic. Applied retrospectively, Method (ii) would have ordered the investigation by descending RPN, reaching the fuse on the seventh step, Method (iii) would have ordered by the lowest per-step decision cost, reaching the fuse within minutes. We compare the three methods on time-to-confirmation and total cost, and found that Methods (i) and (ii) reach the fault with comparable total time and cost despite different orderings, while Method (iii) reaches it approximately an order of magnitude faster and cheaper in this case.

1. INTRODUCTION

Autonomous Mobile Robots (AMRs) are increasingly deployed in industrial, logistics, and research settings, and their reliable operation depends on effective fault diagnosis when failures occur. The Prognostics and Health Management (PHM) literature has produced substantial advances in data-driven (Pan et al., 2025; Li et al., 2024) and knowledge-based (Bouhadra & Forest, 2024) approaches to robotic fault diagnosis. Yet, a common operational scenario remains poorly addressed by these approaches: an operator faced with a legacy or out-of-support robotic asset that has stopped functioning, with limited manufacturer documentation, no instrumented telemetry channel, and a binary failure symptom rather than a degradation trajectory. Data-driven techniques require training data that the operator lacks, and knowledge-based systems require curated rule bases that vendors have not supplied.

In such settings, the operator typically falls back on structured reliability-engineering tools (FTA, FMEA, decision flowcharts) used qualitatively to organise the investigation. A natural question is whether such a qualitative, practical approach can be improved by overlaying quantitative prioritisation, either by RPN scoring or by a per-step decision-cost ordering. The contribution of this paper is to investigate this question by comparing three diagnostic prioritisation methods, applied to the same practical case:

1. *Qualitative power-tracing logic:* combining FTA, FMEA, and a decision flowchart, applied practically to the asset.
2. *FMEA with RPN scoring:* applied analytically to the same case, with S/O/D ratings assigned by the operator based on the experience gained during the practical investigation.
3. *Decision-cost ordering:* applied analytically to the same case, using per-step time and replacement-cost data captured during the practical investigation.

The case study is a no-start fault on a legacy MiR100 AMR deployed in a university laboratory. Method (i) was the actual investigation. Methods (ii) and (iii) are applied retrospectively to enable a controlled comparison on the same case. The three methods are then compared on time-to-confirmation of the fault and total diagnostic cost.

We make no claim that the single case validates any of the three methods in general. We claim it demonstrates how the three methods order the same underlying set of failure modes when applied to the same authentic, limited-information scenario, and we identify the operating conditions under which the differences between them matter most.

The remainder of the paper is organised as follows. Section 2 reviews relevant PHM and reliability-engineering literature and positions the contribution. Section 3 describes the asset, the condition before investigation, and the operating constraints. Section 4 presents the qualitative diagnostic workflow [Method (i)] and its underlying artefacts. Section 5 reports the practical investigation as it was carried out. Section 6 introduces the two analytical methods (FMEA-with-RPN and decision-cost ordering) and the parameters used to apply them. Section 7 reports the comparative analysis of the three methods. Section 8 discusses the comparison analysis and the limitations. Section 9 concludes and identifies future work.

2. BACKGROUND AND RELATED WORK

2.1. Data-Driven PHM for Robotic Systems

Recent surveys have catalogued the rapid growth of data-driven PHM methods for industrial and robotic systems. Deep neural network architectures have been applied to fault diagnosis of industrial manipulators with strong reported performance (Pan et al., 2025), and Bayesian network methods continue to be developed for reliability assessment of complex autonomous systems (Han et al., 2024; Soleimani, Shahbeigi, & Esfahani, 2024). A persistent challenge for this class of methods on legacy or low-volume assets is data scarcity: training reliable detectors typically requires labelled fault data that legacy deployments rarely accumulate. Recent work on small-data PHM partially addresses this gap (Li et al., 2024), but the lower bound on usable telemetry remains substantial.

2.2. Knowledge-Based PHM

Knowledge-based PHM approaches encode expert understanding of failure mechanisms into rule bases, ontologies, or model-based diagnostic graphs (Bouhadra & Forest, 2024). These methods circumvent the data-scarcity problem of data-driven approaches but introduce a knowledge-curation problem: the rule base must be authored and maintained, and the authoring effort scales with asset complexity. For assets outside their original support window, vendor-curated knowledge bases are often unavailable.

2.3. Integrated FTA-FMEA Approaches

The combination of FTA and FMEA is established in reliability engineering (Shafiee, Enjema, & Kolios, 2019; Peeters, Basten, & Tinga, 2018; Yazdi, 2023). FTA provides top-down decomposition of an undesired event; FMEA enumerates failure modes bottom-up with Severity (S), Occurrence (O), and Detection (D) ratings; the Risk Priority Number (RPN) combines these ratings into a single ranking. Integrated approaches use FTA to scope the FMEA's failure-mode set and FMEA to populate the basic events of the fault tree. The principal limitation of RPN as a prioritisation mechanism is documented in the reliability literature.

2.4. AMR-Specific Reliability and Fault Diagnosis

Reliability engineering for AMRs has been addressed across a spectrum of methods, including Bayesian network methods for autonomous systems generally (Han et al., 2024), predictive maintenance approaches for mobile robots via IIoT instrumentation (Krot, Iskierka, Poskart, & Gola, 2022), and broader reviews of fault detection and diagnosis for industrial robots and multi-axis machines (Sabry & Amirulddin, 2024; Luo, Li, Bai, Tang, & Jin, 2024; Mikołajczyk, 2023). Battery and power-system faults are a recurring theme; Battery Management System (BMS) diagnostics in particular have received recent attention (Zhao et al., 2024). These approaches generally assume either instrumented telemetry or documented system architecture; neither is reliable in the legacy-AMR setting.

