Journal of Computers, Mechanical and Management
Journal homepage: www.jcmm.co.in
e-ISSN: 3009-075X
AN PUBLISHING

Volume 5 Issue 1
Article Number: 25206

Advances In Adaptive Machine Learning Algorithms for Enhanced Security In IoT Networks: A Comprehensive Review

Jayashri Jayesh Patil and Ramkumar Solanki*

Department of Computer Science and Engineering, Sandip University, Nashik, Maharashtra, India 422213

Abstract

The accelerated growth and diversification of Internet of Things (IoT) environments have compounded security issues that are caused by non-stationary traffic, dynamic attacker tactics, and extreme resource constraints. In this case, the failure of the existing process for preventing intrusion through conventional, manually configured intrusion detection systems is becoming less sustainable, while adaptive machine learning (ML)-based security systems are becoming more popular. Nonetheless, the survey literature is generally more biased toward algorithmic enumeration or single-detection accuracy and provides very little critical evaluation of adaptation mechanisms, dataset realism, real-time capability, and implementation constraints. This paper provides a systematic, deployment-conscious review of adaptive machine learning methods for IoT security, conducted through a PRISMA-directed, systematic literature review. An integrative taxonomy of learning paradigms and adaptation methods is presented, a critical review of popular datasets on IoT security is conducted, and a discussion of performance metrics beyond accuracy, such as false positives, robustness, latency, and energy overhead, is provided. The review also examines real-time deployment issues related to edge-cloud, resource constraints, retraining costs, and orchestration complexity, as well as the security of adaptive models against adversarial manipulation and privacy leakage. The literature review, by reconceptualizing the problem of IoT intrusion detection as an adversarially exposed, adaptive learning task with deployment constraints, identifies gaps in the field and provides directions for future work to build a scalable, reliable, and dependable IoT security system.

Keywords: IoT Security; Adaptive Machine Learning; Intrusion Detection System; Federated Learning; Adversarial Machine Learning; Real-Time IoT Security

1. Introduction

The Internet of Things (IoT) has been rapidly growing, driving widespread adoption of connected devices across key sectors, including healthcare, industrial automation, smart cities, and intelligent transportation systems. Although this connection enables round-the-clock monitoring and data-driven decision-making, it has also greatly increased the attack surface of contemporary networks. According to recent research, IoT ecosystems have become targets of large-scale, adaptive cyberattacks, especially distributed denial-of-service (DDoS) attacks, orchestrated by fast-adapting botnets that exploit device heterogeneity, lax authentication, and persistent traffic-generation patterns [1, 2].

*Corresponding Author: Ramkumar Solanki (ramkumar.solanki@sandipuniversity.edu.in)
Received: 26 Mar 2025; Revised: 15 Jan 2026; Accepted: 25 Jan 2026; Published: 28 Feb 2026

DOI: 10.57159/jcmm.5.1.25206.

In contrast to conventional enterprise networks, IoT environments are characterized by non-portable traffic streams, large numbers of devices, and severe limitations in latency, power consumption, and processor performance. The described features make the IoT network particularly susceptible to evolving attacks that can evade implemented defense measures. Recent surveys underscore that the growing complexity of IoT-based attacks renders static, manually configured security solutions insufficient, especially when real-time detection and mitigation are needed [3]. This leads to an increasing desire for security solutions that can independently respond to evolving threat environments and perform their tasks within the realistic limits of IoT deployments.

Traditional IoT security systems are primarily based on signature- or rule-based intrusion detection systems (IDSs). Although this is a good method for identifying known attack patterns, it has inherent weaknesses when used in dynamic, large-scale IoT environments. An IDS based on signatures is not very effective at detecting zero-day attacks, encrypted malicious traffic, and previously unknown attack patterns, leading to high false-negative rates in practice [4].

Traditional security solutions also suffer from scalability and adaptability constraints, limiting their effectiveness. The wide range of IoT devices, combined with limited memory, computing power, and energy, limits the availability of regularly updated rules and centralized traffic analysis. Experimental studies indicate that the performance of static IDS systems degrades with time due to dynamic network conditions and attack schemes, especially in long-running IoT applications [5, 6]. Such limitations point to the necessity of security systems capable of learning from data, extrapolating beyond predefined rules, and autonomously adapting to new threats.

