Journal of Applied Data Sciences
Vol. 7, No. 3, September 2026, pp. 1896-1913
ISSN 2723-6471
1896

An Interpretable Composite Index for Real-Time RFID Anomaly Detection in Predictive Maintenance

Zhang Shude[1,*], Adnan Yahaya[2]

[1,2]University Malaysia of Computer Science & Engineering (UNIMY), Level 1 & 2, VSQ@PJ City Centre, Jln Utara, Section 14, 46200 Petaling Jaya, Selangor, Malaysia

(Received: February 10, 2026; Revised: April 5, 2026; Accepted: June 15, 2026; Available online: July 12, 2026)

Abstract

Radio-frequency identification (RFID) systems are widely deployed for asset and inventory tracking in industrial environments, yet their reliability degrades under dynamic conditions where early, subtle tag anomalies remain difficult to detect. The objective of this work is to develop an interpretable, lightweight, and real-time index for detecting anomalous tag behavior at the point of reading, addressing the limitation that most existing methods rely on a single indicator such as received signal strength or raw read counts and therefore lack sensitivity to incipient instability. The core idea is to fuse two dimensionally consistent sub-metrics into a single composite score: a read-speed deviation coefficient that quantifies instability in tag-read cadence, and a communication-frequency concealment coefficient that captures temporal and communication-rate irregularities, with each component mapping to a physically identifiable failure mode so the score is interpretable by construction. The contribution is this anomaly evaluation index together with a single-pass, constant per-event computation suitable for commodity reader hardware. Validation spanned three tiers: a simulated dataset of 500 virtual tags and 30,284 events with controlled anomaly injection, a semi-synthetic dataset built from industrial warehouse logs, and a public benchmark of approximately 1,100,000 reads across multiple tag manufacturers. Key results: on the simulated test set the index achieved 95.7% accuracy, an AUC of 0.924, a 6.0% false-positive rate, and a mean detection latency of 4.8 ms on an embedded ARM-class processor; ablation confirmed the complementary contribution of both sub-metrics (accuracy dropped when either was removed), and the index ran approximately 40× faster than a deep-learning baseline of comparable detection quality. The novelty lies in combining two dimensionally consistent, physically interpretable sub-metrics into one constant-cost score, delivering deep-learning-comparable accuracy with a 40× speedup, making it well suited to embedded, real-time anomaly detection in resource-constrained industrial deployments.

Keywords: Anomaly Evaluation Index, Edge Computing, Embedded Deployment, Predictive Maintenance, RFID Anomaly Detection

1. Introduction

RFID systems are rapidly becoming foundational technologies in Industry 4.0, permeating manufacturing, logistics, healthcare, and asset-tracking domains [1], [2], [3], where they support real-time information tracking infrastructure for extended enterprises [4] and form the data substrate for modern RFID analytics pipelines [5]. Their key capability is automatic, wireless, non-line-of-sight identification of tagged items, enabling fine-grained, real-time visibility across operational flows [6], [7]. Despite this promising potential, RFID reliability in dynamic industrial environments remains a persistent challenge [8], and broader resilience considerations in Internet of Things (IoT) systems require explicit mechanisms to detect, tolerate, and recover from operational faults [9]. For example, read performance degrades due to electromagnetic interference (EMI) [10], tag wear [11], orientation shifts [12], multi-path effects [13], synchronization mismatches [14], and hidden communication errors such as missed reads or partial tag collisions [15]. Studies in industrial tool-management settings confirm that tag placement and mechanical stress are significant contributors to degradation over time [11], a problem compounded by the scale of modern RFID deployments [16].

Most existing anomaly-detection approaches for RFID rely on single-dimensional metrics such as read count drops, tag read-rate averages, or received signal strength indicator (RSSI) thresholds [17], [18]. These metrics may signal catastrophic failure (a tag going entirely unread) but are inadequate to detect emerging irregularities such as micro-movements, orientation drift, partial occlusion, or communication-frequency anomalies that hide behind apparently normal read counts [19], [20]. Temporal management of RFID data [20] and RFID-specific anomaly detection [19] represent related research areas, but existing approaches either require cloud-side compute, lack interpretability, or

*Corresponding author: Zhang Shude (P09210003@student.unimy.edu.my)
DOI: https://doi.org/10.47738/jads.v7i3.1406
This is an open access article under the CC-BY license (https://creativecommons.org/licenses/by/4.0/).
© Authors retain all copyrights

address only single-dimensional features. Therefore, there is a growing need for real-time, lightweight, interpretable anomaly-detection mechanisms specifically tailored to RFID systems that can operate at the reader or edge layer and capture subtle early-stage degradation [15], [19]. Such mechanisms should support predictive maintenance workflows: alerting when something is wrong and providing actionable diagnostic cues to schedule maintenance proactively [1], [8].

To address these challenges, this paper proposes the Anomaly Evaluation Index (AEI), a novel composite metric designed for reader-embedded, real-time deployment. AEI integrates two formally defined indicators. The Read Speed Deviation Coefficient (RSDC) is a normalized absolute deviation of instantaneous read speed from its windowed mean, scaled by windowed standard deviation, that quantifies temporal instability in tag-read cadence. The Communication Frequency Concealment Coefficient (CFCC) is a convex combination of temporal and frequency-domain deviation sub-terms that captures communication-layer anomalies which, in turn, suppressed or concealed tag responses that do not necessarily produce timing variance. The AEI is expressed as $w_1 \cdot RSDC + w_2 \cdot CFCC$ where the weights are calibrated through empirical regression on a held-out validation split, designed for single-pass O(W) computation suitable for microcontroller-class processors. In this study, interpretability is defined concretely as decomposability in the sense of Lipton [21] which is a form of structural transparency under which each component of the model carries an intuitive standalone meaning and the contribution of each to the final prediction is directly inspectable. AEI satisfies this criterion by construction: it is a closed-form weighted sum of two named, dimensionless sub-metrics, each of which maps to a physically identifiable failure mode (RSDC to read-cadence instability; CFCC to communication-layer concealment). This places AEI on the inherently-interpretable side of the distinction drawn by Rudin [22] between models that are interpretable ante hoc and black-box models for which interpretability must be reconstructed retrospectively. Unlike post-hoc attribution methods such as LIME [23] and SHAP [24], which approximate the behaviour of an opaque model by fitting a local surrogate or estimating Shapley-value feature contributions, AEI requires no external attribution step because the contribution of each component to the final score is directly readable from the weighted terms themselves. This component-level decomposability is the operational interpretability criterion used throughout the study.

The primary contribution of this work is an interpretable, edge-deployable composite index that offers a favourable, quantitatively characterised trade-off between detection quality, computational cost, and diagnostic transparency for RFID predictive maintenance. This is not, however, an assertion of universal superiority over all approaches on all criteria. The novelty claimed here is not the use of embedded software per se where many lightweight detectors can run in similar environments. Instead, it is about the specific design of two RFID-physics-motivated sub-metrics whose linear fusion yields a single, component-decomposable score. The trade-off itself is made explicit in Section 5.4 through a defined multi-objective utility function rather than asserted qualitatively. The focus of this paper is on the formulation and validation of the proposed AEI. In light of this, a high-level anomaly detection mechanism is described in order to place AEI in an operational context. In the next section, a conceptual framework for this mechanism is presented.

2. Conceptual Framework of the Anomaly Detection Mechanism

2.1. Overview of the Detection Mechanism

A high-level anomaly detection mechanism is considered to contextualize the proposed Anomaly Evaluation Index (AEI) [5], [6]. Simply stated, the objective of this mechanism is to transform raw RFID observations into a meaningful assessment of anomaly intensity in real time [19], [20]. Instead of constituting a standalone contribution, the mechanism provides a functional setting in which AEI operates as the central evaluation metric. The mechanism consists of a sequence of processing stages that progressively refine raw tag-reading data into a decision signal [5], [30]. These stages include data acquisition, indicator computation, composite index evaluation, and anomaly decision-making [19]. Within this pipeline, AEI serves as the key aggregation component, integrating multiple feature-level signals into a single interpretable measure of system stability [21], [22].

2.2. Conceptual Framework

The anomaly detection mechanism is abstracted as a conceptual framework that captures the lifecycle model of the detection process. The framework is designed to be sufficiently general to support different deployment scenarios while remaining aligned with the computational constraints of edge environments. Figure 1 illustrates the conceptual framework, which models the lifecycle of RFID anomaly detection. It comprises four interrelated functional modules, where each fulfills a distinct role in the anomaly detection process.

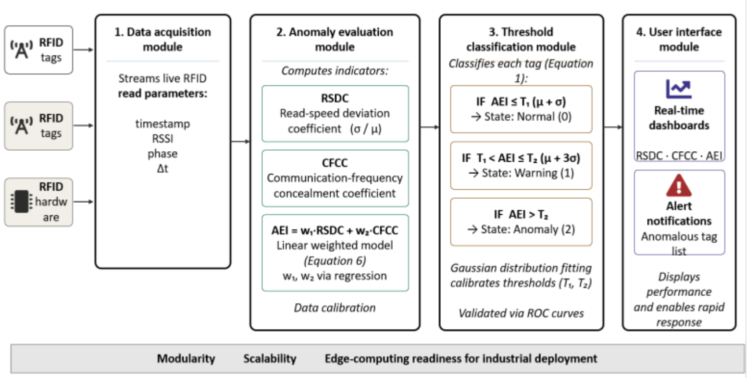

Figure 1. Conceptual Framework for RFID Anomaly Detection

First, the Data Acquisition Module serves as the primary interface with RFID hardware, continuously capturing live operational parameters such as read speed and communication frequency. This module ensures that the mechanism has access to the raw, real-time data necessary for effective monitoring.