2.5. Position of the Present Contribution

The framework presented in this paper sets in the operational gap between data-driven and knowledge-based PHM; it requires neither training data nor a curated knowledge base, and is operator-driven rather than model-driven. Its contribution is comparative and methodological rather than algorithmic. We do not propose a new diagnostic algorithm. Instead, we apply three established prioritisation strategies (qualitative power-tracing logic, FMEA with RPN scoring, and decision-cost ordering) to the same authentic legacy-AMR case and report the differences in time and cost they produce. To the best of our knowledge, a controlled comparison of these three strategies on a single practical AMR no-start case has not previously been reported.

3. THE MIR100 CASE

3.1. Asset and Deployment

The asset is a MiR100 autonomous mobile robot in service in a university research laboratory at the University of Sun-

derland. The MiR100 is treated as a representative legacy AMR; the workflow is not specific to this make or model. The robot is approximately seven years old, has been in low-duty research use, and is outside its original manufacturer's support window. Available documentation is limited to publicly accessible user manuals; full electrical schematics and component-level parts lists are not in the operator's possession. The asset has a built-in hardware health diagnostic that reports high-level subsystem status (Computer, Motors, Power System, Sensors, and so on), but this diagnostic is accessed through the robot's web interface and was therefore initially inaccessible because the robot would not power on.

3.2. Condition Before Investigation

At the start of the case study, the only confirmed problem was that the MiR100 did not start or power on, and no maintenance history or diagnostic report was available. A preliminary visual inspection of the asset found no obvious physical fault. Dust accumulation and cobwebs on the chassis suggested the robot had not been used recently. After the investigation was complete and the web interface was again accessible, an internal diagnostic scan dated from 2023 was found, but this was not available at the start. The starting point of the investigation was therefore the observed no-start symptom rather than any known historical fault record. Because no maintenance history or telemetry could be relied on, the diagnostic approach had to be built around logical power-chain tracing, starting from accessible power points and moving systematically through the relevant power-chain stages.

3.3. Operating Constraints

The diagnostic episode was constrained as follows: no manufacturer schematic was available. The operator (under academic supervision) had access to standard laboratory tools (multimeter, screw kit, and power supply generator) and could access the internal power-chain related components by removing the top panel. No specialised vendor diagnostic tooling was available. The operating context permitted a multi-hour diagnostic episode following the accessible safety guidelines manual of the robot. The investigation was carried out under academic supervision and followed a documented pre-work risk assessment covering electrical exposure, battery-handling constraints, and the use of external power sources.

4. DIAGNOSTIC WORKFLOW

4.1. Workflow Architecture of the Practical Investigation

The qualitative workflow that was actually applied to the asset integrates three reliability-engineering artefacts into a single, sequential, auditable process performed by an operator at the point of failure. The workflow has three stages, executed in order:

1. *System decomposition via FTA:* A top-event fault tree is constructed from the operator's working understanding of the asset and publicly accessible documentation, decomposing the observed symptom into branches that distinguish state-of-system causes, battery-side causes, and downstream power-path causes. The fault tree is constructed without requiring the manufacturer's schematic; its branches reflect the operator's functional decomposition of the asset.
2. *Failure-mode enumeration via FMEA:* Failure modes corresponding to the basic events of the fault tree are enumerated in a descriptive FMEA table. The table is used as a component-level diagnostic tool to identify how individual components within the power chain could fail, what effect each failure would have on the robot, what causes could lead to that failure, and how the fault could be detected during practical investigation. In this qualitative form the FMEA does not assign numerical S/O/D scores; its function is to organise the likely faults and support fault isolation, not to calculate which fault is statistically most likely.
3. *Investigation path selection via decision flowchart:* A flowchart operationalises the power-tracing logic, converting the possible fault areas and detection methods identified in the FTA and FMEA into a step-by-step diagnostic procedure that can be followed during the physical investigation of the robot. Each decision node corresponds to an observable condition at the asset, and each branch corresponds to a fault-tree pathway. The flowchart terminates either with a confirmed fault and applied fix, or with a documented escalation if the operator's accessible tests are exhausted.

The workflow architecture is summarised in Fig. 1. The relationship between the three artefacts is complementary rather than independent: the FTA scopes the failure space, the FMEA both scopes and drives the component-level investigation, and the decision flowchart drives the sequence of physical tests. The flowchart is therefore developed last in the model because it depends on the outputs of the other two.

4.2. Fault Tree Analysis

The FTA (shown in Fig. 2) decomposes "Robot Does Not Start" into three intermediate events combined under an OR gate:

- *Robot Not In Startup State*: user-side and interlock causes (active emergency stop, disengaged battery disconnect button, breakers in the off state, charging-mode switch left engaged).
- *No Usable Power In Battery*: battery-side causes including deep discharge, internal cell fault, and upstream

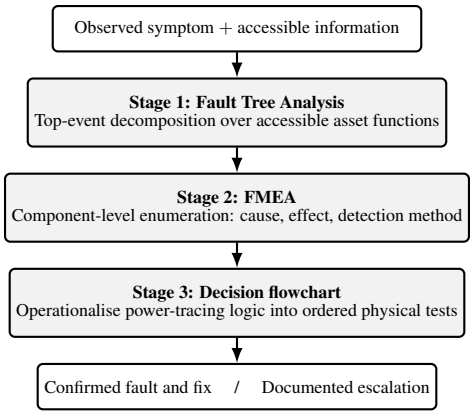

Figure 1. Architecture of the qualitative diagnostic workflow applied practically to the asset. The artefacts produced at each stage feed the next: FTA scopes, FMEA scopes and drives, and the decision flowchart drives the sequence of physical tests.

charging failure. Charging failure is itself decomposed into mains supply, charger fuse, and BMS-related faults.