1.1. Adaptive Machine Learning for IoT Security

Adaptive machine learning (ML) has emerged as a promising paradigm for addressing the dynamic nature of threats to the security of IoT-connected devices. In contrast to static models, adaptive ML-based methods can continuously retrain their parameters, features, or decision boundaries when network traffic or attack behavior changes. Recent studies show that IoT traffic is non-stationary and exhibits concept drift, with its statistical properties varying with factors such as device mobility, firmware updates, workload changes, and adversarial manipulations [7, 8].

Adaptive machine learning is a concept used in this review to refer to security models that dynamically update their parameters, features, or decision boundaries in response to concept drift, evolving attack patterns, or changes in IoT traffic. This flexibility is essential for real-time intrusion detection, where slow or fixed response times can cause serious service outages. According to recent research, ML-based intrusion detection systems can learn more complex, high-dimensional traffic patterns and identify more advanced attack types with greater flexibility than rule-based systems [9].

Moreover, distributed and federated learning paradigms allow adaptive model updates without requiring central access to raw data, helping overcome privacy concerns and reducing communication load in sensitive areas of IoT, such as healthcare and industrial control systems [10]. Nonetheless, such methods also present novel difficulties in terms of computational overhead, latency, resilience, and dataset biases, and require critical, deployment-conscious consideration rather than solely focusing on accuracy metrics.

1.2. Contributions and Novelty of This Review

The review addresses adaptive mechanisms, real-time viability, and deployment-constrained evaluation in heterogeneous IoT environments [2, 6], in comparison to more recent surveys on IoT security.

The primary findings of this review can be outlined as follows:

- Theoretical description of adaptive machine learning for IoT security, including disparities between online learning, incremental retraining, federated adaptation, and drift-sensitive detection models.
- An evaluation addressing deployment factors driven by real-time limitations, computational overhead, and resource limitations of IoT devices and edge-driven environments.
- Critical examination of datasets and performance assessment, focusing on issues of class imbalance, outdated traffic trends, and the inadequacy of accuracy-based analysis.
- A systematic review of the latest research (2022–2024) that presents the gaps in the study that must be filled, including resistance to adversarial attacks, interpretability for IoT operators, and scalable adaptation.

Exploring these issues, this review will not be reduced to descriptive overviews but will provide a critical synthesis that can assist in comprehending the existing limitations in research and future directions for learning-based security solutions.

2. Review Methodology

This review has employed a systematic, transparent approach to identify, screen, and discuss the current literature on adaptive machine learning-based security mechanisms for Internet of Things (IoT) networks. This research approach is intended to ensure that high-quality, peer-reviewed publications are well-researched and, at the same time, applicable to practice-related concerns in the area of IoT security and implementation constraints.

2.1. Literature Search Strategy

To identify recent research on adaptive machine learning in computer science and engineering, a systematic literature review was conducted across four large bibliographic databases: IEEE Xplore, Scopus, Web of Science (WoS Core Collection), and ScienceDirect. The rationale for selecting these databases is that they cover high-quality journal articles and conference papers on IoT security, machine learning, and adaptive intrusion detection systems.

The searches were conducted between 1 January 2019 and 31 December 2024, and the last search was conducted on 12 January 2025. This large search window was used to guarantee complete literature retrieval. After the title/abstract screening and full-text analysis, only studies published in 2022–2024 were retained in the final qualitative synthesis, to capture the latest state of the art in adaptive machine learning-based IoT security.

Peer-reviewed journal articles and conference papers in the English language were taken into consideration only. The exclusion criteria were editorials, book chapters, theses, preprints, patents, and non-peer-reviewed technical reports.

To ensure reproducibility, database-specific Boolean search strings were built using a controlled vocabulary and free-text keywords related to IoT security, adaptive machine learning, and intrusion detection systems. Title, abstract, and keyword searches, when supported by the database, and document-type filters were used to filter the search to journal articles and conference proceedings.

The initial database searches yielded a total of 312 records, distributed as follows:

- IEEE Xplore: 96 records
- Scopus: 88 records
- Web of Science: 71 records
- ScienceDirect: 57 records

After excluding 84 duplicate records, 228 unique studies remained and were screened based on titles and abstracts. Follow-up screening and eligibility examinations were performed according to the PRISMA 2020 guidelines, as shown in Figure 1.