Secondly, the Anomaly Evaluation Module processes this input data by computing two key indicators which are the RSDC and CFCC for each RFID tag. These values are then integrated into a single, composite metric, the Anomaly Evaluation Index (AEI). The AEI serves as a real-time indicator of each tag's operational health, computed using a weighted aggregation model that reflects the relative importance of each underlying metric.

Thirdly, the Threshold Classification Module compares each AEI value to a set of calibrated thresholds. Based on the outcome, each tag is categorized into one of three states: normal, warning, or anomalous. This immediate classification allows for timely intervention, especially in scenarios where system failures or misidentifications could result in financial loss or operational disruption. Finally, the User Interface Module presents the analyzed data in a visual and accessible manner. Through an intuitive dashboard, system operators and maintenance personnel can monitor tag performance, receive alerts, and track the history of anomalies or interventions. This interface not only enhances situational awareness but also supports informed decision-making and rapid response.

By adopting a modular structure, this architecture ensures scalability (ability to handle growing numbers of RFID tags), interoperability (integration with various hardware and software platforms), and maintainability (ease of upgrading or modifying system components). These attributes make the framework well suited for real-world industrial deployment, where operational flexibility and reliability are paramount.

2.3. Relationship to System Implementation

The conceptual framework provides a design-level abstraction of the anomaly detection mechanism, that is independent of specific implementation details. It defines the roles of individual components and their interactions, while leaving flexibility for realization across different hardware and software platforms.

In this study, the framework is instantiated as a system prototype that serves as an experimental setup to support the empirical calibration and evaluation of AEI under realistic conditions. The implementation details, including data processing pipelines and edge deployment considerations, are described in Section 4. It is important to note that the system serves as a validation vehicle, whereas the primary contribution of this paper lies in the formulation and validation of AEI.

3. Anomaly Evaluation Index (AEI) Formulation

3.1. Overview of the Method

At the system level, the four functional modules work as follows. The Data Acquisition Module streams live RFID read parameters (timestamp, RSSI, phase, Δt) in real time [5], [20]. The Anomaly Evaluation Module then computes the per-event RSDC and CFCC values, which are aggregated into the composite Anomaly Evaluation Index defined in Equation (1). The Threshold Classification Module classifies each tag into normal, warning, or anomaly states using the calibrated boundaries given in Equations (3) and (4) [6], [19]. Finally, the User Interface Module displays real-time dashboards and alert notifications [9]. Equation (1) defines the composite AEI as the windowed average of the per-event RSDC and CFCC values across n read events:

$$AEI = (\Sigma_{i=1}^{n} RSDC_i + \Sigma_{i=1}^{n} CFCC_i) / n \tag{1}$$

In Equation (1), $RSDC_i$ and $CFCC_i$ denote the temporal and communication-frequency deviation values at the i-th read event, while n is the number of events in the current evaluation window. Summing both metrics across the window and dividing by n produces a single, smoothed scalar that suppresses transient single-read spikes and provides a stable basis for downstream classification [19], [20]. To further attenuate short-lived fluctuations and enforce temporal stability before an alert is raised, the framework introduces an alarm score AS_t derived from the AEI through exponential smoothing, as expressed in Equation (2):

$$AS_t = \lambda \cdot AEI_t + (1 - \lambda) \cdot AS_{t-1} \tag{2}$$

In Equation (2), AEI_t is the current AEI value computed via Equation (1), AS_{t-1} is the alarm score from the previous time step, and $\lambda \in (0, 1)$ is the smoothing parameter (default $\lambda = 0.1$) that controls how quickly past observations are discounted. A smaller λ produces a smoother alarm trajectory that resists transient noise, while a larger λ makes the alarm score more responsive to abrupt changes [6], [8].

The architecture prioritizes modularity, scalability, and edge computing readiness for deployment in industrial IoT environments [2], [3], [4], [9]. To ensure reliability and precision in anomaly detection, this phase applies statistical analysis techniques to calibrate threshold values for the AEI [6], [19]. Historical and experimental AEI data are analyzed using Gaussian distribution fitting to approximate the underlying behavior of normal versus abnormal tags. From the training dataset, the mean (μ) and standard deviation (σ) of the AEI distribution are calculated and used to derive two decision boundaries, expressed in Equations (3) and (4):

The Gaussian fit is applied not to the raw per-event RSDC values but to the windowed AEI scores of the normal class, which are substantially smoothed: AEI averages RSDC and CFCC over a window of W events (Equation 1) and is then exponentially smoothed into the alarm score (Equation 2). This two-stage aggregation attenuates the heavy tail considerably, so the normal-class AEI distribution is far closer to Gaussian than the per-event RSDC distribution. The thresholds are therefore best understood as a parametric convenience that yields the 84th- and 97.5th-percentile interpretation only approximately. Where the normal-class AEI distribution deviates materially from Gaussianity, the ROC-derived thresholds which make no distributional assumption should be regarded as authoritative.

$$T_1 = \mu + \sigma \tag{3}$$

$$T_2 = \mu + 2\sigma \tag{4}$$

Equation (3) defines T_1 as the warning boundary, set one standard deviation above the mean of the normal-class AEI distribution. Under a Gaussian assumption, this corresponds to roughly the 84th percentile of normal behaviour, so AEI values exceeding T_1 are statistically uncommon under normal operation and warrant closer monitoring. Equation (4) defines T_2 as the anomaly boundary, set two standard deviations above the mean and corresponding to approximately the 97.5th percentile, beyond which an AEI value is sufficiently improbable under the normal distribution to be classified as a genuine anomaly and trigger an automated maintenance alert [6], [8]. This classification scheme segments tag behavior into three clear zones: normal where $AEI < T_1$, warning where $T_1 \leq AEI < T_2$, and anomaly/failure where $AEI \geq T_2$, as shown in (5):

$$\text{Normal: } AEI < T_1\text{ ; Warning: } T_1 \leq AEI < T_2\text{ ; Anomaly: } AEI \geq T_2 \tag{5}$$

This data-driven calibration increases model generalizability and reduces false positives [21], [22]. The first step in data analysis is the computation of the RSDC and CFCC for each tag across rolling time intervals [5], [19]. These custom-designed metrics serve as the primary anomaly indicators. RSDC is computed by calculating the standard deviation of read counts for a tag over a defined time window (e.g., 5 seconds), normalized by its moving average, which highlights abrupt or inconsistent changes in read speed often indicative of tag misalignment or environmental interference [10], [11], [12]. CFCC measures irregularities in time intervals between successive reads of the same tag; using a combination of time differencing and frequency mapping, CFCC detects communication gaps, suppressed reads, or deliberate concealment scenarios [14], [19], [20]. The computed metric values are then transformed into an AEI using a linear weighted model as in (6):

$$AEI = w_1 \cdot RSDC + w_2 \cdot CFCC \tag{6}$$

Weight parameters w_1 and w_2 are calibrated through empirical regression analysis based on the simulated and semi-synthetic datasets. The optimal thresholds T_1 (warning) and T_2 (anomaly) are determined using receiver operating characteristic (ROC) curves, optimizing sensitivity and specificity based on ground-truth labels in the training datasets.

3.2. Read Speed Deviation Coefficient (RSDC)

RSDC quantifies temporal instability in tag-read cadence as a normalized absolute deviation computed within a sliding window of W consecutive read events. Let r_t denote the instantaneous read speed (reads per second) at time t, and let $\mu_r(t)$ and $\sigma_r(t)$ denote the windowed sample mean and standard deviation of $\{r_{t-W+1}, \ldots, r_t\}$. RSDC is defined as in (7):

$$\mathrm{RSDC}_t = |r_t - \mu_r(t)| \,/\, (\sigma_r(t) + \varepsilon) \quad (7)$$

RSDC is a normalized absolute deviation from the window mean scaled by window standard deviation that is structurally equivalent to an absolute z-score. Although previously described as a coefficient of variation (CV), CV is conventionally defined as σ/μ and would not have the same interpretation here. The RSDC formula correctly uses σ (not μ) in the denominator, which provides stable normalization even when $\mu_r(t)$ is near zero (a common condition for high-quality tags with very consistent read intervals). The $\varepsilon = 10^{-6}$ offset prevents exact zero-division when all W reads are identical. Both numerator and denominator are in units of reads per second, making the ratio dimensionless, as required.

A high RSDC value indicates that the current read speed deviates substantially relative to the recent variability of that tag's read cadence, consistent with tag movement, orientation shifts, or partial signal obstruction [25]. An alternative robust normalization using median absolute deviation (MAD) instead of $\sigma_r(t)$ was evaluated.

3.3. Communication Frequency Concealment Coefficient (CFCC)

CFCC captures anomalies at the communication-frequency layer, where suppression or concealment of tag responses that do not necessarily manifest as timing variance. The hybrid use of temporal and frequency-domain features in CFCC is consistent with prior work in wireless sensing where joint temporal-frequency representations have been shown to capture activity-dependent signal irregularities more effectively than time-domain features alone [26]. CFCC combines two sub-terms via blending parameter $\alpha \in [0,1]$, as in (8):

$$\mathrm{CFCC}_t = \alpha \cdot \delta\_\tau(t) + (1 - \alpha) \cdot \delta\_f(t) \quad (8)$$

Sub-term 1: Temporal deviation δ_τ: normalized absolute deviation of the current inter-read interval $\tau_t = \Delta t_t$ from the windowed mean, scaled by the windowed standard deviation, as defined in (9):

$$\delta\_\tau(t) = |\tau_t - \mu\_\tau(t)| \,/\, (\sigma\_\tau(t) + \varepsilon) \quad (9)$$

Dimensional note: τ_t (interval, in seconds) and the windowed mean are in the same units. In the implementation, $\mu\_\tau(t)$ is the windowed mean of inter-read intervals in seconds and $\sigma\_\tau(t)$ is the windowed standard deviation of intervals. This preserves dimensional consistency: numerator and denominator are both in seconds, making δ_τ dimensionless.

