Where and When: How Continuous Lift Monitoring Answers the Two Questions That Drive Service Efficiency

Polina Rukin, Christoph Ruediger

CEDES AG, Kantonsstrasse 14, 7302 Landquart, Switzerland

Keywords: lift monitoring, IoT, adaptive baseline, anomaly detection, predictive maintenance, sensor fusion, fault detection

Abstract. Effective lift service has always depended on two questions: where is the problem, and when did it start? In the absence of continuous monitoring, answers to both questions arrive late, filtered through tenant callbacks, manual inspection logs, and technician recall. This paper examines how a controller-agnostic IoT platform, deployed across a heterogeneous fleet of commercial, transit, and healthcare lift installations, changes the answer to both questions in practice. Central to the approach is the absence of modelled baselines. Instead, two algorithms based on isolation forests and adaptive windowing are presented which detect anomalies by different methodologies to cover different kinds of anomalies. This approach is contrasted with fixed global threshold alerting, which the authors implemented and subsequently replaced following significant false positive rates. Operational case studies illustrate the practical outcomes. The paper discusses the known limitations of adaptive baseline methods, including the initialisation period required before reliable detection is possible, and the latency in alert resolution once a fault has been cleared.

1 INTRODUCTION

The lift service industry operates under a fundamental information problem. A fault develops gradually, often over days or weeks, but becomes visible to the service organisation only when a passenger reports it, typically by telephone, after the lift has already stopped or degraded to the point of inconvenience. By that point the fault has progressed further than necessary, the response is reactive rather than planned, and the dispatched technician may arrive to find the lift running again, having recovered on its own.

Estimates of fault-driven callout volumes illustrate the scale of this problem. Customer interviews concerning a 500-lift fleet operated by an independent service provider indicated approximately 1,000 callback visits annually. Approximately 47% of these were attributable to temporary shutdowns from which the lift had already recovered by the time a technician was dispatched, representing avoidable trips that could be eliminated with real-time status visibility. A further 12% were attributable to environmental causes addressable through physical monitoring of the shaft environment. The combined cost of avoidable dispatch alone was estimated at a minimum of USD 176,000 per year for a fleet of this size, based on an average callout duration of 1.5 hours at USD 250 per hour.

The technical precondition for addressing this problem is continuous, structured data from each lift unit. Various approaches have been explored in the literature. Accelerometer-based logging without controller access, as demonstrated by Peters and Peters [1], provides kinematics data including trip distance, velocity, and floor position. Sensor fusion approaches combining optical, accelerometric, and environmental sensors have been applied to door fault detection [2]. Machine learning methods trained on vibration signatures have demonstrated high classification accuracy for specific fault classes under laboratory conditions [3] [4]. Simulation-based approaches have been used to generate synthetic fault data for training condition monitoring models [5].

The laboratory approaches, while technically rigorous, share a common constraint: they require labelled examples of fault conditions in advance. In a real heterogeneous fleet, this constraint is significant. Fault events are rare relative to normal operation; labelled fault data is therefore scarce. More importantly, what constitutes normal behaviour varies substantially between units. A door closing time of 3.4 seconds may be unremarkable for one lift and indicative of developing friction in another. A gap in rides of 40 minutes may be expected overnight and anomalous at midday. Fixed global thresholds cannot accommodate this variability without either generating excessive false positives or missing genuine degradation.

This paper describes a different approach. Rather than training a model on labelled fault examples, the system constructs a per-unit baseline from each lift's own operational history and detects deviations from that baseline using an adaptive windowing method. No controller access is required. The method is applied across three signal domains: ride frequency (for shutdown detection), door behaviour (for door health and damage detection), and ride quality (for mechanical degradation detection). The approach is grounded in the adaptive windowing framework of Bifet and Gavaldà [6] and the barometric floor-level positioning methods described by Li et al. [7].