- *Power Path Interruption*: downstream-of-battery causes including the main power-path fuse, wiring faults, contactors, breakers and relays in the power path, and internal Power Supply Unit (PSU) faults.

The decomposition was constructed iteratively, informed by accessible observation, reasoning, and publicly accessible documentation rather than by a manufacturer's schematic. Where component-level granularity could not be assumed (e.g., inside the battery), basic events are named at the level of observable behaviour rather than presumed internal mechanism.

4.3. Failure-Mode Enumeration

The second part of the diagnostic workflow is the FMEA. It is used here as a component-level diagnostic tool to identify how individual components within the power chain could fail, what effect each failure would have on the robot, what causes could lead to that failure, and how the fault could be detected during practical investigation. The aim is to cover the main reasons identified in the FTA that would lead to the robot not starting and to use the FMEA alongside the FTA to develop the diagnostic flowchart.

Both FMEA and FTA are complementary rather than competing tools. The FTA identifies the main fault branches, and the FMEA expands these branches into a detailed component-level table. The component column of the table was developed by analysing the lower-level branches in the FTA and identifying the components and functions linked to each event within the power chain. Only components related to the power chain were selected; the table is not intended to represent every component in the robot. The remaining columns were developed from publicly available documentation and inspection findings, supported by the literature.

Table 1 lists twelve failure modes derived from the FTA. The first five columns (Component / Function, Failure Mode, Possible Cause, Effects on Robot, Detection Method) are the descriptive columns used by the qualitative workflow and are sufficient to drive the practical investigation reported in Section 5. The remaining columns (Severity S, Occurrence O, Detection D, RPN, and the per-component decision cost C_{step}) are not part of the qualitative workflow; they were added retrospectively to support the analytical methods compared in Section 6 and Section 7, and they are explained where they are used. The ratings reflect this asset class and operating context (a legacy commercial AMR in a university laboratory) and are presented as auditable, based on the operator's experience after the practical investigation, rather than canonical. Different operating contexts (e.g., a continuously-operated logistics floor) would warrant adjusted weights. The *BMS Protection State* row carries the highest RPN, driven by the BMS's opaque internal state. The *Li-ion Battery output* row carries the next-highest RPN, reflecting both the operational severity of battery loss and the moderate detection difficulty under the BMS opacity condition.

4.4. Investigation Flowchart

The flowchart (Fig. 3) operationalises the power-tracing logic developed through the FTA and FMEA into a sequence of physical tests. It is the practical, action-based part of the workflow, in contrast to the FTA and FMEA, which are analytical. Its purpose is to convert the possible fault areas and detection methods identified in the FTA and FMEA into a step-by-step diagnostic procedure that can be followed during the physical investigation, starting from the earliest accessible point in the power chain.

The flowchart's entry condition is the observed top-event symptom; each subsequent decision node maps either to a direct observable (voltage, continuity, LED state) or to a higher-level question (e.g., "is usable power available?"). Two structural features are notable. First, an early decision point "is power restoration possible?" separates faults the operator can resolve directly from faults that require escalation. Second, an inner upstream/downstream loop supports re-investigation if a candidate fault is not confirmed within the power-chain, and an outer loop-back re-attempts the startup condition after each successful fix, ensuring the diagnostic concludes with verified operation rather than with an unconfirmed repair. The flowchart terminates either with a confirmed fault and applied fix, or with a documented escalation if the operator's accessible tests are exhausted.

4.5. Tools and Measurement Methods

The investigation used standard laboratory tools; a digital multimeter for voltage and continuity measurements, a hand tool kit for opening the asset, and a bench-top variable DC