For transparency and reproducibility, the exact search queries, applied filters, and database-specific search configurations are provided in Appendix A.

2.2. Inclusion and Exclusion Criteria

Inclusion and exclusion criteria were applied to select the literature, ensuring that the reviewed literature is relevant, high-quality, and technically rigorous. This review relied on peer-reviewed journal articles and conference papers on machine learning-based security mechanisms for Internet of Things (IoT) networks. To be included, the study had to access real-world, simulated, or benchmark IoT data, and it had to be technical enough to allow reproduction and comparison. For this reason, the literature search was limited to 2019–2024 to capture recent contributions and advancements in adaptive IoT security.

Further studies were excluded when they were not related to IoT environments, did not use machine learning-based security methods, or were not peer-reviewed (e.g., opinion pieces, editorials, and technical reports). Moreover, papers that lacked experimental validation, lacked technical clarity, or described non-adaptive security mechanisms were excluded from final analysis. The criteria were applied to obtain technically sound, sufficiently relevant studies that align with the objectives of this review.

In addition to topical relevance, included articles had to be characterized by technical transparency and full reporting to ensure reproducibility and comparative analysis. To quantify this requirement, during full-text assessment, a structured quality appraisal rubric emphasizing dataset reporting, methodological clarity, evaluation rigor, and deployment relevance was used. The appraisal criteria and scoring process are presented in Section 2.5.

2.3. Study Selection and Screening Process (Revised)

The records from the chosen databases were exported to BibTeX and CSV formats and combined into a single reference library. A two-stage process was adopted to identify and delete duplicate records. First, reference management software that performs exact matching of title, author list, publication year, and Digital Object Identifier (DOI) was used to detect duplicates. Secondly, a manual check was performed to detect near-duplicates due to indexing errors, e.g., abbreviated titles, conference extensions, or missing DOI.

After duplicate removal, 228 unique records remained and were subjected to title and abstract screening to assess their relevance to adaptive machine learning–based security in Internet of Things (IoT) environments. The primary reviewer conducted screening in accordance with the predefined inclusion and exclusion criteria described in Section 2.2. To reduce the risk of premature exclusion, studies with ambiguous relevance were retained for full-text assessment.

Subsequently, 72 full-text articles were evaluated considering their eligibility. The primary reviewer performed full-text screening based on explicit criteria: (i) relevance to IoT settings, (ii) use of machine learning-based security schemes, and (iii) presence of adaptive, learning-based, or drift-aware features. The screening protocol involved a single reviewer, as is typical in systematic reviews in the field of computer science and engineering when objective technical criteria are explicitly defined and clearly applied.

A total of 34 full-text articles were excluded, and the reasons for exclusion were documented and classified to facilitate transparency and auditing. The qualitative synthesis and thematic analysis were conducted on the final set of 38 studies. Quality appraisal was carried out alongside full-text eligibility assessment, as described in Section 2.5.

One reviewer performed screening of studies, data extraction, and taxonomy classification, based on clearly defined inclusion and exclusion criteria, formalized taxonomy-mapping guidelines, and data-extraction templates. This formalized, rule-oriented workflow was created to guarantee uniformity, repeatability, and auditability of screening and coding decisions and to reduce subjective interpretation. Ambiguous cases were retained for full-text review to prevent premature exclusion of potentially pertinent studies.

2.4. Reasons for Full-Text Exclusion

To maintain transparency in the PRISMA flow and facilitate independent evaluation of the selection process, the rationale for excluding full-text articles was documented during their eligibility evaluation. The 34 omitted articles were eliminated for the following reasons:

Table 1: Reasons for Full-Text Exclusion ($n = 34$)

Exclusion Category	**Description**	**Number of Studies**
Not IoT-focused	Study addressed general networks or cybersecurity without an IoT context	9
Non-adaptive ML	ML-based IDS without adaptation, drift handling, or learning updates	8
No experimental validation	Conceptual or architectural work lacking empirical evaluation	6
Non-ML-based security	Rule-based, cryptographic-only, or policy-driven approaches	5
Insufficient technical detail	Inadequate methodological or dataset description	4
Duplicate/extended version	Overlapping or earlier version of an included study	2
Total		34

This classification provides a trace of the PRISMA flow and aligns with best practice in systematic reviews.