The paper is structured as follows. Section 2 describes the sensor architecture. Section 3 describes the isolation forest and the adaptive detection methods. Section 4 contrasts the adaptive approach with the fixed threshold approach previously in use. Section 5 presents two operational case studies. Section 6 discusses limitations and directions for further development. Section 7 concludes.

2 SYSTEM ARCHITECTURE

The monitoring platform described in this paper is hardware-agnostic. It operates alongside the lift controller without modification to the certified safety chain, consistent with the architectural principles described by Smans [8]. Data collection is performed by two primary sensing elements installed on the lift cabin: a light curtain and a gateway device. The light curtain is a safety sensor already present on most modern lifts for door obstruction detection. In the monitoring context, its beam interruption pattern provides a continuous record of door open and close events, passenger presence inference, and door cycle counts. The gateway device contains an accelerometer, a barometric pressure sensor, and a temperature and humidity sensor, and communicates via LTE or Ethernet to a cloud analytics platform.

Floor position is determined from barometric pressure using the method described by Li et al. [7], which exploits the measurable pressure differential between floors in enclosed building shafts. This approach provides floor-level resolution without requiring encoder access or controller integration. The accelerometer provides vertical acceleration data from which ride kinematics are derived: trip direction, peak speed, jerk, and vibration intensity per floor, following the analytical framework of Peters and Peters [1]. The light curtain provides door cycle timing, panel-level vibration, obstacle counts, and a direct measure of whether each door cycle completed its full range of travel.

Together, these three sensing elements – the light curtain, accelerometer, and barometer – produce a continuous structured data stream characterising each lift unit across three domains: what the cabin is doing (ride kinematics), what the doors are doing (door events), and where the cabin is (floor position). All data is timestamped at the gateway and transmitted to the cloud platform where baseline

construction and anomaly detection are performed. No proprietary controller protocol is required at any stage.

3 ANOMALY DETECTION

In this section, we are presenting two approaches: isolation forests as an unspecified anomaly detection means, and adaptive windowing as a means of detecting specific anomalies.

3.1 Adaptive windowing for anomaly detection

Windowed anomaly detection is a tool for detecting ad hoc changes on single-dimensional time series, as they often appear when things break or spring out of adjustment. Windowed anomaly detection does not need any baselines.

This feature is derived from the ADWIN algorithm [6] and works by comparing the left and right sides of a window via the Mean Absolute Error (MAE) method. In case of an anomaly within the window, the MAE becomes quite large as the mean is no longer a good approximation of the data in the window. When slicing the window into a right and left side, the MAE for both windows becomes minimal when the slicing happens exactly at the location of the anomaly. This causes the window to split into two sides, with each being stable in its own regard.

Given a window W of size n with t_i being the timeseries data at index i as

$$W_n = \{t_i | \forall i \in \mathbb{N}\ with\ 0 \leq i < n\}. \tag{1}$$

Splitting the window at position $k < n$ into a left and a right side reveals

$$\begin{aligned} W_{n,k} &= W_{left,n,k} \cup W_{right,n,k} \\ &= \{t_i | \forall i < k\} \cup \{t_i | \forall i \geq k\} \end{aligned} \tag{2}$$

and subsequently the MAE of window size n at splitting point k as

$$MAE_{n,k} = \left\{ \left| x_i - \underset{\forall i}{\text{mean}}(x_i) \right| \mid \forall i, 0 \leq i < k \right\} \cup \left\{ \left| x_j - \underset{\forall j}{\text{mean}}(x_j) \right| \mid \forall j, k \leq j < n \right\}. \tag{3}$$

Finally, the MAE of the entire window comes to

$$MAE_n = \left\{ \text{mean}\left(MAE_{n,k}\right) \middle| \forall k, 0 \leq k < n \right\} \tag{4}$$

The range between the minimum and maximum MAE within a window is a good measure for whether an anomaly is present within the window. Depending on the type of data that is analysed, both static and dynamic (data-driven) thresholds show good results. The position k with the minimal MAE is the temporal position of the anomaly.