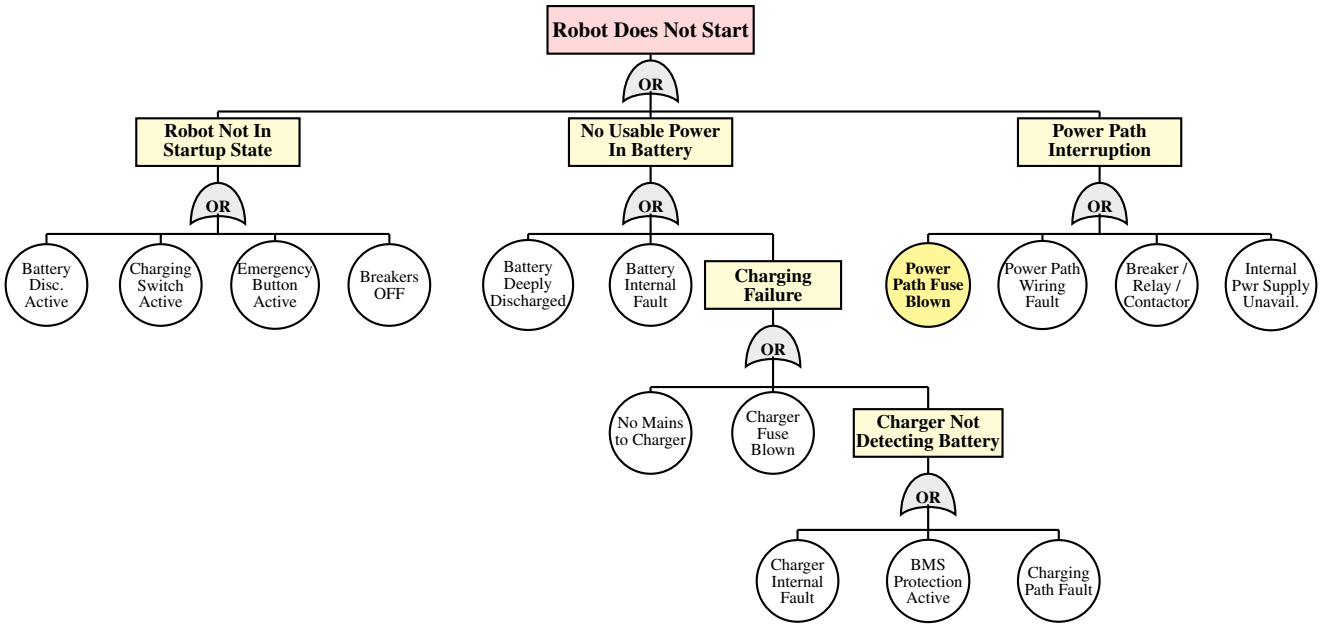

Figure 2. Fault tree for the no-start fault on a legacy MiR100. Canonical OR-gate symbols (curved sides meeting at apex, concave bottom); yellow rectangles are intermediate events; white circles are basic events. The highlighted basic event (*Power Path Fuse Blown*) was identified as the actual cause by the workflow applied in Section 5.

power supply for direct power injection where the Original Equipment Manufacturer (OEM) charger could not be used. The asset's own built-in hardware health diagnostic was used post-resolution to verify that the repair did not introduce collateral faults. No specialised vendor diagnostic tooling was used.

5. PRACTICAL INVESTIGATION

5.1. Initial Startup Attempt and Fault Confirmation

The practical investigation began with a visual inspection of the asset (no physical damage) and preparation of the startup conditions specified in the MiR100 User Guide: connect the battery, set the battery disconnect safety button to the correct state, set the charging switch to the correct state, set the emergency button to the correct state, and set the breakers to the correct state. With these conditions prepared, the power button was used to attempt normal startup. The robot did not power on, confirming the top-level fault used in the diagnostic model: *Robot Does Not Start*. The four startup-related basic events under the *Robot Not In Startup State* branch of the fault tree (Table 1, rows 6, 9, 10, and 11) were eliminated by this preparation. The first decision point in the flowchart of Fig. 3 therefore became *is usable power available?*; since the robot gave no startup response, the first measurement targeted the battery.

Fig. 4 shows the operator's working view of the asset after the top panel was removed. In order to get access to the dense, multi-vendor layout that includes all of the robot's components in which the subsequent measurements were performed.

5.2. Battery Voltage Measurement

The voltage across the battery terminals was measured and read approximately 3.3 V, which is significantly lower than the 24 V specified for the Li-ion battery used in the MiR100. The 3.3 V reading was not taken as direct evidence of the internal cell condition; Li-ion BMS can activate protection circuits when cells are over-discharged or outside the safe operating zone, making the pack appear to have no usable output even if the cells themselves are not at 3.3 V. The result was therefore treated as evidence of no usable output at the battery terminals, with the internal cell and BMS condition unknown because the battery pack was not opened (due to safety concerns). The investigation moved next to the upstream power path, as directed by the flowchart.

5.3. AC Supply and OEM Charger Checks

Before testing the charger output, the AC input supply was verified within the required input range for the charger. An attempt to charge the robot was initiated. The charger's LED illuminated, but the charging/status LEDs did not. Output voltage measurements at the charger, both connected and disconnected from the robot, read approximately 1.8 V (against a nominal output of 24 V). As a result, the investigation moved to the robot-side input path.

5.4. External DC Power Supply and Power Tracing

A bench-top DC power supply rated at 24 V and 5 A was connected to the same robot-side connection point that the charger uses (not directly to the battery), to test whether controlled power applied to the robot's charging input would

Table 1. FMEA for the MiR100 no-start fault. Columns (2 to 6) form the descriptive (qualitative) FMEA used in the practical investigation; columns (7 to 11) are added for the analytical methods of Section 6. Severity, Occurrence, and Detection scores are on a 1 (lowest) to 10 (highest) scale, with rationales recorded in Section 4.3. RPN $= S \times O \times D$. $C_{\text{step}} = h_s \cdot r + p_s \cdot c_s$ is the per-component decision cost, where h_s = technician time (hours) per test, $r = £35/\text{hr}$ labour rate, p_s = renormalised RPN prior, and c_s = component replacement cost. Values of h_s and c_s are given in Section 6.2