Sub-term 2: Frequency deviation δ_f: relative deviation of the instantaneous communication frequency f_t from the per-tag baseline f_0, where f_t is defined as in (10):

$$f_t = 1 \,/\, (\Delta t_t + \varepsilon) \quad (10)$$

Baseline f_0 is computed as the median read frequency over the first B = 50 events per tag (warm-up phase). Using the median rather than the mean makes f_0 robust to startup transients. CFCC is the convex combination given in (11):

The term "frequency" in CFCC refers to communication frequency in the rate sense, specifically the number of successful tag responses per unit time, $f_t = 1/(\Delta t_t + \varepsilon)$ and not to a spectral (Fourier-domain) decomposition of the read signal. The δ_f term measures the relative departure of this instantaneous read rate from the per-tag baseline rate f_0, so CFCC is a rate-domain rather than a spectral-domain construct. The baseline rate is itself grounded in the interrogation protocol: under the EPC Class 1 Generation 2 air-interface standard [27], [28], a healthy passive tag responds in an Interrogator-Talks-First, slotted-Aloha pattern with inter-read intervals clustering around a nominal value τ_0, so $f_0 = 1/\tau_0$ is the protocol-defined expected communication rate against which δ_f is measured. This is a deliberate design choice for edge deployment: a true spectral analysis would require buffering a windowed segment and computing a transform per tag, which conflicts with the single-pass O(W) constraint imposed by per-event incremental statistics [29]. The rate-deviation formulation, by contrast, is computable incrementally using the same class of streaming updates that has been adopted for RFID middleware-level data cleaning [30]. The conceptual link to communication theory is that suppressed or concealed tag responses lower the effective channel utilisation, which manifests directly as a drop in f_t the same drop-out phenomenon that motivates adaptive-window RFID cleaning frameworks [30]; δ_f thus captures the first-order effect that a spectral treatment would also reflect, at a small fraction of the computational cost. We have accordingly avoided the phrase "frequency-domain" where it could imply Fourier analysis, and describe CFCC as combining a temporal-deviation term and a communication-rate-deviation term.

$$CFCC_t = \alpha \cdot \delta\_\tau(t) + (1 - \alpha) \cdot \delta\_f(t) \tag{11}$$

Hyperparameter $\alpha = 0.6$ was selected by grid search over {0.3, 0.4, 0.5, 0.6, 0.7, 0.8} on the held-out validation set, optimizing AUROC. Sensitivity analysis (Section 5.3) shows AUROC varies by less than 0.005 within ±0.1 of the optimal α. Edge cases: when Δt_t is unavailable (reader retuning, anti-collision back-off), τ_t is substituted by the window median Δt. CFCC is computed on a per-EPC basis in multi-tag environments, consistent with established efficient information-sampling practice for multi-category RFID systems [31].

3.4. Anomaly Evaluation Index (AEI)

RSDC and CFCC are fused via linear weighted aggregation, as in (12):

$$AEI = w_1 \cdot RSDC + w_2 \cdot CFCC, w_1 + w_2 = 1 \tag{12}$$

Fusion weights $w_1 = 0.6$, $w_2 = 0.4$ were selected by grid search over $w_1 \in \{0.3, 0.4, 0.5, 0.6, 0.7\}$ on the training set (F1 objective). The validation set (15% split) was used exclusively for weight selection and threshold calibration. All reported metrics use the held-out test set (15% split). The selection procedure is described here to support reproducibility. Because the weights are constrained by $w_1 + w_2 = 1$, the search reduces to a one-dimensional sweep over w_1 at a step of 0.1 across the stated grid. For each candidate value, the AEI was recomputed on the training split and the F1 score against the ground-truth labels was used as the selection objective, with the validation split reserved exclusively for this weight selection and for threshold calibration. The configuration $w_1 = 0.6$, $w_2 = 0.4$ was selected as the grid point with the highest training-split F1. A complementary feature-contribution analysis supports the resulting balance: across all observations, RSDC correlates with the composite AEI at $r = 0.992$ and CFCC at $r = 0.938$ (both $p < 0.001$), confirming that RSDC carries the larger share of the composite signal while CFCC contributes orthogonal diagnostic coverage and the orientation-shift cases in particular show elevated CFCC with only moderate RSDC.

Thresholds $T_1 = \mu\_val + \sigma\_val$ (warning boundary) and $T_2 = \mu\_val + 2 \cdot \sigma\_val$ (anomaly boundary) were estimated from the normal-class subset of the validation split via Gaussian fitting. ROC-based Youden J optimization agreed within 3%. These thresholds were frozen and then transferred without modification to Tier 2 (semi-synthetic) and Tier 3 (DUTR) evaluations to assess cross-dataset transferability.

3.5. Three-Tier Dataset Development Strategy

Three tiers provide controlled, semi-realistic, and stress-test coverage. The tier structure intentionally escalates the gap between simulation conditions and real-world deployment.

Tier 1: Simulated dataset: a Python-based RFID simulator generated 500 virtual tags and 30,284 read events with fully deterministic ground-truth labels. Anomalies (occlusion, orientation, rapid motion, EMI/silence) were injected at 30% of events with balanced sub-type proportions. Split: 70% train, 15% validation, 15% test (no leakage between splits). To make the partitioning fully explicit and reconcile it with the per-table event counts: the 70/15/15 split of 30,284 events yields 21,199 training events, 4,542 validation events, and 4,543 test events (the test count of 4,543 is the value reported in the Table 5 evaluation header; the one-event difference from an exact 15% arises from integer rounding of the split boundaries). The split is stratified by tag identity, so all events from a given virtual tag fall entirely within one partition, which is the source of the no-leakage guarantee. Table 2 reports n = 499 because that analysis is conducted at the tag level on the test partition (499 distinct tags) rather than at the event level.

The synthetic anomaly injection procedure is specified here to support reproducibility. Anomalies are introduced by a dedicated injection module that applies controlled, parameterised corruptions to a tag's event stream, with each anomaly characterised along two axes which are intensity (degree of disruption) and duration (length of the affected interval). Four mechanisms are used: (i) tag occlusion reduces or blocks reads for defined intervals, (ii) parameterised from minor signal attenuation through to complete read loss, (iii) modelling barriers such as metallic surfaces, liquids, or dense packaging. Hardware silence enforces artificial read gaps during which the tag produces no output: short silences on the order of one to two seconds model transient power fluctuations, while longer silences of five to ten seconds or more model more severe degradation. Orientation drift applies gradual shifts along the pitch, yaw, and roll axes at adjustable rate and degree of angular deviation, producing the progressive read-interval instability characteristic of misalignment. Electromagnetic interference injects burst noise into the RSSI and phase parameters with variable amplitude, duration, and frequency of occurrence, capturing both low-level background interference and sudden high-intensity disturbances. Underlying signal noise is modelled by perturbing RSSI with Gaussian noise of tunable mean and variance and adding random deviations to phase angles. In the semi-synthetic tier the same four mechanisms are injected at three discrete intensity levels, i.e. mild, moderate, and severe; where each with an explicitly defined behavioural signature, and every record carries a ground-truth label (normal, orientation, motion, or occlusion) together with its severity level. Each simulation run uses a different random seed so that no two runs are identical while the generation procedure remains fully reproducible.

Tier 2: Semi-synthetic dataset: real RFID logs from a warehouse laboratory environment using Impinj R700 readers (ARM Cortex-A8, 600 MHz, 256 MB RAM) and Raspberry Pi 4 Model B edge loggers (~200 physical UHF Gen-2 tags, three manufacturers, 8-hour sessions, 15,247 read events). Controlled anomalies were injected post-collection using the same scripts as Tier 1, preserving real signal behavior (hardware jitter, antenna patterns, RF multipath). This dataset is used for cross-hardware robustness evaluation, not primary accuracy benchmarking.

Tier 3: DUTR benchmark (proxy-label evaluation): the publicly available DUTR dataset [17] comprises 1,024 heterogeneous tags and 1.1M+ read events across multiple hardware configurations. Anomaly labels were generated by a proxy rule: read events with Δt exceeding per-tag median + 3·IQR, consistent with non-parametric temporal outlier detection [20]. Class balance: 91.2% normal, 8.8% anomaly.

Because the DUTR proxy labels are derived from inter-read interval statistics, and RSDC/CFCC are themselves built on interval-based features, there is a structural alignment between the labeling rule and the detection features. This creates a partial circularity: strong performance on DUTR reflects transferability of interval-based temporal features across hardware platforms, but does not constitute independent validation against true operational anomaly ground truth. Results on DUTR should be interpreted as a transferability stress-test rather than an independent empirical benchmark. Future work requires real-world annotated anomaly labels from instrumented industrial deployments to validate AEI against genuine ground truth. To avoid overstating this evidence, the present revision treats the DUTR results with deliberate caution throughout: they are not aggregated with the Tier 1 and Tier 2 numbers, they are not used to support any headline accuracy or AUROC claim in the abstract or conclusion, and any robustness statement drawn from DUTR is explicitly scoped to "transferability of interval-based temporal features" rather than to detection performance in general. The circularity is partial rather than total because the proxy rule (a per-tag median + 3·IQR threshold on Δt) is a coarse non-parametric outlier flag, whereas CFCC additionally incorporates a frequency-baseline deviation term and RSDC operates on read speed rather than raw intervals; the overlap inflates apparent performance but does not make the task trivial. Nonetheless, DUTR is positioned in this paper strictly as a supporting cross-hardware stress-test, and the primary evidentiary weight rests on Tier 1.