This method just looks at data cohesion within a given window. A baseline of any sort is not required.

The key property of the modified ADWIN algorithm relevant to this application is its mathematical guarantee on false positive and false negative rates. The algorithm detects change when two sufficiently large sub-windows of the current window exhibit means that are statistically distinct

beyond a configurable bound. This provides a principled basis for the alert threshold that is specific to each unit's own variance, rather than a global engineering judgement.

3.2 Isolation forest

An isolation forest [9] is a model-based classification method that can isolate anomalies in a multi-dimensional dataset. It has the advantage that one does not need to model normal behaviour upfront, which is not known on a per-lift basis. The isolation forest is used to answer the question: is the most recently observed behaviour different to the behaviour observed throughout the last weeks?

Once this question gets answered positively for a point in time, a detailed statistical analysis on a per-lift basis is carried out to determine the relevance of the anomaly and determine the offending lift component. This was modelled by expert knowledge to look for certain behavioural anomalies to identify the most likely faulty component and the location of the fault.

Both the isolation forest and the statistical analysis have problems which make them hard to use in isolation. The statistical analysis is hard to parameterise between false positives and false negatives. In contrast, the isolation forest lacks explainability. With SHAP, the isolation forest can reveal the most dominant dimensions in the anomaly. However, this still does not reveal the concrete problem in terms of lift domain knowledge.

The combination of both methods has shown good results for detecting door failures. In a fleet of 178 lifts observed over four months, 1.7 million tumbling windows of 15 minutes each have been analysed. In that dataset, the isolation forest reported 1.24% of them as outliers. The statistical analysis was loosely trimmed and reported 1.7% of the bins as anomalies. The intersection of these two sets left only 264 windows which were flagged by both algorithms. These windows were then combined into consecutive events and alerted to customers as door behaviour anomalies. The backing of the expert-modelled statistical analysis made the alerts understandable to the lift mechanics.

4 CONTRAST WITH FIXED THRESHOLD ALERTING

Fixed thresholds are inherently limited because they cannot adequately capture either differences between lifts or changes in the behaviour of the same lift over time. For this reason, an adaptive thresholding approach was adopted.

An earlier implementation relied on global static thresholds, whereby alerts were triggered when a metric exceeded a predefined limit (for example, excessive door cycling between rides or vibration above a fixed amplitude). In practice, this approach produced a high false-positive rate, with the system generating up to 30 alerts per day for fewer than a hundred lifts. Most alerts reflected normal operational variability rather than emerging faults, reducing user confidence and limiting practical value.

Following implementation of adaptive baseline detection, alert frequency fell to approximately one alert per month and lift. This reduction reflected improved specificity rather than reduced sensitivity. As shown in section 5, genuine fault conditions were detected earlier and with greater precision under the adaptive approach.

5 FIELD OBSERVATIONS

Two field observations are presented, each illustrating a different signal domain and maintenance outcome. Lift manufacturers are not revealed, and customer information has been anonymised.

5.1 Door component degradation: roller off-track detection

The first case concerns a commercial lift installation monitored over a five-month period from August to December 2025. The door health analysis for this unit is shown in Figure 1, which presents two metrics tracked at daily resolution. The dashed blue vertical line marks the date on which a door health alert was detected by the system (7 September 2025). The dashed red vertical line marks the date which the alert lists as the behaviour change date (3 September 2025). The four days in between were used by the system to gain confidence on whether this is a permanent change.