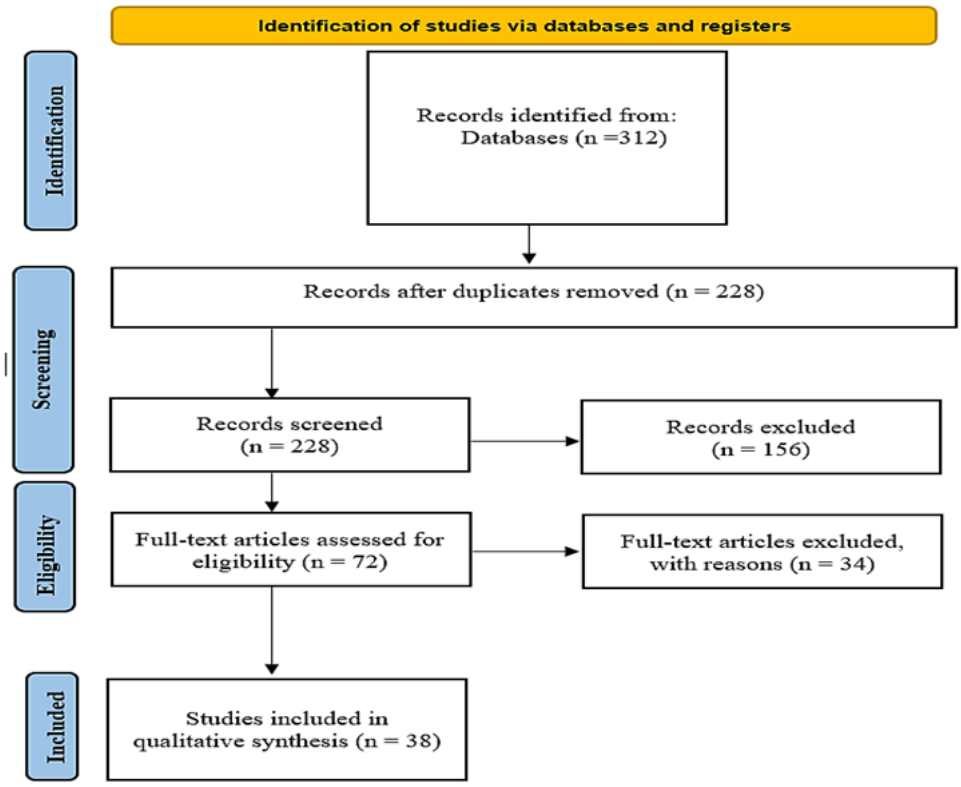

Figure 1: PRISMA 2020 flow diagram illustrating the study identification, screening, eligibility assessment, and inclusion process.

Note: Duplicate detection procedures, screening protocol, and categorized reasons for full-text exclusion are reported in Sections 2.3 and 2.4.

2.5. Quality Appraisal and Reproducibility Assessment

To ensure that the inclusion criterion of technical adequacy and reproducibility was met consistently and transparently, a formal quality appraisal was conducted for all 72 full-text articles evaluated against the eligibility criteria. Since the focus of this review is on engineering and computer science, conventional clinical risk-of-bias assessment tools could not be used. Rather, a domain-specific appraisal rubric was created based on widely accepted reporting standards in machine learning and IoT security research.

All studies were assessed across five quality dimensions relevant to reproducibility and technical rigor. Each criterion was rated on a binary scale (0 = not satisfied, 1 = satisfied), yielding a total of 5 possible points per study.

Only studies achieving a minimum quality score of ≥ 3 out of 5 were retained for inclusion in the final synthesis. This cutoff was chosen to allow for at least some baseline level of methodological rigor and to avoid being too restrictive in excluding recent or exploratory research. A score of three indicates that the research demonstrates sufficient data transparency, clear methodological rigor, and evaluation rigor, meeting the requirements for reproducibility and comparative analysis in IoT security studies. A more stringent sensitivity check ($\geq 4/5$) did not alter the thematic patterns or the general conclusions of the review, suggesting that the results are robust to plausible changes in the quality appraisal cutoff.