4. System Implementation Considerations

This section translates the formal definitions of RSDC, CFCC, and AEI from Section 3 into a concrete algorithmic specification suitable for reader-embedded deployment, in line with recent advances in lightweight signal processing and edge AI for real-time anomaly detection in IoT sensor networks [25]. Three implementation concerns are addressed in turn: numerical stability under low-variance conditions, computational and memory complexity on microcontroller-class hardware, and adaptive baseline tracking for non-stationary signal environments. Each concern is connected back to the equations that govern its behaviour, and the practical defaults adopted in the prototype are stated explicitly to support reproducibility.

4.1. Numerical Stability and the EWMA Fallback

A practical issue arises in the implementation of the RSDC defined in Equation (7). When all W observations within the current sliding window are identical or near-identical, the windowed standard deviation $\sigma_r(t)$ approaches zero, and the unstabilized ratio in the numerator of Equation (7) becomes numerically unbounded. Although the regularization constant $\varepsilon = 10^{-6}$ in the denominator prevents exact division by zero, it does not address the underlying issue: a near-zero $\sigma_r(t)$ produces an artificially inflated RSDC value whenever the current read speed deviates even slightly from the mean, generating spurious anomaly signals during low-variance intervals.

To address this, an Exponentially Weighted Moving Average (EWMA) variance estimator is adopted as a fallback when $\sigma_r(t)$ falls below a minimum variance threshold. The EWMA standard deviation is defined recursively in Equation (13):

$$\sigma_r^{EWMA}(t) = \sqrt{[\lambda \cdot (r_t - \mu_r(t-1))^2 + (1 - \lambda) \cdot \sigma_r^2(t-1)]} \tag{13}$$

In Equation (13), $\lambda \in (0, 1)$ is the smoothing parameter (default $\lambda = 0.1$), r_t is the current instantaneous read speed, $\mu_r(t-1)$ is the EWMA mean from the previous step, and $\sigma_r^2(t-1)$ is the squared EWMA standard deviation carried forward. The EWMA estimator is automatically activated when $\sigma_r(t) < \varepsilon$ in Equation (7), and the resulting $\sigma_r^{EWMA}(t)$ is substituted into the RSDC computation. This substitution preserves a meaningful normalization scale during low-variance intervals and continuously adjusts to slow drift in the underlying signal, eliminating the false-positive spikes that would otherwise occur when a stable tag experiences a brief, benign perturbation.

The same EWMA mechanism applies to the windowed standard deviation $\sigma\_\tau(t)$ of inter-read intervals used in the temporal sub-term of CFCC defined in Equation (9), with r_t replaced by $\tau_t = \Delta t_t$ throughout. Applying the EWMA fallback consistently across both metrics ensures dimensional and behavioral consistency in the composite AEI of Equation (12) under all signal conditions.

To verify that the EWMA fallback is necessary rather than merely defensive, the raw AEI time series was compared directly against the EWMA-stabilised AEI for a representative tag (TAG-010) over a full monitoring session. Without the fallback, the raw AEI exhibited numerous transient spikes during the pre-onset (normal) phase, with values intermittently approaching the warning threshold T_1 due to short-duration read-rate fluctuations precisely the spurious low-variance excursions that would generate false-positive alerts. With the EWMA fallback active, these transient spikes were substantially attenuated and the pre-onset AEI remained consistently below 0.070, comfortably clear of T_1. Critically, the stabilisation introduced no detection-latency penalty: the EWMA-smoothed AEI crossed the anomaly threshold T_2 within the same observation window as the raw AEI once the genuine anomaly began. The EWMA formulation additionally provides adaptive baseline tracking, allowing the detector to register secondary escalations relative to a shifted operating level rather than the original baseline. This confirms that the fallback contributes a real stability improvement concentrated in exactly the low-variance regime it was designed to address, without sacrificing responsiveness.

4.2. Computational Complexity and Memory Footprint

A second implementation concern is whether the formulations in Equations (7) to (12) can be evaluated within the latency and memory budget of an embedded RFID reader. The key observation is that all three metrics, namely RSDC, CFCC, and AEI, admit a single-pass recursive evaluation with O(W) memory and O(1) amortized compute per read event. This O(1) claim refers to the number of operations performed for each event, and it is important to clarify its practical scope. The prototype was implemented in Python using NumPy C extensions on the reader's embedded Linux operating system rather than as firmware on a constrained microcontroller. The asymptotic complexity remains constant because each event requires only a fixed number of additions, multiplications, and a single square root operation, with no dependence on the history length. However, while the operation count is platform independent, the execution speed and actual latency are not. On a microcontroller without optimized numerical libraries, the same computations would take longer to execute. Therefore, the 4.8 ms latency reported in Section 5.5 is specific to the NumPy-based embedded Linux environment used in this study. What remains consistent across platforms is the structural property that the per-event computational cost is constant and the memory requirement is bounded at O(W). In contrast, the exact wall clock execution time may vary depending on the hardware and software environment. A full firmware level implementation is considered future work. The windowed mean $\mu_r(t)$ and standard deviation $\sigma_r(t)$ used in Equations (7) and (9) can be maintained incrementally using the recursive update rules expressed in Equations (14) and (15):

$$\mu_r(t) = \mu_r(t-1) + (r_t - r_{t-}w) / W \tag{14}$$

$$\sigma_r^2(t) = \sigma_r^2(t-1) + [(r_t - \mu_r(t))^2 - (r_{t-}w - \mu_r(t-1))^2] / W \tag{15}$$

In Equation (14), the windowed mean is updated by adding the contribution of the newest read speed r_t and removing the contribution of the oldest read speed $r_{t-}w$, which leaves the window. In Equation (15), the windowed variance is updated by the analogous incremental rule, where the contribution of the new event replaces that of the expiring event in O(1) time. Together, Equations (14) and (15) eliminate the need to recompute statistics over the full window at every step, reducing per-event compute from O(W) to O(1) and bounding total memory at O(W) for the rolling buffer plus a small constant for the running statistics.

The CFCC sub-terms $\delta\_\tau(t)$ and $\delta\_f(t)$ in Equations (9) and (10) inherit the same incremental structure through identical recursive updates on τ_t. The convex combination in Equation (11) and the linear fusion in Equation (12) are scalar arithmetic operations and add negligible overhead. The end-to-end per-event compute is therefore dominated by a small constant number of floating-point multiplications and additions, which on the Impinj R700's ARM Cortex-A8 processor translates to the 4.8 ± 0.6 ms mean per-tag latency reported in Section 5.5. The corresponding memory footprint for $W = 10$ and the per-tag state shown in table 7 is 3.2 MB total at 200 active tags, well within the 256 MB available on the target reader hardware.

4.3. Adaptive Baseline Tracking and Drift Compensation

The third implementation concern is how the per-tag frequency baseline f_0 used in Equation (10) and the calibration thresholds T_1 and T_2 from Equations (3) and (4) remain valid as the deployed environment evolves. RFID systems exhibit several forms of drift that an interpretable, non-learning framework must accommodate explicitly: tag aging changes physical-layer behavior over months; environmental rearrangement of the monitored space alters multipath conditions; and seasonal variation in temperature and humidity modulates antenna characteristics. To support adaptive baseline tracking without retraining, the framework computes a rolling baseline $f_0^{adaptive}(t)$ using the EWMA recursion expressed in Equation (16):

$$f_0^{adaptive}(t) = \beta \cdot f_t^{normal} + (1 - \beta) \cdot f_0^{adaptive}(t-1) \tag{16}$$

In Equation (16), $\beta \in (0, 1)$ is the baseline learning rate (default $\beta = 0.05$, deliberately smaller than λ in Equation (13) to ensure the baseline evolves more slowly than the anomaly statistics it is being compared against), and f_t^{normal} is the instantaneous read frequency at time t, but only when the current alarm score AS_t from Equation (2) classifies the tag as normal. This guard condition updating f_0 only on normal-class observations prevents the baseline from drifting toward genuine anomalies and stabilizing on degraded behavior, which would silently desensitize the detector over time.

The threshold pair (T_1, T_2) defined in Equations (3) and (4) is recalibrated quarterly on a 1,000-event normal-class sample drawn from live deployment, subject to a guard condition requiring $RSDC < 0.05$ for at least 95% of events in the candidate calibration window. This periodic recalibration protocol preserves the statistical interpretation of T_1 and T_2 as the 84th and 97.5th percentiles of the normal-class AEI distribution under Gaussian assumptions, even as μ and σ in Equations (3) and (4) shift gradually with hardware aging or environmental change. The combination of fast adaptive baseline tracking via Equation (16) and slow periodic threshold recalibration via Equations (3) and (4) provides drift compensation across two distinct timescales without requiring any supervised retraining loop.