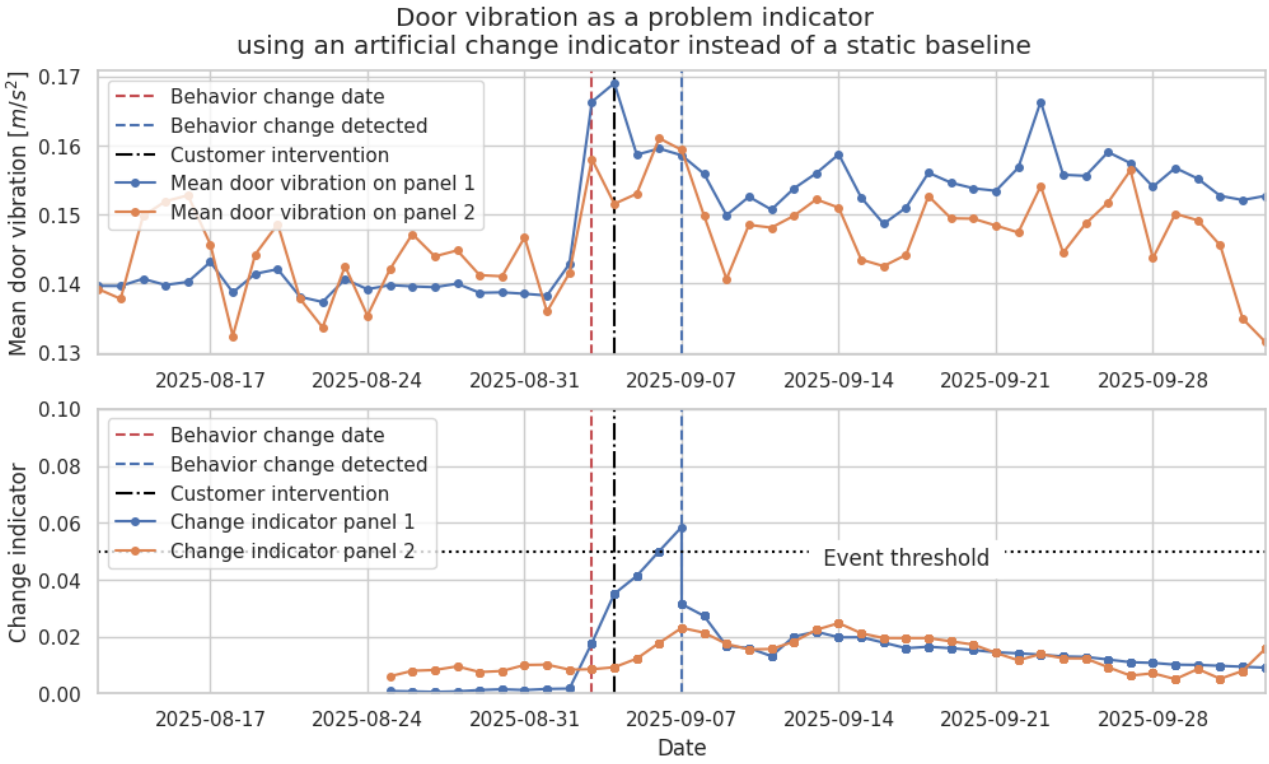


Figure 1 Multi-signal door health analysis for a commercial installation.

The dashed red line marks the detected behaviour change date. The dash-dotted black line marks the unrelated replacement of the roller. The blue dashed line marks the point in time when the anomaly was detected. The bottom plot shows the increase of the change indicator, easily understood as the system building confidence. Following roller replacement (black dashed line), the vibration signal stabilises at a new operating level while the adaptive change indicator returns below the detection threshold after sufficient post-repair data have been accumulated.

Based on our alert, a technician inspected the lift on 4 September 2025 and found that one of the door rollers had moved off track following an unreported impact on the door panel. The roller was replaced, preventing further wear progression and avoiding a potential door jam that would require emergency attendance.

The upper graph in Figure 1 shows that the mean door vibration changes significantly after the maintenance intervention. In particular, both door panels exhibit a new stable vibration level with increased fluctuations compared to the pre-fault state. Although the repaired door no longer behaves identically to the pre-fault condition, it settles into a consistent operating regime characteristic of the new roller. This change in normal operating behaviour demonstrates why fixed baselines would require manual recalibration after maintenance.