#	Component / Function	Failure Mode	Possible Cause	Effects on Robot	Detection Method	S	O	D	RPN	C_{step} (£)
1	Li-ion Battery	No usable output	Deep discharge; internal cell fault	No usable power source for startup	Voltage at pack terminals; under-load test	8	5	6	240	178.85
2	BMS Protection State	Output inhibited	Deep-discharge trigger; protection-flag latch	Battery appears unavailable; charger cannot charge	Pack terminal voltage; charger LED state	7	6	8	336	256.00
3	AC Power Source	No mains supply	Outlet de-energised; plug fault	Charger cannot operate	Verify outlet voltage	6	2	2	24	2.91
4	Charger	No usable output	Blown fuse; no input; detection failure	Cannot charge battery	Continuity on fuse; output voltage	6	4	4	96	47.90
5	Charging Path	Path interrupted	Loose connection; connector fault; blown link	Battery does not receive charge	Continuity test; visual inspection	6	5	4	120	13.50
6	Breaker	Tripped	Overcurrent; downstream fault	Power path open	Visual state; continuity	7	3	2	42	3.57
7	Power Supply (downstream)	No stable output	Internal PSU fault; no battery feed	PC/PLC not powered; robot does not start	Voltage at PSU input/output	8	4	5	160	27.65
8	Contactor	Open / not closing	No control signal; coil fault	Power path open	Continuity across power path; coil response	8	3	5	120	18.25
9	Battery Disconnect Button	Battery isolated	Disconnect engaged; button failure	Battery cannot supply robot	Button position; effect on startup	7	2	1	14	3.24
10	Emergency Button	E-stop active / latched	Button engaged; switch fault	Startup inhibited	Continuity test in ON/OFF state	7	2	1	14	3.24
11	Charging Switch	Robot in charging mode	Switch left engaged; switch fault	Robot does not enter normal startup	Continuity test in ON/OFF state	7	3	1	21	3.42
12	Power Path Fuse	Open / blown	Overcurrent; ageing; vibration fatigue	Input path interrupted; battery not charged	Continuity test across fuse	8	5	2	80	3.10

reach the battery and enable startup. After connection, a startup attempt was made; the robot still did not power on, and the battery terminal voltage still measured 3.3 V. This indicated that power was reaching the robot-side input but was not reaching the battery through the internal input path.

Power was then traced from the charging connector into the robot. Voltage was present at the connector, indicating the connector itself was not the point of interruption. The wiring from the connector was followed further into the chassis, leading to an in-line fuse located in the traced input path.

5.5. Fuse Fault in the Robot-Side Input Path

The fuse identified in the robot-side input path was rated at 20 A / 32 V. Removed from its holder and tested for continuity, it was found to be open (blown), confirming that the input path was interrupted. The fuse was replaced with an equivalent.

After fuse replacement, a startup attempt using the DC power supply, resulted in the robot powering on successfully. This confirmed that the original no-start condition was caused by an interruption in the input path; the fuse is therefore most accurately described as the root cause of that condition, since it was also linked to the lack of usable battery power.

5.6. Charging Recovery and Web-Interface Verification

After the successful startup, the DC power supply was left connected for approximately half an hour to allow partial charging of the battery and revival of the pack. The OEM charger was then reconnected, and it detected the battery and began normal charging indication, with the first yellow state-of-charge LED active. The battery voltage was later measured at approximately 28 V after charging.

With the robot powered on, the web interface became accessible and the built-in hardware-health diagnostic could be read. Fig. 5 shows that all major subsystem groups report nominal status, including the Power System group. This confirmed that after recovery, the no-start condition was not caused by a general failure of the other groups. The confirmed physical fault was the blown fuse in the robot-side input path, and the diagnostic outcome was a restored, operational robot.

The total active diagnostic time for this practical investigation was approximately 120 minutes, broken down approximately as: 25 min for the initial visual inspection, startup-condition preparation, and startup attempt (Section 5.1); 10 min for the battery voltage measurement (Section 5.2); 45 min for the AC supply and charger checks (Section 5.3); 30 min for the bench DC supply connection and the power-tracing into the chassis

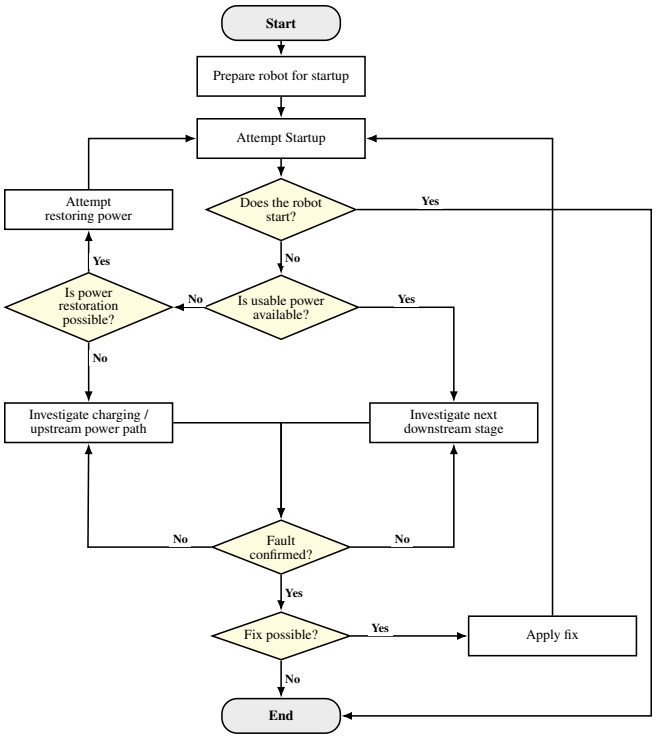

Figure 3. Diagnostic flowchart operationalising the power-tracing logic developed through the FTA and FMEA (Table 1). Yellow diamonds are decision points; white rectangles are operator actions or investigations.

(Section 5.4); and 10 min for the fuse removal, continuity test, replacement, and verification (Section 5.5). The charging recovery (Section 5.6) added approximately 30 minutes of mostly-passive elapsed time.

6. ANALYTICAL DIAGNOSTIC METHODS

The qualitative workflow of Section 4 drove the practical investigation reported in Section 5. We introduce two analytical methods that can be applied to the same FMEA (Table 1) to produce a different ordering of the candidate investigations. Both methods are applied retrospectively in Section 7; their inputs are the operator's experience and the time/cost data captured during the practical investigation.