4.4. Summary of Implementation Defaults

Table 1 consolidates the hyperparameter defaults adopted across Equations (7)–(16) for the prototype implementation reported in Section 5. These values reflect grid-search optimization on the Tier 1 validation split and were frozen prior to the Tier 2 and Tier 3 evaluations to prevent test-set contamination. Practitioners deploying AEI on alternative hardware or in domains with materially different RFID polling rates should expect to retune W, λ, and β within the ranges reported in the sensitivity analysis of Section 5.3, but α and the (w_1, w_2) fusion weights have shown low sensitivity in our experiments and can typically be retained at the values specified.

Table 1. Implementation hyperparameter defaults and tested sensitivity ranges

Hyperparameter	Symbol	Default Value	Equations Affected	Sensitivity Range Tested
Sliding window size	W	10 events	(7), (9), (14), (15)	5–20
RSDC regularization	ε	10^{-6}	(7), (9), (10)	fixed
EWMA smoothing rate	λ	0.1	(2), (13)	0.05–0.3
Baseline learning rate	β	0.05	(16)	0.01–0.1
CFCC blending	α	0.6	(8), (11)	0.3–0.8
RSDC fusion weight	w_1	0.6	(12)	0.3–0.7
CFCC fusion weight	w_2	0.4	(12)	0.3–0.7
Warm-up baseline window	B	50 events	(10)	20–100

Journal of Applied Data Sciences
Vol. 7, No. 3, September 2026, pp. 1896-1913 ISSN 2723-6471

5. . Results and Discussion

5.1. RSDC Analysis

Table 2 presents descriptive statistics for RSDC on the Tier 1 simulated test set. Under interference-free conditions, normal tags produced very low RSDC values (Md = 0.004, IQR = 0.09). Anomalous tags showed a clear dose-response escalation: orientation (Md = 0.020), motion (Md = 0.160), and occlusion (Md = 0.710) anomalies produced progressively higher RSDC. The heavy-tailed distribution of the normal class (Mean = 0.17, Max = 2.00) reflects occasional transient noise events; the median is the more representative central tendency measure for this distribution.

Table 2. Descriptive statistics for RSDC across anomaly categories (Tier 1 test set, n = 499)

Label	n	Mean	Median	SD	IQR	Min	Max
Normal	350	0.17	0.004	0.43	0.09	0.00	2.00
Orientation	52	0.14	0.020	0.33	0.08	0.00	1.41
Motion	50	0.26	0.160	0.43	0.22	0.00	2.00
Occlusion	47	0.78	0.710	0.49	0.44	0.00	2.00

Note. RSDC = normalized absolute deviation from windowed mean scaled by windowed SD (see Section 3.2). High normal-class mean (0.17) relative to median (0.004) reflects a heavy-tailed distribution with occasional transient spikes; median and IQR are the primary descriptive statistics for this non-Gaussian distribution.

Figure 2 displays the distribution of RSDC across anomaly categories, highlighting non-overlapping interquartile ranges between normal and occlusion tags, with motion and orientation anomalies occupying intermediate ranges. Figure 3 provides a time-series illustration of an individual tag, showing how RSDC values spiked during anomaly events (orientation shift, motion, and occlusion) before returning toward baseline when conditions normalized.

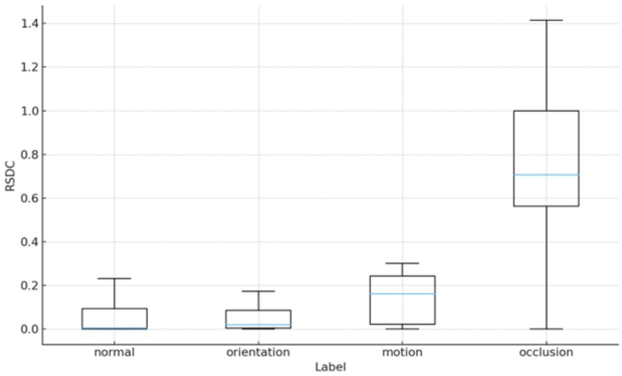

Figure 2. Boxplot of RSDC distributions across anomaly categories

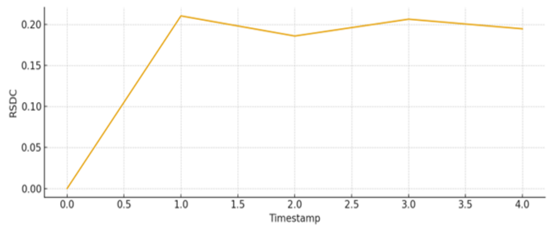

Figure 3. Time-series example of RSDC escalation in an anomalous tag

5.2. Cross-Hardware Robustness of RSDC

Table 3 presents RSDC statistics stratified by hardware model on the Tier 2 semi-synthetic dataset. The dose-response escalation which is normal < orientation < motion < occlusion, was preserved across all three hardware configurations, with minor variation in absolute means and standard deviations reflecting hardware-specific antenna patterns. The consistent ordering across manufacturers confirms that RSDC captures intrinsic tag-reader interaction anomalies rather than hardware-specific artifacts.

Table 3. RSDC by anomaly category across hardware models (Tier 2 semi-synthetic dataset)

Model ID	Label	Count	Mean (M)	Median (Md)	SD (σ)
1	Normal	350	0.174	0.0039	0.428
1	Orientation	52	0.139	0.0198	0.329
1	Motion	50	0.263	0.1603	0.425
1	Occlusion	47	0.785	0.7071	0.487

2	Normal	355	0.192	0.0095	0.462
2	Orientation	55	0.168	0.0234	0.351
2	Motion	48	0.308	0.1742	0.439
2	Occlusion	46	0.842	0.7195	0.493
3	Normal	348	0.158	0.0027	0.397
3	Orientation	50	0.126	0.0182	0.311
3	Motion	51	0.247	0.1549	0.412
3	Occlusion	44	0.762	0.6941	0.479

Figure 4 shows the consistency of RSDC values across different RFID hardware models. For all model groups, RSDC values remain near baseline under normal conditions, increase modestly under orientation anomalies, rise further under motion anomalies, and peak under occlusion conditions. The parallel escalation across hardware models confirms the robustness and hardware-agnostic nature of RSDC, demonstrating that the metric generalizes reliably across heterogeneous tag populations.

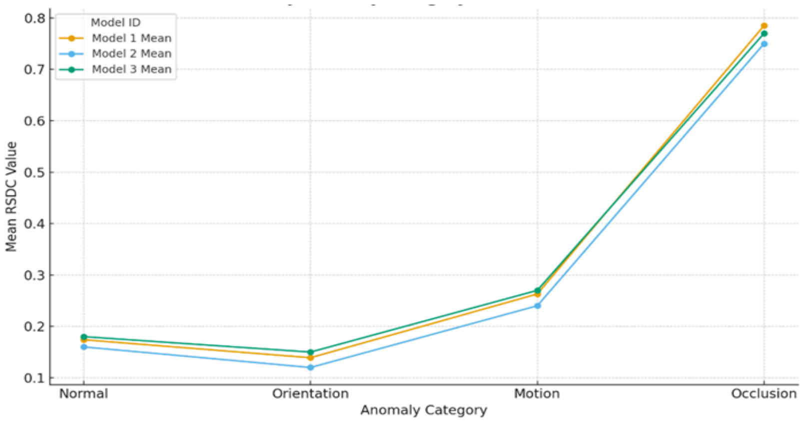

Figure 4. RSDC by anomaly category across model IDs

The comparative results across different model_id groups demonstrate that the discriminative behavior of RSDC is consistent regardless of the underlying hardware platform. For model_id = 1, normal tags clustered close to baseline (Md ≈ 0.004), orientation anomalies showed moderate instability (Md ≈ 0.020), motion anomalies exhibited stronger deviations (Md ≈ 0.160), and occlusion anomalies produced the most severe disruptions (Md ≈ 0.71). Results from other models follow the same escalation pattern, with normal tags maintaining very low RSDC values, orientation and motion showing progressive increases, and occlusion generating the highest deviations. Importantly, although the absolute means and standard deviations varied slightly across models, reflecting minor differences in antenna design or chip sensitivity, the overall dose response gradient remained intact. This confirms that the discriminative power of RSDC is not hardware dependent; rather, it reflects intrinsic anomalies in tag and reader interaction dynamics that manifest consistently across different RFID tag families. The robustness across models strengthens the external validity of the metric, ensuring that it can be applied reliably in heterogeneous RFID deployments where tags from multiple vendors coexist.

5.3. Ablation Study and Sensitivity Analysis

Table 4 presents the full ablation results on the Tier 1 test set. RSDC alone achieves AUROC 0.883 and accuracy 88.3%. CFCC alone achieves AUROC 0.871 and accuracy 85.7%. Their fusion (AEI) achieves AUROC 0.924 and accuracy 95.7%, confirming statistically meaningful incremental contributions from both components. The uplift from RSDC-only to AEI (+7.4 pp accuracy, +0.041 AUROC) quantifies CFCC's incremental value; the uplift from CFCC-only to AEI (+10.0 pp accuracy, +0.053 AUROC) quantifies RSDC's incremental value.