The lower graph illustrates how the adaptive windowing algorithm responds to this change. The change indicator rises as evidence for the behavioural shift accumulates, triggering the anomaly alert. Following the roller replacement, the indicator gradually decreases below the event threshold as sufficient post-repair observations become available and the algorithm establishes the repaired behaviour as the new normal.

This case illustrates the core advantage of trend detection over threshold detection. The key was an early shift in the statistical distribution of vibration measurements, not a violation of an absolute limit. The adaptive windowing captured the shift; a fixed threshold for all lifts would not have.

5.2 Door damage detection: Multiple anomalies detected with one algorithm

With the combination of isolation forest and statistical post-processing, multiple different kinds of anomalies can be detected. Figure 2 illustrates how the two-stage detection process identifies operational anomalies over several consecutive days.

Each data point in the upper graph represents the anomaly score computed for a 15-minute tumbling window. Windows identified as anomalous by the isolation forest (brown dots) are subsequently evaluated using the expert-defined statistical analysis described in Section 3.2. The shown graphs of ride and door movement counts, and the number of obstacles in the light curtain, are the main input factors for the statistical post-processing. Only windows classified as anomalous by both methods are retained. Consecutive anomalous windows are merged into a single event, which is displayed as a green shaded region. The width of each green region therefore represents the duration of the detected abnormal behaviour rather than a single timestamp.