6.1. FMEA with RPN Scoring

The first analytical method assigns Severity (S), Occurrence (O), and Detection (D) ratings to each failure mode on a 1-10 scale and computes a Risk Priority Number RPN $= S \times O \times D$ for each row. Candidate investigations are then ordered by descending RPN. The ratings used in this paper (columns S, O, D, RPN in Table 1) were assigned by the operator after the practical investigation, based on the experience gained during that investigation. This is a retrospective assignment and is acknowledged, as such, an operator might in general assign different values.

The limitation of RPN as a prioritisation mechanism is that

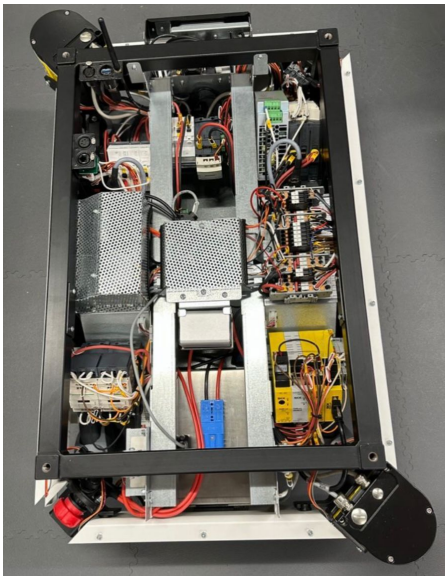

Figure 4. Top-down view of the MiR100 internal layout. It illustrates the robot's components and connections/wiring being in a start-up ready state,

Figure 5. Built-in hardware-health diagnostic on the MiR100 immediately following fuse replacement and battery revival. All subsystems within the Power System group report nominal.

it ranks by risk magnitude rather than by diagnostic cost-effectiveness. High-RPN items are not necessarily the most cost-effective items to investigate first (e.g., the Detection component scales inversely with diagnostic ease). Items that are hardest to detect (and therefore have a high D) are also those whose investigation tends to be expensive in technician time.

6.2. Decision-Cost Ordering

The second analytical method orders candidate investigations by the lowest per-step expected cost. For each failure mode s in Table 1, the per-step decision cost is expressed as

$$C_{\text{step},s} = h_s \cdot r + p_s \cdot c_s \tag{1}$$

where h_s is the technician time required to test for that failure mode, r is the hourly labour rate, p_s is the prior probability that the fault is at s, and c_s is the replacement cost of the com-

Table 2. Technician time per failure-mode test (h_s, hours) used in the decision-cost computation. Values are provisional and reflect typical effort for the indicated detection method.

Failure-mode test	h_s (hr)
Visual / continuity (rows 3, 6, 9, 10, 11, 12)	0.083
Voltage at terminals (rows 1, 5, 7, 8)	0.25
Charger output and disassembly (row 4)	0.50
BMS protection state check (row 2)	0.50

ponent if the fault is confirmed. Candidate investigations are then ordered by ascending $C_{\text{step},s}$. The expected total cost of following a path i through the flowchart up to the confirmed-fault step k is

$$E[C_i] = \sum_k P(\text{fault} = k) \left[r \sum_{j \le k} h_j + c_k \right] \qquad (2)$$

The conditional probabilities p_s used here are derived from the FMEA RPN values renormalised across the twelve candidate failure modes ($p_s = \text{RPN}_s / \sum_s \text{RPN}_s$); other priors are possible and are out of scope for this paper. The cost-per-component values in Table 1 (column C_{step}) are computed using $r = £35/\text{hr}$ (Average UK university laboratory technician rate), the h_s values listed in Table 2, the renormalised p_s values from RPN, and approximate replacement costs c_s; power-path fuse £3, AC source £0 (no replaceable component for the operator), breaker £20, button-style components (battery disconnect, emergency stop, charging switch) £30, charging-path wiring or connectors £50, contactor £100, downstream PSU £150, charger £400, battery pack £900. The BMS unit is not separately replaceable from the battery pack and is therefore assigned $c_{\text{BMS}} = £900$.

The framework is deliberately modest. It does not model downtime cost, lost throughput, or safety penalties, which are out of scope for the university-laboratory operating context. For commercial deployments, those terms can be added to $C_{\text{step},s}$ and $E[C_i]$ without altering the structure of the analysis.

7. COMPARATIVE ANALYSIS OF THE THREE METHODS

This section applies the two analytical methods of Section 6 retrospectively to the same case and compares them with the practical investigation of Section 5 and using Table 1.

7.1. Order Produced by Each Method

Method (i) (Power-tracing logic, Practical): The order actually followed by the operator was: startup-condition visual checks (rows 6, 9, 10, 11) → battery terminal voltage (row 1) → AC mains check (row 3) → charger output (row 4) → bench DC supply and power tracing into the chassis (effectively rows 5 and 7) → power-path fuse continuity (row 12). The order was driven by the power-tracing logic of the flowchart, not by a specific scoring of the failure modes.

Table 3. Comparison of the three diagnostic prioritisation methods on the MiR100 no-start case. Time is rounded to the nearest 5 minutes. Cost combines labour at £35/hr with the fuse replacement (£3).