Table 2: Quality Appraisal Criteria for Study Inclusion

Criterion	Description
Q1. Dataset transparency	Dataset source, characteristics, and class composition are clearly described
Q2. Methodological clarity	Model architecture, features, and training procedure are adequately specified
Q3. Evaluation rigor	Use of appropriate metrics beyond accuracy and a clear experimental protocol
Q4. Reproducibility support	Sufficient detail to allow replication (parameters, splits, setup)
Q5. Deployment relevance	Discussion of computational cost, latency, or real-world IoT constraints

Scoring rule:

- Score range: 0–5
- Inclusion threshold: ≥ 3

The quality appraisal rubric was used to select 38 studies that met or surpassed the minimum quality threshold (score ≥ 3). The other 34 papers were not included due to failure to meet the predefined eligibility and quality appraisal criteria, including lack of methodological description, absence of experimental validation, or insufficient information to ensure reproducibility. The results of the quality appraisal are in line with the exclusion criteria outlined in Section 2.4 and are consistent with the PRISMA flow diagram depicted in Figure 1.

The quality appraisal criterion of deployment relevance was added to ensure the review met its stated objective of assessing practical adaptive machine learning strategies for real-world IoT settings. As IoT intrusion detection systems are strict on latency, computational, and resource constraints, factors of deployment, including the operational environment, computational overhead, and real-time feasibility, were considered indicators of practical applicability rather than indicators of a limited review scope. Including this criterion may introduce selection bias toward studies that discuss system-level aspects more extensively; however, the scope of the review was limited to deployment-ready, operationally practical IoT security solutions.

2.6. Data Extraction and Coding Procedure

A standardized protocol for data extraction and coding was used across all 38 included studies to achieve consistency and transparency in the comparative analysis before populating Table 8.

An extraction-data form was standardized to include the study characteristics considered during the adaptation process, dataset realism, and deployment feasibility. Each study was coded according to the following fields:

- Target threat or application domain
- Machine learning technique(s) employed
- Adaptation type (static, incremental, online, federated, drift-aware)
- Dataset(s) used and dataset provenance
- Dataset quality indicators (e.g., synthetic vs. real traffic, class imbalance, dataset aging)
- Reported evaluation metrics
- Resource and deployment considerations (e.g., latency analysis, energy cost, edge feasibility)
- Stated limitations relevant to real-time IoT deployment

The primary reviewer performed extraction and coding using explicit coding rules. Table 8 used qualitative descriptors (e.g., Synthetic environment, no latency analysis) only when they were mentioned or could be inferred directly from the study's experimental design (e.g., simulation-only testing; latency results were not reported). Unclear cases were conservatively coded as "Not reported."

Since a single-reviewer protocol was used, inter-rater agreement statistics could not be calculated; nevertheless, a predefined extraction template and coding guide were used to reduce subjective interpretation and ensure internal consistency.

3. Adaptive Machine Learning in IoT Security: Concepts and Taxonomy

Internet of Things (IoT) environments are dynamic, large-scale, and heterogeneous and require security mechanisms that can constantly respond to traffic statistics, attack techniques, and operational limitations. Offline-trained models that are deployed in a static environment are becoming incapable of maintaining reliable detection performance in these non-stationary environments, and their performance declines with time, increasing false alarms and unnoticed attacks. This section formalizes the idea of adaptive machine learning in the context of IoT security, defines a consistent taxonomy of how adaptation can be executed, and clearly states the conditions that must be met to determine when and how adaptation should take place.

This review includes statements describing methods observed, performance characteristics, or limitations that are directly supported by the studies incorporated into the final synthesis ($n = 38$). Generalized taxonomies, cross-sectional meanings, and forward-looking observations in later passages are syntheses based on comparison of multiple studies, rather than empirical results of a particular study. This distinction is maintained to clearly separate evidence-based observations from analytical interpretation.

3.1. Adaptive Machine Learning for IoT Security

Adaptive machine learning can be defined as a concept that refers to learning systems whereby model parameters, structure, or decision boundaries are dynamically changed in response to changes in data distributions or operating conditions. Dynamic traffic, device churn, firmware variability, workload variability, and adversarial activity are factors that have resulted in such changes in IoT security environments. Unlike fixed ML models, where the relationship between features and labels is considered static, adaptive models explicitly consider non-stationarity and change their behavior during deployment rather than merely retraining offline.