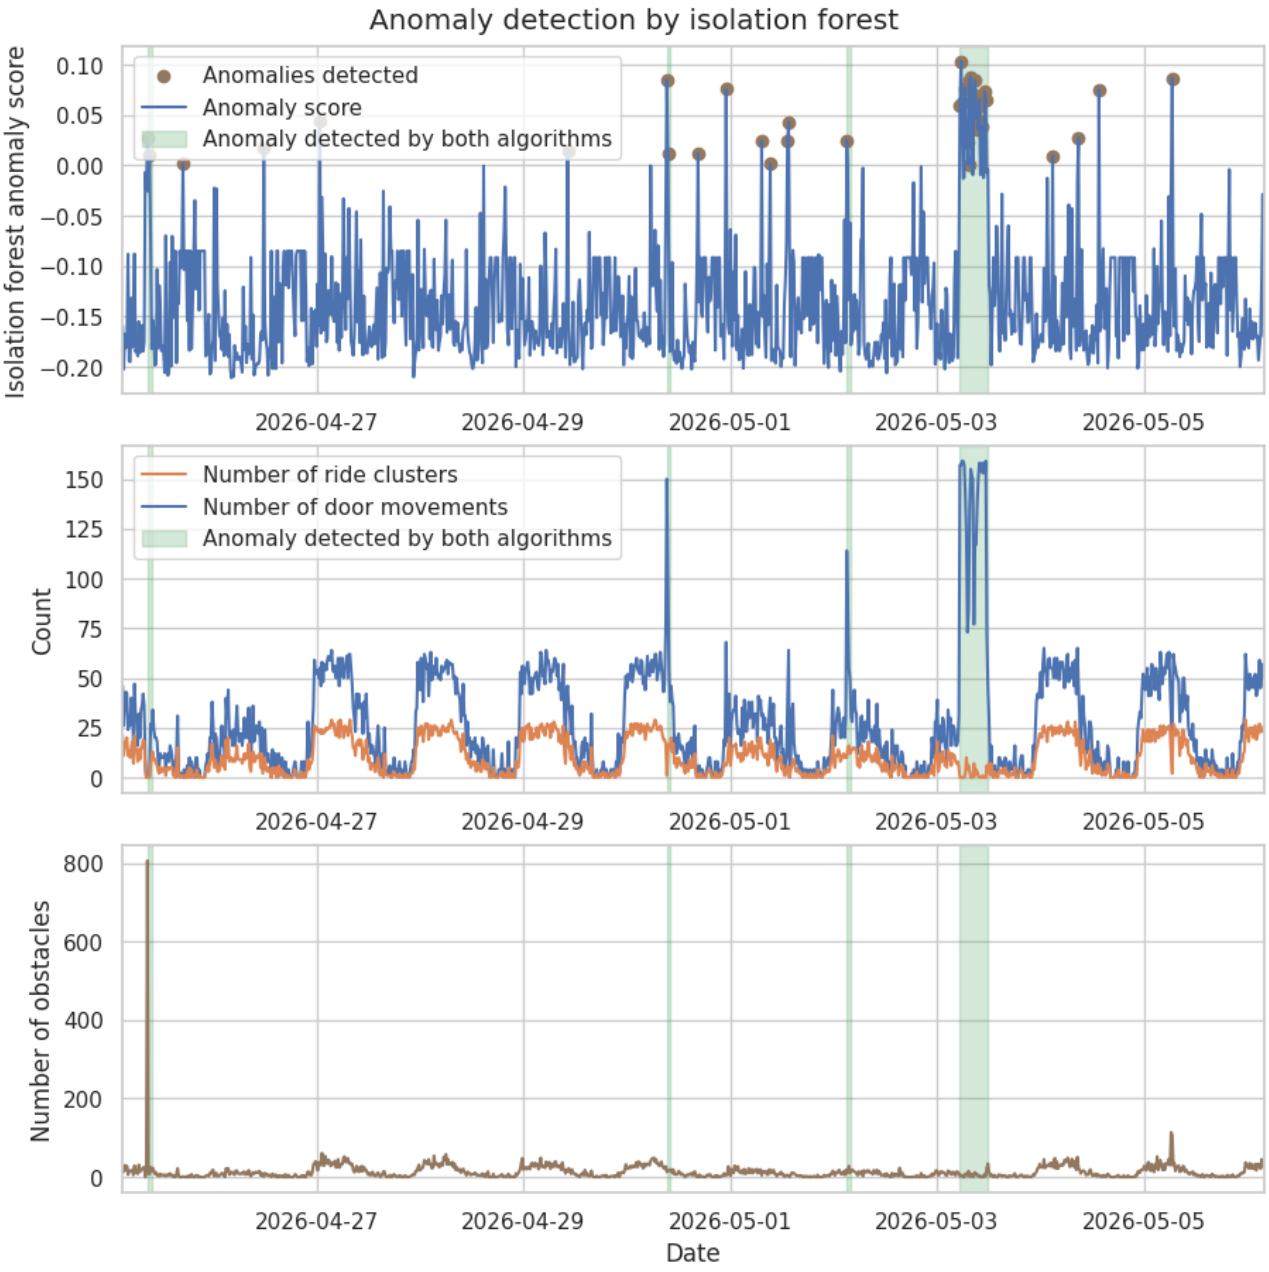


Figure 2 Results of the described anomaly detector

The top graph shows the isolation forest anomaly score, with brown markers indicating windows classified as anomalous by the isolation forest alone. The middle graph shows the number of door movement clusters (blue) and ride clusters detected by the light curtain (orange). The bottom graph shows the number of obstacle detections recorded by the light curtain. Green shaded regions indicate anomaly events triggered by the isolation forest and confirmed by the statistical post-processing algorithm.

The first detected anomaly was a blocked door. The number of obstacles in the light curtain jumps up, followed by no further door or car movements. The simultaneous increase in obstacle counts and absence of subsequent activity indicates that the door failed to complete normal operation, resulting in a sustained interruption of service.

The latter three anomalies are different kinds of repeated door cycling without rides (second anomaly, 30 April) and with lift rides (anomalies on 2 and 3 May). In these cases, repeated door movements occur without the expected correspondence between passenger activity, obstacle detections and car movement, producing an abnormal multi-sensor pattern that is recognised by both the isolation forest and the statistical post-processing.