Method	Time (min)	Cost (£)
i. Power-tracing logic (practical)	120	73.00
ii. FMEA + RPN (retrospective)	125	75.91
iii. Decision-cost (retrospective)	10	8.83

Method (ii) (FMEA with RPN Scoring, Retrospective): The descending RPN order over Table 1 would prioritise: BMS Protection State (336) → Li-ion Battery (240) → Power Supply downstream (160) → Charging Path (120) → Contactor (120) → Charger (96) → Power Path Fuse (80) → Breaker (42) → AC (24) → Charging Switch (21) → Battery Disconnect (14) → Emergency Button (14). On this case, the fault would have been confirmed at the seventh step (Power Path Fuse).

Method (iii) (Decision-Cost Ordering, Retrospective): Ascending order by C_{step} from Table 1 would prioritise: AC (£2.91) → Power Path Fuse (£3.10) → Battery Disconnect (£3.24) → Emergency Button (£3.24) → Charging Switch (£3.42) → Breaker (£3.57) → Charging Path (£13.50) → Contactor (£18.25) → Power Supply downstream (£27.65) → Charger (£47.90) → Li-ion Battery (£178.85) → BMS Protection State (£256.00). On this case, the fault would have been confirmed at the second step (Power Path Fuse), after a brief AC sanity check that has no replacement cost.

The three orderings reach the confirmed fault at different positions: practical investigation reached it through the power-tracing logic; pure RPN would have reached it after six unproductive higher-RPN investigations; decision-cost ordering would have reached it after one preliminary AC check.

7.2. Time and Cost on This Case

Table 3 summarises the three methods on time-to-fault-confirmation and total cost. Time figures sum the h_s values along each method's path up to the step that confirms the fuse; cost is labour at £35/hr plus the fuse replacement (£3). The numbers for Method (i) are taken from the practical investigation and are provisional pending review by the authors. The numbers for Methods (ii) and (iii) are computed retrospectively from Table 1 and the h_s values of Table 2.

Two observations follow from Table 3. First, the practical power-tracing logic and pure RPN ordering produce very similar performance in this case, even though their orderings of the failure modes are quite different. Method (i) reaches the fuse via the power-chain trace from the battery upward; Method (ii) reaches it after sequentially eliminating six higher-RPN candidates. The total time and cost are nearly

the same because the labour-time-per-step values along both paths are dominated by similar investigations of the upper power chain. Second, the decision-cost ordering is dramatically faster and cheaper in this case, by roughly a factor of ten in both time and cost, because the per-step cost calculation puts the fuse as a second in the queue, immediately after an AC sanity check. The pattern of failure on legacy power-path components, where the failed item is often cheap to test and replace, is exactly the pattern for which decision-cost ordering produces its largest benefit.

7.3. Sensitivity

A sensitivity check varying the technician rate from £25/hour to £55/hour and the replacement costs by $\pm 50\%$ preserves the ordering of the three methods across the explored range. The relative advantage of Method (iii) over Methods (i) and (ii) is largely insensitive to the labour rate because the path of Method (iii) is short; its dominant cost is the labour for the short path plus the fuse, both of which are small. Methods (i) and (ii) have longer paths whose total labour scales linearly with r, so their cost grows with the labour rate, while the cost of Method (iii) grows only marginally. The differential between Method (iii) and the other two, therefore, widens at higher labour rates.

8. DISCUSSION

8.1. Comparison Analysis

The three methods produce three different orderings of the same twelve candidate failure modes, but only two distinct performance levels. Methods (i) and (ii) (the practical power-tracing logic and pure RPN ordering) arrive at the fuse with similar total labour and cost, even though they take different routes through the failure space. Method (iii) (decision-cost ordering) arrives at the fuse almost immediately because the fuse is cheap to test and to replace, and its low per-step decision cost places it near the top of the queue. The practical takeaway is that a structured qualitative workflow (Method i), guided by power-tracing logic and the operator's experience, performs similarly to a strict RPN-based prioritisation in this case. Additionally, a simple per-step cost overlay (Method iii) can recover much of the unnecessary investigation effort, especially when the failed item is a low-cost, easy-to-test power-path component. For maintenance teams considering which strategy to adopt, Method i remains the natural default choice; Method ii adds an auditable ranking for comparison; while Method iii is most useful when the component-cost spread across candidate failure modes is large and the failure distribution is not obviously skewed toward high-RPN items.

We emphasise that this is a single-case observation, not a general claim. The relative advantage of Method (iii) is concentrated on the realised fault in this case; the expected-cost benefit under the prior of Table 1 (averaged across all possible fault locations) is more modest. The case shows what decision-cost ordering can do; it does not show what it always does.

8.2. Limitations

We encountered several limitations in the present work. The comparison covers three methods applied to a single fault on a single asset. Generalisation to other no-start faults in AMR robotic systems requires further evaluation. In particular, the advantage of Method (iii) in this case stems from the fault being a cheap-to-test, cheap-to-replace power-path component; but for a different realised fault, the comparative ranking could change.

The S, O, D ratings and the time/cost parameters used by Methods (ii) and (iii) were assigned by the operator after the practical investigation, based on the experience gained during that investigation. This is a controlled comparison but not an a priori test. A genuine forward test would require an operator to score the FMEA, apply the analytical methods, and then perform the practical work; this is left to future work.

RPN scores and the time estimates used in h_s are derived from operator judgement, and three different operators may produce three different rankings. The audit lever of recording the rationale for each score and parameter (Section 4.3, Section 6.2) addresses this by making the inputs reviewable, but it does not eliminate subjectivity itself. We do not claim the values in Table 1 and Table 2 are canonical; we claim they are auditable.

The decision-cost formula (1) deliberately omits downtime cost, lost throughput, and safety penalties. For low-stakes contexts, these terms are negligible; for commercial deployments, they will dominate, and the relative ranking of the methods may change. Extending the cost model is straightforward but requires deployment-specific values that we cannot generalise.