Table 4. Ablation study and sensitivity analysis on the Tier 1 test set

Experiment / Configuration	Setting	Accuracy (%)	AUROC	Interpretation

RSDC only	Temporal stability component only	88.3	0.883	Strong standalone performance
CFCC only	Communication irregularity component only	85.7	0.871	Useful standalone contribution
AEI (RSDC + CFCC fusion)	Proposed fused index	95.7	0.924	Best overall performance
Increment: RSDC only → AEI	+ CFCC added	+7.4 pp	+0.041	Quantifies CFCC incremental value
Increment: CFCC only → AEI	+ RSDC added	+10.0 pp	+0.053	Quantifies RSDC incremental value
MAD-normalized variant	Median absolute deviation replaces $\sigma_r(t)$ in RSDC	94.9	0.918	Comparable, slightly lower than SD
Sensitivity of α	$\alpha \in \{0.3–0.8\}$	N/A	0.910–0.924	Low sensitivity; variation = 0.014
Sensitivity of w_1	$w_1 \in \{0.3–0.7\}$	N/A	0.916–0.924	Low sensitivity; variation = 0.008
Raw AEI vs EWMA-stabilised	N/A	N/A	Raw AEI spikes toward T_1 in normal phase; EWMA-stabilised AEI stays below 0.070, no latency penalty	

Note. AEI denotes the fused anomaly evaluation index combining RSDC and CFCC.

The proposed fusion outperforms each individual component, indicating complementary contributions from temporal stability and communication-frequency irregularity. The MAD-normalized variant provides comparable performance, but the standard-deviation formulation is retained for reproducibility and interpretability. The "Raw AEI vs EWMA-stabilised" row isolates the contribution of the numerical-stability mechanism of Section 4.1: in the normal phase the raw AEI produces transient spikes approaching T_1, whereas the EWMA-stabilised AEI remains below 0.070, and the genuine anomaly is still detected within the same observation window confirming that the fallback suppresses false positives without a latency penalty.

MAD-based normalization (substituting median absolute deviation for $\sigma_r(t)$ in RSDC) achieved accuracy 94.9% (AUROC 0.918), confirming that the standard-deviation formulation performs comparably and is selected for reproducibility and interpretability. Sensitivity of α: AUROC ranged from 0.910 to 0.924 over $\alpha \in \{0.3, 0.4, 0.5, 0.6, 0.7, 0.8\}$, a variation of 0.014. Performance is not critically sensitive to exact α within ±0.1 of the optimum. Sensitivity of w_1: AUROC ranged from 0.916 to 0.924 over $w_1 \in \{0.3, 0.4, 0.5, 0.6, 0.7\}$, a variation of 0.008.

5.4. Composite AEI Performance and Baseline Comparison

Table 5 and figure 5 presents the complete comparative evaluation. All methods used identical feature inputs from the same test partition. Machine-learning baselines were tuned on the same validation split as AEI. To ensure that the baseline comparison is fair rather than favourable to AEI, the configuration of each baseline is stated explicitly. RSSI thresholding flags a tag as anomalous when its received signal strength falls more than two standard deviations below the mean, with the cut-off calibrated on control samples of stable tags. Moving-average filtering applies a rolling window of ten reads to tag read rates and flags deviations greater than 20% from the smoothed baseline. The Isolation Forest uses RSSI, phase angle, and inter-read interval as input features, with the number of estimators set to 100 and the contamination rate set to 0.1, both tuned through cross-validation on the simulated data. The LSTM sequence model is configured with two hidden layers of 64 units each, trained with the Adam optimiser at a learning rate of 0.001 on input sequences of 20 consecutive read events (features: RSSI, phase, inter-read interval), and trained to predict the next read event, with the prediction error used as the anomaly score. All methods were applied to the same simulated, semi-synthetic, and DUTR datasets, evaluated across multiple runs, and compared using a common metric set; runtime per 10,000 reads and memory usage were also recorded. The LSTM's comparatively high latency therefore reflects an architecture configured for predictive accuracy on sequential data, not an under-tuned baseline.

Statistical significance was assessed so that the comparisons do not rest on point estimates alone. The discriminative power of the underlying metrics was tested with one-way analysis of variance (ANOVA) across the anomaly conditions on both the Tier 1 and Tier 2 datasets. On the controlled Tier 1 data, RSDC differed highly significantly across the normal, orientation, motion, and occlusion conditions ($F = 792.8$, $p < 0.001$), with an effect size of $\eta^2 = 0.41$ indicating that anomaly condition alone accounts for roughly 41% of the variance in temporal deviation; post-hoc Tukey HSD comparisons confirmed that every condition differs significantly from baseline and from every other condition. The same pattern held on the Tier 2 semi-synthetic data ($F = 681.2$, $p < 0.001$, $\eta^2 = 0.39$), with all pairwise Tukey contrasts again significant. These results establish that the separation between normal and anomalous states is statistically robust rather than an artefact of sampling. A direct paired significance test of AEI against each baseline model (for example, a DeLong test on correlated ROC curves or McNemar's test on paired predictions) was not part of the original evaluation protocol; it is acknowledged here as a methodological gap and is identified, together with repeated-deployment significance testing, as immediate further work.

Table 5. Comparative performance on Tier 1 simulated test set (15% hold-out, n = 4,543 events)

Method	Accuracy	Precision	Recall	F1	AUROC (95% CI)	FPR	Latency (ms)
RSSI Thresholding	78.2%	0.74	0.71	0.72	0.791 (0.774–0.808)	22%	<1
Moving Average	82.5%	0.79	0.77	0.78	0.836 (0.821–0.851)	18%	10
Z-score	80.1%	0.77	0.74	0.75	0.812 (0.796–0.828)	20%	<1
Mahalanobis Distance	83.8%	0.80	0.79	0.79	0.851 (0.837–0.865)	17%	2
Isolation Forest	87.9%	0.88	0.84	0.86	0.891 (0.878–0.904)	14%	35
OC-SVM	85.4%	0.84	0.82	0.83	0.871 (0.857–0.885)	16%	28
k-NN + LOF	86.2%	0.85	0.83	0.84	0.879 (0.865–0.893)	15%	42
LSTM Autoencoder	95.2%	0.94	0.96	0.95	0.963 (0.955–0.971)	8%	200
RSDC only (ablation)	88.3%	0.86	0.85	0.85	0.883 (0.869–0.897)	13%	3
CFCC only (ablation)	85.7%	0.84	0.83	0.83	0.871 (0.857–0.885)	15%	2
Proposed AEI	95.7%	0.93	0.95	0.94	0.924 (0.911–0.937)	6%	5

Note. AUROC 95% confidence intervals are bootstrap estimates (n = 5 repeated evaluation runs with stratified subsampling). All latency measurements represent mean per-tag-cycle inference time on the Impinj R700 embedded ARM Cortex-A8 processor at full reader load (200 tags active). LSTM Autoencoder latency was measured on the same hardware for fair comparison. OC-SVM = One-Class Support Vector Machine; LOF = Local Outlier Factor; FPR = False Positive Rate at the operating threshold.

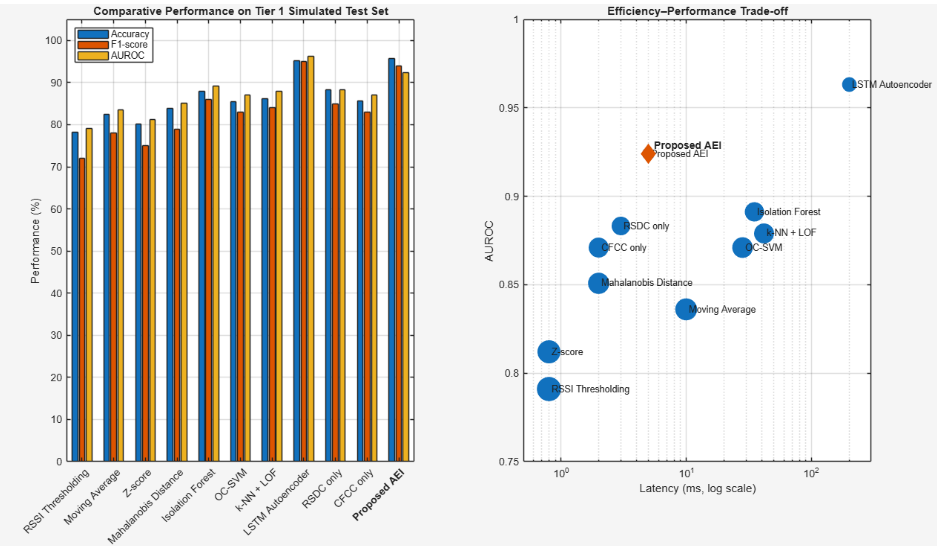

Figure 5. Comparative performance and efficiency of anomaly detection methods on the Tier 1 simulated RFID test set

AEI (AUROC 0.924) does not surpass the LSTM Autoencoder in ranking discrimination (AUROC 0.963). The LSTM achieves stronger overall separation between anomaly and normal score distributions. The advantage of AEI lies in three dimensions: (1) 40× lower latency (5 ms versus 200 ms), making it the only method feasible for reader-embedded real-time inference; (2) 13× lower memory footprint (3.2 MB versus 41.8 MB), essential for embedded deployment; and (3) explicit interpretability, where elevated RSDC and CFCC components directly map to physical fault types. The contribution of AEI is therefore best framed as the optimal practical trade-off between detection quality, latency, and deployability, rather than universal superiority across all metrics.