Also, the slightly unusual behaviour on 24 April was not flagged, as this constitutes only an above-average usage period where all monitored signals remain mutually consistent. Although passenger traffic increased, the relationships between door movements, passenger clusters and obstacle counts remained consistent with normal operation. Consequently, the statistical analysis rejected the isolation forest output, preventing a false positive.

6 DISCUSSION

6.1 Relationship to laboratory-based approaches

The isolation forest approach described in this paper addresses a different problem from the supervised machine learning methods applied in laboratory settings by Kaczmarczyk et al. [3] [4] and Lin and Roivainen [5]. Laboratory approaches achieve high accuracy in fault classification because they control the conditions: a known fault is introduced into a test rig, labelled data is collected, and a classifier is trained to distinguish fault classes from normal operation. The resulting models can identify specific fault types, such as inner bearing ring damage or door lock mechanism failure, with high precision.

The operational constraint is that labelled fault data at scale is difficult to obtain in real fleet deployments. Faults are rare relative to normal operation. The specific failure modes that occur in the field may not match those represented in the laboratory dataset. And the variability of real-world installations, in terms of manufacturer, age, configuration, and operating environment, limits the transferability of models trained on a specific test rig.

The isolation forest approach, on the other hand, is an unsupervised methodology and detects that something has changed from normal for a specific lift, localises the change to a component and floor, and timestamps it. This is complementary to, rather than in competition with, fault classification approaches. In a mature monitoring architecture, baseline-based anomaly detection could serve as a first-stage filter that flags candidate events for subsequent classification, with the classification stage applied only to the small subset of events that cross the detection threshold.

6.2 Known limitations

Three limitations of the current approach warrant explicit discussion.

The initialisation period means that the system cannot provide anomaly detection immediately upon installation. The length of the training period depends on the regularity of the unit's operating pattern. Service organisations should be informed of this constraint at the point of onboarding.

Alert resolution latency is a second limitation. The adaptive window shrinks when a change is detected, but grows again only as new stationary data accumulates. When a fault is repaired and the monitored signal returns to its baseline level, the system may continue to carry an active alert for some time while the post-repair window accumulates sufficient observations to confirm stability. This

can be mitigated by manual alert clearance at the time of the technician visit, which requires integration between the monitoring platform and the service organisation's workflow.

The third limitation is that anomaly detection identifies change but does not identify the specific cause. The alert directs a technician to the right location and the right signal domain, but root cause determination remains a human task. This is appropriate for a safety-critical industry where maintenance decisions carry liability implications, but it means the system functions as an aid to technician judgement rather than a replacement for it.

6.3 Directions for further development

Several directions for further development are under active consideration. Further work on distinguishing vibration sources for specific components can lead to better characterisation of the malfunction. Maintenance event detection, the ability to identify from sensor data that a maintenance visit has occurred, would allow the system to automatically reset baselines following maintenance and reduce alert resolution latency. Multi-unit comparison, using aggregated fleet data to identify units whose baseline behaviour is systematically different from similar units in similar buildings, could surface latent issues that per-unit detection would miss. Integration with field service management systems would close the loop between alert generation and technician dispatch, enabling automated work order creation and structured feedback on alert accuracy.

7 CONCLUSION

This paper has described a controller-agnostic IoT monitoring architecture for lift installations that constructs per-unit adaptive baselines across three signal domains and detects anomalies as deviations from those baselines using an adaptive windowing method. The approach was motivated by the operational failure of fixed global threshold alerting, which generated up to 30 false positive alerts per day before being replaced, and was subsequently reduced to approximately one alert per month following the transition to adaptive baselines.