8.3. Relationship to Existing PHM Approaches

Data-driven PHM techniques (Pan et al., 2025; Li et al., 2024) require labelled fault data and substantial telemetry to train reliable detectors, neither of which is typically available on legacy or sealed assets. Knowledge-based PHM systems (Bouhadra & Forest, 2024) apply but require expert-curated rule bases that may not exist for assets outside their original support window. The qualitative diagnostic workflow fills the operational gap between these conditions: it requires no training data, minimal documentation, and is operator-driven rather than model-driven, while the analytical methods provide complementary quantitative prioritisation that can be applied retrospectively to the same FMEA. The comparative observation we contribute (the use of three methods produces different orderings but only two distinct performance levels

in this case, with decision-cost ordering most efficient when the fault is on a cheap-to-test component) is, to the best of our knowledge, not previously reported for AMR corrective diagnosis.

9. CONCLUSION AND FUTURE WORK

This paper has presented a structured framework for diagnosing legacy AMR startup failures under limited-information conditions, and has compared three diagnostic prioritisation methods applied to the same case: a qualitative power-tracing logic combining FTA, FMEA, and a decision flowchart (applied practically); FMEA with RPN scoring; and a per-step decision-cost ordering (both applied analytically). The case is a no-start fault on a legacy MiR100 AMR in a university laboratory.

The practical investigation, following the qualitative power-tracing logic, traced the fault through the upstream and downstream branches of the power chain to a blown 20 A / 32 V fuse in the robot-side input path; replacement of the fuse and a brief charging recovery restored normal operation. Applied retrospectively to the same case, FMEA-with-RPN would have ordered investigations by descending RPN and reached the fuse at the seventh step, with total time and cost comparable to the practical investigation. Decision-cost ordering would have reached the fuse at the second step, with total time and cost approximately an order of magnitude lower than the other two methods. The comparison is summarised in Table 3.

The principal observation is that the qualitative power-tracing logic and the analytical RPN ordering produce different ordering paths, but very similar total time and cost on this case, while decision-cost ordering produces a substantially lower time and cost because the realised fault was on a cheap to test and to replace component (the fuse). The relative advantage of decision-cost ordering is concentrated on this pattern of failure and is more modest in expectation across all possible fault locations.

Future work will include applying the same three-method comparison on different AMR platforms to establish how often the cheap-component pattern dominates and how often it does not. Also, incorporating downtime, throughput, and safety-penalty terms for deployment in commercial settings, and quantifying the conditions under which the relative ranking of the three methods changes.

ACKNOWLEDGEMENTS

The authors would like to thank the technical staff at the University of Sunderland for their specialised knowledge and laboratory access support.

REFERENCES

Bouhadra, A., & Forest, F. (2024). A review of knowledge-based prognostics and health management approaches. *International Journal of Prognostics and Health Management*, *15*(2). doi: 10.36001/ijphm.2024.v15i2.3986

Han, P., et al. (2024). Dynamic Bayesian network for risk assessment of autonomous ships. *Accident Analysis and Prevention*, *194*, 107342.

Krot, P., Iskierka, G., Poskart, B., & Gola, A. (2022). Predictive maintenance of mobile robots based on IIoT and AI technologies. *Materials*, *15*(19), 6561. doi: 10.3390/ma15196561

Li, Y., Li, Y., Feng, K., Gryllias, K., Gu, F., & Pecht, M. (2024). Small-data prognostics and health management: A survey. *Artificial Intelligence Review*, *57*(8), 214. doi: 10.1007/s10462-024-10820-4

Luo, H., Li, Y., Bai, W., Tang, H., & Jin, H. (2024). A review of reliability engineering approaches for autonomous robotic systems. *Proceedings of the Institution of Mechanical Engineers, Part O: Journal of Risk and Reliability*. doi: 10.1177/1748006X231171449

Mikołajczyk, T. (2023). Reliability analysis and failure mode identification in mobile robotic systems. *Sensors*, *23*(15), 6890.

Pan, H., et al. (2025). Deep convolutional neural network-based fault diagnosis for industrial manipulators. *Journal of Field Robotics*, *42*(7), 3308–3322.

Peeters, J. F. W., Basten, R. J. I., & Tinga, T. (2018). Improving failure analysis efficiency by combining FTA and FMEA in a recursive manner. *Reliability Engineering and System Safety*, *172*, 36–44. doi: 10.1016/j.ress.2017.11.024

Sabry, A. H., & Amirulddin, U. A. U. (2024). A review on fault detection and diagnosis of industrial robots and multi-axis machines. *Results in Engineering*, *21*, 101888.

Shafiee, M., Enjema, E., & Kolios, A. (2019). An integrated FTA-FMEA model for risk analysis of engineering systems: A case study of subsea blowout preventers. *Applied Sciences*, *9*(6), 1192. doi: 10.3390/app9061192

Soleimani, M., Shahbeigi, V., & Esfahani, M. N. (2024). A methodology for reliability assessment using Bayesian networks. *Mechanical Systems and Signal Processing*, *216*, 111459. doi: 10.1016/j.ymssp.2024.111459

Yazdi, M. (2023). A bibliometric review of fault tree analysis methods and applications. *Quality and Reliability Engineering International*, *39*, 1234–1255. doi: 10.1002/qre.3271

Zhao, W., et al. (2024). Battery management system diagnostics for mobile robotic platforms. *Journal of Energy Storage*, *84*, 110826.