Data stream learning studies demonstrate that concept drift is an inherent characteristic of actual IoT traffic and renders static intrusion detection models unsuitable in long-term applications [11, 12]. Resistance to new attacks, detection accuracy, and false-positive control decline because of changing network behavior when fixed models are used. To address these limitations, adaptive ML systems incorporate resource constraints, drift awareness, and update mechanisms into the learning process, and their reliability in detection can be maintained over time [13, 14]. The distinction between fixed and adaptive ML is fundamental from a security perspective. Whereas fixed models are oriented toward offline optimization, adaptive models are oriented toward operational resilience under continuous change, which is a significant feature in IoT environments, where traffic evolution and adversarial adaptation are regular occurrences.

All sources included in Table 1 were chosen from the 38 studies identified and selected after full-text screening and quality appraisal to avoid mixing evidence levels in the table. They were representative studies that implemented the respective taxonomy category, described their methodology explicitly, and were relevant to adaptive learning in the context of IoT security. Table 1 is not an exhaustive list of all studies addressing adaptation mechanisms but is intended to be illustrative. Non-included literature and general background surveys are mentioned in the narrative discussion and support conceptual framing but are not utilized as primary evidence in the taxonomy table.

3.2. Taxonomy of Adaptation Mechanisms

Adaptive machine learning approaches to IoT security can be grouped into broad categories based on the location and manner of adaptation. The literature identifies five predominant categories.

Online learning solutions update model parameters as new data samples emerge and enable real-time response to dynamic trends in IoT traffic. They are particularly suitable for streaming systems where it is not possible to store historical data and low-latency adjustments are needed [15]. However, they are prone to interruption by noise and manipulation by attackers, which is undesirable in security-related applications.

Incremental retraining processes are updated regularly with new data batches and offer a trade-off between flexibility and uniformity. This strategy can be pursued in edge-assisted IoT security architectures, where retraining is triggered when performance degradation or workload variation is observed [11]. Incremental retraining is more stable than pure online learning, although training cost and latency must be carefully managed.

Federated adaptive learning enables distributed learning across IoT devices or edge nodes without relaying raw traffic information, addressing privacy and communication constraints. Recent studies focus on adaptive federated methods in which the frequency of aggregation, learning rates, or client participation are modified in real time to accommodate system heterogeneity and non-IID data distributions [16, 17]. Client drift and communication overhead continue to be major challenges despite these advances.

Semi-supervised and self-supervised adaptive learning are used to address the lack of labeled attack data in IoT environments. Typically, unsupervised methods learn traffic representations from unlabeled data, whereas semi-supervised methods leverage a few labeled samples alongside a large amount of unlabeled data to learn continuously [18, 19]. These methods reduce dependency on annotations, although they typically do not guarantee reliable discrimination against sophisticated attacks without careful calibration.

Drift-aware learning systems explicitly monitor data distributions or model behavior to detect concept drift and may induce adaptation accordingly. These are necessary for isolating benign traffic growth from actual security hazards, particularly when gradual or abrupt changes are introduced to IoT network operation [20].

3.2.1 Taxonomy Mapping Rules and Classification Procedure

To prevent subjective classification and ensure consistent categorization of adaptive machine learning methods, explicit mapping rules were established to determine the taxonomy category to which each study would be assigned. Categorization was performed at the method level (not solely based on authors' claims) and was based on observable technical features of each study.

The reviewed papers were assigned to one or more taxonomy categories based on the following operational criteria:

- **Online Learning:** A study was considered online learning when model parameters were updated at every deployment step or data instance in a stream, without discrete retraining cycles or full access to historical data.
- **Incremental Retraining:** Research was classified as incremental retraining when model updates were made through periodic or batch retraining driven by the accumulation of new data, performance degradation, or a predefined update schedule.
- **Federated Adaptive Learning:** A study was considered federated adaptive learning when it used distributed model training among IoT devices or edge nodes, where model updates (e.g., gradients or weights) were aggregated without raw-data sharing, and incorporated adaptive factors such as dynamic aggregation, client selection, or update frequency.
- **Self-Supervised / Semi-Supervised Learning:** A study was categorized in this class when it directly used unlabeled or weakly labeled data for representation learning, anomaly detection, or adaptation, and involved limited or partial supervision in either the training or update process.
- **Drift-Aware Learning:** A study was categorized as drift-aware when it included explicit concept drift detection mechanisms (e.g., statistical tests, distribution monitoring, performance-driven alarms) and when drift signals were used to trigger model adaptation or retraining.