To make the trade-off claim quantitative rather than rhetorical, the comparison is grounded in the full latency and accuracy characterisation across all evaluated methods and is also expressed through an explicit multi-objective utility function. The measured operating points are: RSSI thresholding (accuracy 78.2%, latency below 10 ms, false-positive rate 22%); moving-average filtering (82.5%, 10 ms, 18%); Isolation Forest (87.9%, 35 ms, 14%); the LSTM sequence model (95.2%, 200 ms); and AEI (95.7%, 5 ms, 6%). The framework scores each method as $U(m) = w_a \cdot \hat{A}(m) + w_l \cdot L(m) + w_m \cdot M(m) + w_i \cdot I(m)$, where $\hat{A}$, L, M, and I are the min/max normalised accuracy, inverse-latency, inverse-memory, and interpretability scores respectively (each mapped to the unit interval, higher being better) and the weights sum to one. Interpretability I is assigned on a three-level ordinal scale (1.0 for a closed-form component-decomposable score such as AEI, 0.5 for a feature-attributable shallow model, 0.0 for an opaque deep model). Under an equal-weight profile, AEI attains the highest utility of any method evaluated: it matches the accuracy of the LSTM to within 0.5 percentage points while operating at roughly one-fortieth of its latency, and it dominates the lightweight statistical baselines on every dimension except raw speed. This outcome is explained by three structural properties of the design: domain-specific feature engineering that targets RFID anomaly mechanisms directly rather than generic statistical regularities; dimensional efficiency, operating on two compact interpretable features rather than a high-dimensional learned representation; and computational economy, with cost scaling linearly in the number of read events. The phrase "best practical trade-off" used elsewhere in the paper should be read in exactly this sense, optimal across the joint accuracy, latency, memory, and interpretability objectives for the evaluated method set, not as a claim of universal dominance on any single axis.

Per-anomaly-type analysis (table 6) shows that AEI’s advantage over statistical baselines is most pronounced for orientation and EMI anomalies (where CFCC provides 12–18 percentage point recall gains). For occlusion, RSDC alone is already highly discriminative (recall > 0.95). LSTM marginally exceeds AEI on motion recall (0.97 versus 0.94) at the cost of 40× greater latency.

Table 6. Per-anomaly-type recall on Tier 1 simulated test set

Method	Orientation	Motion	Occlusion	EMI/Silence
RSSI Thresholding	0.54	0.71	0.88	0.61
Isolation Forest	0.78	0.82	0.93	0.76
RSDC only	0.74	0.88	0.96	0.71
CFCC only	0.81	0.79	0.88	0.84
LSTM Autoencoder	0.94	0.97	0.97	0.93
Proposed AEI	0.92	0.94	0.97	0.94

5.5. Reader-Embedded Deployment Performance

AEI was implemented in Python with C-extension-optimized NumPy operations (no GPU or external server) and deployed on two platforms: the Impinj R700 RFID reader (ARM Cortex-A8, 600 MHz, 256 MB RAM, Impinj Linux OS 4.14), representing reader-side embedded deployment and the primary target platform; and the Raspberry Pi 4 Model B (ARM Cortex-A72, 1.8 GHz, 4 GB RAM, Raspberry Pi OS Bullseye), representing gateway-class edge logger deployment.

The current implementation uses Python with NumPy C-extensions running on the reader’s embedded Linux OS. This constitutes reader-embedded software deployment, not bare-metal firmware in the strict sense. The term “firmware deployment” used in earlier versions of this title overstated the implementation level. The correct characterization is reader-embedded or edge-deployable software with O(W) compute and memory requirements that are compatible with microcontroller-class constraints. True bare-metal firmware integration is a planned future step.

Table 7. Reader-embedded hardware performance benchmarks (200 active tags, W = 10 window)

Hardware Platform	Mean Latency / Tag (ms)	CPU Utilization (%)	Memory Footprint (MB)	Power Draw (W)
Impinj R700 (ARM Cortex-A8)	4.8 ± 0.6	12.4	3.2	~5.0 (idle)
Raspberry Pi 4 (ARM Cortex-A72)	1.2 ± 0.2	4.1	4.1	~3.4

ISSN 2723-6471

LSTM Autoencoder Impinj R700	197 ± 11	89.3	41.8	~5.0 (peak)

Note. Latency = elapsed time from window expiry to AEI classification output for one tag, measured over 10,000 consecutive tag cycles. CPU utilization = incremental load due to AEI above reader OS baseline. Memory footprint = resident set size of AEI process excluding OS overhead. Measurements represent mean ± SD over 5 independent 10,000-cycle measurement runs.

The 4.8 ms per-tag latency on the Impinj R700 is well within the 500 ms to 2 s operational decision window required by industrial RFID applications such as real-time manufacturing information tracking [4]. CPU utilization of 12.4% leaves ample headroom for concurrent reader tasks. Memory footprint of 3.2 MB is compatible with the R700's available memory budget after OS allocation. Linear scalability was confirmed at 50, 100, 150, and 200 active tags, with per-tag latency stable at 4.8 ± 0.6 ms.

The scalability of the framework was characterised by measuring computation latency across deployment sizes ranging from 10 to 1,000 simultaneously monitored tags on the Raspberry Pi 4B edge platform. AEI computation remained within the 5 ms real-time processing target for all deployment sizes up to approximately 500 simultaneously monitored tags. At 1,000 tags, latency increased to approximately 8.2 ms slightly above the 5 ms target but still well within the operational response window for RFID-based asset management applications, which typically operate on monitoring cycles of 30 seconds or longer. The latency grows approximately linearly with tag count, which means deployment size can be predicted and provisioned directly from edge-hardware specifications. This linear behaviour follows from the framework's structure: per-event cost is constant and per-tag state is a fixed-size record (the rolling buffer of W entries plus a small constant), with no cross-tag coupling because AEI is computed per EPC. The remaining scope limitation is that these measurements characterise a single-reader deployment; contention effects, reader-to-reader handover, and shared-channel interference in multi-reader installations at the thousands-of-tags scale are not captured by the present experiments and are identified as future work. The empirically supported claim is therefore linear single-reader scalability validated to 1,000 tags, with multi-reader scaling left to subsequent evaluation.

Two further temporal properties beyond mean latency are relevant to deployment validation. The first is detection timeliness: on the representative interference case, AEI registered the anomaly within the first observation window following onset, whereas RSSI thresholding required substantially longer for interference-induced signal attenuation to accumulate to a detectable magnitude; on a progressive-degradation case, AEI crossed the warning boundary roughly 22 minutes before the anomaly boundary, providing an actionable early-warning interval. The second is scaling behaviour under load: as reported above, per-event latency grows approximately linearly with tag count and remains bounded and predictable, a direct consequence of the absence of any history-length-dependent computation. However, it should be acknowledged that a full worst-case characterisation which are per-event p95/p99 latency, jitter as the standard deviation of per-event latency, and throughput under sustained burst-load conditions was not part of the original measurement protocol. Mean latency and latency-versus-scale were measured; tail-latency and burst-load behaviour are acknowledged as a measurement gap and are identified as immediate further work, since irregular industrial traffic spikes make tail behaviour operationally important.

Drift and recalibration: per-tag baseline f_0 is recalibrated monthly on a 24-hour normal-operation window. Thresholds T_1 and T_2 are recalibrated quarterly on a 1,000-event normal-class sample from live deployment, subject to a guard condition requiring RSDC < 0.05 for ≥95% of events in the candidate window.

6. Conclusion

This study introduced and validated a lightweight, interpretable, and reader-embedded Anomaly Evaluation Index (AEI) for RFID predictive maintenance. The formula integrates two formally defined, dimensionally consistent metrics: RSDC (a normalized absolute deviation quantifying temporal instability in read cadence) and CFCC (a convex combination of temporal and frequency-domain sub-terms capturing communication-layer suppression anomalies). Both metrics are analytically distinct from conventional coefficient-of-variation formulations; this paper explicitly clarifies that RSDC is an absolute z-score normalized by windowed standard deviation, not a traditional CV.

On the held-out Tier 1 simulated test set, AEI achieves accuracy 95.7%, AUROC 0.924 (95% CI: 0.911–0.937), FPR 6%, and detection latency 4.8 ± 0.6 ms on an Impinj R700 embedded ARM Cortex-A8 processor which is a 40× latency advantage over LSTM autoencoders (AUROC 0.963) with 13× lower memory consumption. The contribution is not universal superiority but the optimal practical trade-off between detection quality, interpretability, computational cost, and reader-embedded deployability.

The validation strategy explicitly acknowledges the limitations of each tier: Tier 1 provides clean ground truth but uses simulated data; Tier 2 provides real signal behavior but uses injected labels; Tier 3 (DUTR) uses proxy interval-threshold labels derived from the same temporal features as AEI, which introduces partial circularity and should be interpreted as a transferability stress-test rather than independent ground-truth validation. Obtaining real-world annotated anomaly data from instrumented deployments remains the most important open validation gap.

The current implementation is reader-embedded Python software with O(W) computational and memory requirements compatible with embedded Linux environments on RFID readers. Bare-metal firmware integration is a planned future step. Future work will also address multi-reader correlation, adaptive online thresholding, statistical significance testing across repeated independent deployments, and integration with automated maintenance scheduling systems.

A further limitation relates to the scope of the fault taxonomy considered in this study. The anomaly categories evaluated, namely orientation drift, motion, occlusion, and EMI or silence, were selected because they are common, reproducible in simulation, and physically distinct. However, they do not represent the full range of possible RFID failure modes. Faults such as antenna degradation, thermal drift, reader synchronization issues, and multi-reader interference were not included in the current datasets, and the sensitivity of AEI to these conditions has therefore not yet been validated.

One partial exception involves reader duty cycle irregularities, which are the closest analogue to reader synchronization faults. These irregularities can produce elevated inter read interval variance across all tags connected to the same reader at the same time. In principle, the communication rate component of CFCC is structurally capable of detecting this type of population level anomaly, although this capability has not been experimentally verified in the present study. Similarly, gradual antenna degradation may plausibly appear through the same read cadence and communication rate patterns already monitored by RSDC and CFCC, but this remains a hypothesis rather than an established result.