Three operational case studies demonstrate practical outcomes: early detection of door roller degradation through vibration trend analysis two weeks before absolute thresholds would have been crossed; characterisation of a door damage and incipient shutdown event through the fusion of independent barometric and light curtain signals; and detection of three shutdown events on a high-traffic unit, with the second shutdown independently corroborated by zero accelerometer activity across two separate light curtain sensors. A fourth operational outcome, run-on-arrival avoidance, is described as the most consistently quantifiable form of value for independent service providers, with an estimated minimum saving of USD 176,000 per year on a fleet of 500 units.

Three limitations are discussed: the initialisation period before reliable detection is possible, alert resolution latency following fault clearance, and the absence of root cause identification. Each is a known and characterizable property of the adaptive windowing approach rather than a fundamental constraint of the sensor architecture.

The isolation forest approach does not replace laboratory-based fault classification methods. It addresses the complementary problem of detecting that a change has occurred in a real heterogeneous fleet, without requiring labelled fault data or controller access. In doing so, it answers, more reliably

and more specifically than fixed threshold alerting, the two questions that drive lift service efficiency: where is the problem, and when did it start.

The authors would like to thank Guanghan Song and Matt Fabina for their support.

8 REFERENCES

[1] A. Peters and R. Peters, "Logging and Analysis of Lift Journeys Using an Accelerometer," in *Proceedings of the 10th Symposium on Lift and Escalator Technologies*, Northampton, 2019.

[2] D. Smans, "Modelling Failure Detection of Lift Doors Through Sensor Fusion," in *Proceedings of the 15th Symposium on Lift and Escalator Technologies*, Kettering, 2024.

[3] S. Kaczmarczyk, R. Smith and M. Gizicki, "Vibration Signature and the Application of Intelligent Pattern Recognition in Detection and Classification of Damage in Automatic Power Operated Lift Doors," in *Proceedings of the 13th Symposium on Lift and Escalator Technologies*, Northampton, 2022.

[4] M. Gizicki, S. Kaczmarczyk and R. Smith, "Classification and Recognition of Roller Bearing Damage in Lift Installations Using Supervised Machine Learning and Vibration Analysis," in *Proceedings of the 14th Symposium on Lift and Escalator Technologies*, Northampton, 2023.

[5] M. Lin und G. Roivainen, „Dynamic Simulations for Lift Door Health Diagnosis," in *Proceedings of the 13th Symposium on Lift & Escalator Technologies*, Northampton, 2022.

[6] A. Bifet und R. Gavaldà, „Learning from Time-Changing Data with Adaptive Windowing," in *Proceedings of the 2007 SIAM International Conference on Data Mining*, Minneapolis, 2007.

[7] B. Li, B. Harvey and T. Gallagher, "Using barometers to determine the height for indoor positioning," in *International Conference on Indoor Positioning and Indoor Navigation*, Montbeliard-Belfort, 2013.

[8] D. Smans, "Enabling innovation through an Internet of Things (IoT) platform that operates alongside the lift controller without modifying the certified safety chain," in *Proceedings of 16th Symposium on Lift & Escalator Technologies*, Kettering, 2025.

[9] F. T. Liu, K. M. Ting und Z.-H. Zhou, „Isolation Forest," in *2008 Eighth IEEE International Conference on Data Mining*, Pisa, 2008.

[10] T. Ehrl, R. Smith and S. Kaczmarczyk, "Key Dynamic Parameters that Influence Ride Quality of Passenger Transportation Systems," in *Proceedings of the 5th Symposium on Lift and Escalator Technologies*, Northampton, 2015.

9 BIOGRAPHICAL DETAILS



Polina Rukin is a product owner with long-term experience in design and user experience at CEDES. She started her tenure in the lift industry at Relayr when working for the Skyler Elevate product.



Christoph Rüdiger currently occupies the chief engineer role of Elevate at CEDES. He started his career over 15 years ago at TKE in the lift controller R&D and has since worked in different positions in TKE's MAX. He gained extensive knowledge of modern lift emergency call systems when working as the technical lead in TKE's Berlin-based call centre for the DACH region.