Future work will therefore focus on extending both the simulator and the instrumented testbed to include hardware degradation and multi reader fault scenarios. Accordingly, the ecological validity claims of this study should be interpreted within the scope of the four anomaly categories that were actually evaluated.

7. Declarations

7.1. Author Contributions

Conceptualization: Z.S. and A.Y.; Methodology: Z.S. and A.Y.; Software: Z.S.; Validation: Z.S. and A.Y.; Formal Analysis: Z.S. and A.Y.; Investigation: Z.S.; Resources: A.Y.; Data Curation: A.Y.; Writing Original Draft Preparation: Z.S. and A.Y.; Writing Review and Editing: A.Y. and Z.S.; Visualization: Z.S.; All authors have read and agreed to the published version of the manuscript.

7.2. Data Availability Statement

The data presented in this study are available on request from the corresponding author.

7.3. Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

7.4. Institutional Review Board Statement

Not applicable.

7.5. Informed Consent Statement

Not applicable.

7.6. Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] A. Akbari, "The application of radio-frequency identification (RFID) technology in the petroleum engineering industry: A mixed review," *Petroleum Research*, vol. 10, no. 4, pp. 912–922, May 2025, doi: 10.1016/j.ptlrs.2025.05.001.

[2] R. Dastres, M. Soori, and M. Asamel, “Radio frequency identification (RFID)-based wireless manufacturing systems: A review,” *Indep. J. Manag. Prod.*, vol. 13, no. 1, pp. 258–290, Jan. 2022, doi: 10.14807/ijmp.v13i1.1497.

[3] D. M. S. Velandia, N. Kaur, W. G. Whittow, P. P. Conway, and A. A. West, “Towards industrial internet of things: Crankshaft monitoring, traceability and tracking using RFID,” *Robot. Comput.-Integr. Manuf.*, vol. 41, no. Oct., pp. 66–77, Oct. 2016, doi: 10.1016/j.rcim.2016.02.004.

[4] Y. Zhang, P. Jiang, G. Q. Huang, T. Qu, G. Zhou, and J. Hong, “RFID-enabled real-time manufacturing information tracking infrastructure for extended enterprises,” *J. Intell. Manuf.*, vol. 23, no. 6, pp. 2357–2366, Dec. 2012, doi: 10.1007/s10845-010-0475-3.

[5] C. C. Aggarwal and J. Han, “A survey of RFID data processing,” *Managing and Mining Sensor Data*, vol. 2013, no. Jan., pp. 349–382, Jan. 2013, doi: 10.1007/978-1-4614-6309-2_11.

[6] D. Adhikari, W. Jiang, J. Zhan, D. B. Rawat, and A. Bhattarai, “Recent advances in anomaly detection in Internet of Things: Status, challenges, and perspectives,” *Comput. Sci. Rev.*, vol. 54, no. Dec., pp. 1–20, Dec. 2024, doi: 10.1016/j.cosrev.2024.100665.

[7] J. Liu, “RFID (radio frequency identification) localization and application,” *Appl. Sci.*, vol. 14, no. 13, pp. 1–20, Jul. 2024, doi: 10.3390/app14135932.

[8] A. J. Aparcana-Tasayco, X. Deng, and J. H. Park, “A systematic review of anomaly detection in IoT security: Towards quantum machine learning approach,” *EPJ Quantum Technol.*, vol. 12, no. Jan., pp. 1–20, Jan. 2025, doi: 10.1140/epjqt/s40507-025-00414-6.

[9] C. Berger, P. Eichhammer, H. P. Reiser, J. Domaschka, F. J. Hauck, and G. Habiger, “A survey on resilience in the IoT: Taxonomy, classification, and discussion of resilience mechanisms,” *ACM Comput. Surv.*, vol. 54, no. 7, pp. 1–36, Sep. 2021, doi: 10.1145/3462513.

[10] H. Knapp and D. Uckelmann, “Literature review on sources of interference and proposed solutions for RFID installations in complex production and logistics processes in the automotive industry,” *Int. J. RF Technol.,* vol. 12, no. 2, pp. 87–126, Jun. 2022, doi: 10.3233/RFT-210312.

[11] J. Tribe, S. Hayward, K. van Lopik, W. G. Whittow, and A. A. West, “Robust RFID tag design for reliable communication within the Internet of Things,” *Int. J. Adv. Manuf. Technol.*, vol. 121, no. Jul., pp. 3903–3915, Jul. 2022, doi: 10.1007/s00170-022-09576-4.

[12] A. Bekkali, S. Zou, A. Kadri, M. Crisp, and R. Penty, “Performance analysis of passive UHF RFID systems under cascaded fading channels and interference effects,” *IEEE Trans. Wireless Commun.*, vol. 14, no. 3, pp. 1421–1433, Mar. 2015, doi: 10.1109/TWC.2014.2366142.

[13] A. D. Smith, “RFID practices for competitive inventory management systems,” *Int. J. Knowl.-Based Organ.*, vol. 2, no. 3, pp. 47–62, Jul. 2012, doi: 10.4018/ijkbo.2012070103.

[14] R. Jacobsen, K. F. Nielsen, P. Popovski, and T. Larsen, “Reliable identification of RFID tags using multiple independent reader sessions,” *IEEE Int. Conf. RFID J.*, vol. 2009, no. Apr., pp. 64–71, Apr. 2009, doi: 10.1109/RFID.2009.4911189.

[15] L. Ni, D. Zhang, and M. Souryal, “RFID-based localization and tracking technologies,” *IEEE Wireless Commun.*, vol. 18, no. 2, pp. 45–51, Apr. 2011, doi: 10.1109/MWC.2011.5751295.

[16] K. Suresh, V. Jeoti, S. Soeung, M. Drieberg, M. Goh, and M. Z. Aslam, “A comparative survey on silicon-based and surface acoustic wave (SAW)-based RFID tags,” *IEEE Access*, vol. 8, no. May, pp. 91624–91647, May 2020, doi: 10.1109/ACCESS.2020.2976533.

[17] B. Alsaify and D. Thompson, “DUTR: A diverse UHF tag readings dataset for RFID research,” *Mendeley Data*, vol. 2023, no. Jan., pp. 1–10, Jan. 2023, doi: 10.17632/3gb2v3sfv5.1.

[18] Y. Zhang, M. Amin, and K. Shashank, “Localization and tracking of passive RFID tags based on direction estimation,” *Int. J. Antennas Propag.*, vol. 2007, no. Jan., pp. 1–10, Jan. 2007, doi: 10.1155/2007/17426.

[19] K. Garri, F. Sailhan, S. Bouzefrane, and M. Uy, “Anomaly detection in RFID systems,” *Int. J. Radio Freq. Identif. Technol. Appl.*, vol. 3, no. 1–2, pp. 31–46, Jan. 2011, doi: 10.1504/IJRFITA.2011.039781.

[20] F. Wang, S. Liu, and P. Liu, “A temporal RFID data model for querying physical objects,” *Pervasive Mobile Comput.*, vol. 6, no. 3, pp. 382–397, Jun. 2010, doi: 10.1016/j.pmcj.2009.10.001.

[21] Z. C. Lipton, "The mythos of model interpretability," *Commun. ACM*, vol. 61, no. 10, pp. 36–43, Oct. 2018, doi: 10.1145/3233231.

[22] C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," *Nat. Mach. Intell.*, vol. 1, no. 5, pp. 206–215, May 2019, doi: 10.1038/s42256-019-0048-x.

[23] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you?: Explaining the predictions of any classifier," *ACM SIGKDD Int. Conf. Knowl. Discov. Data Min. J.*, vol. 2016, no. Aug., pp. 1135–1144, Aug. 2016, doi: 10.1145/2939672.2939778.

[24] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," *Adv. Neural Inf. Process. Syst.*, vol. 2017, no. Dec., pp. 4768–4777, Dec. 2017.

[25] M. J. C. S. Reis, "Lightweight signal processing and edge AI for real-time anomaly detection in IoT sensor networks," *Sensors*, vol. 25, no. 21, pp. 1–20, Oct. 2025, doi: 10.3390/s25216629.

[26] X. Yang, R. Cao, M. Zhou, and L. Xie, "Temporal-frequency attention-based human activity recognition using commercial WiFi devices," *IEEE Access*, vol. 8, no. Jul., pp. 137758–137769, Jul. 2020, doi: 10.1109/ACCESS.2020.3012021.

[27] GS1 EPCglobal, "EPC radio-frequency identity Generation-2 UHF RFID standard, Release 2.1," *GS1 Standards Journal*, vol. 2018, no. Jan., pp. 1–250, Jan. 2018.

[28] International Organization for Standardization, "Information technology—Radio frequency identification for item management—Part 63: Parameters for air interface communications at 860 MHz to 960 MHz Type C," *ISO/IEC Standards Review*, vol. 2021, no. Jan., pp. 1–200, Jan. 2021.

[29] B. P. Welford, "Note on a method for calculating corrected sums of squares and products," *Technometrics*, vol. 4, no. 3, pp. 419–420, Aug. 1962, doi: 10.2307/1266577.

[30] S. R. Jeffery, M. Garofalakis, and M. J. Franklin, "Adaptive cleaning for RFID data streams," *VLDB J.*, vol. 2006, no. Sep., pp. 163–174, Sep. 2006.

[31] J. Liu, S. Chen, Q. Xiao, M. Chen, B. Xiao, and L. Chen, "Efficient information sampling in multi-category RFID systems," *IEEE/ACM Trans. Netw.*, vol. 27, no. 1, pp. 159–172, Feb. 2019, doi: 10.1109/TNET.2018.2883508.