Multi-labeled studies that used various mechanisms of adaptation were allowed rather than being forced into mutually exclusive categories.

3.2.2 Classification Protocol and Bias Considerations

The above-described mapping rules were followed by the primary reviewer when assigning taxonomy categories. Classification decisions were based on methods and mechanisms clearly stated in each paper, not on authors' terminology or self-proclaimed model labels.

Since the study was technical and review-based, the classification protocol involved a single reviewer. Although it was therefore impossible to compute inter-rater agreement measures, explicit, rule-based mapping criteria were used to reduce subjectivity and ensure internal consistency across classifications. Ambiguous cases were conservatively classified only when sufficient methodological evidence was present in the text.

3.3. Adaptation Criteria

Adaptive machine learning-based IoT security systems are based on explicit trigger-action mechanisms to decide when and how model updates are made. Based on the analysis of the 38 included studies, three major categories of adaptation triggers were identified. Where explicit trigger-action loops were applied in the evaluated literature, they are mentioned by reference. Where the evidence was fragmented or indirect, the criterion is provided as a synthesis based on comparison of several studies.

3.3.1 Traffic-Driven Triggers

Many studies that implement adaptation respond to apparent changes in network traffic, including packet rates, protocol distributions, and communication patterns. Such mechanisms usually track traffic statistics or feature distributions and take action to adapt when deviations exceed predefined limits.

Drift-aware intrusion detection systems explicitly track distributional changes and trigger retraining whenever drift detectors indicate substantial behavioral changes relative to a baseline. The literature on these mechanisms shows that traffic-triggered events are useful for sustaining detection accuracy in non-stationary IoT traffic. However, the sensitivity of threshold selection and drift detection varies widely across implementations, and no consistent measure or cutoff has been reported across studies.

Evidence-backed: Several included studies specifically apply traffic-distribution monitoring and drift-based retraining.

Limitation: Thresholds and statistical tests are heterogeneous and are seldom empirically justified.

3.3.2 Performance-Driven Triggers

The second category of adaptation criteria uses model performance degradation as the trigger for adaptation. These methods involve performance measures such as detection rate, false positive rate, or classification confidence, which are continually or periodically monitored. When performance falls below predefined acceptable levels, the adaptation process is initiated.

Some studies note retraining or updating models as false-positive rates rise or detection accuracy decreases over time, especially in long-term IoT deployments. Performance-driven triggers are typically employed in incremental retraining and federated learning-based intrusion detection systems, where retraining is driven by performance degradation rather than continuous online updates.

Evidence-backed: Performance-based retraining is described in several included studies, particularly in incremental and federated learning settings.

Limitation: In most studies, explicit numeric thresholds are not provided, and trigger sensitivity is not justified, which limits reproducibility.

3.3.3 Resource-Aware Triggers

Resource-aware adaptation treats computational capacity, memory, and energy as primary triggers in IoT security systems. Some studies directly limit adaptation frequency or model complexity based on resource availability at IoT devices or edge nodes, such as deferring training to edge or cloud infrastructure when local resources are exceeded.

However, explicitly defined closed-loop trigger-action mechanisms determined by resource thresholds (e.g., using an energy budget to make retraining decisions) are less commonly implemented in practice. The majority of the literature discusses resource awareness qualitatively or evaluates resource consumption ex post but does not integrate resource metrics directly into adaptation logic.

Partially evidence-backed: Resource constraints are widely acknowledged and measured.

Conceptual synthesis: Explicit resource-threshold–driven adaptation loops remain largely conceptual and underexplored.

3.3.4 Summary of Evidence vs. Synthesis

In general, traffic-driven and performance-driven triggers have strong empirical support in the literature considered, whereas resource-aware triggers are either conceptual or only indirectly addressed. This discrepancy highlights one of the main research gaps: the absence of unified trigger-action models that jointly consider traffic evolution, detection reliability, and resource sustainability in real-time IoT settings.

Table 3: Evidence Mapping of Adaptation Triggers in Included Studies

Adaptation Trigger	**Explicit Trigger–Action Loop Reported**	**Metrics Used**	**Evidence Status**
Traffic shifts	Reported in multiple studies	Feature distributions, drift statistics	Empirical
Performance degradation	Reported in multiple studies	Accuracy, FPR, DR	Empirical
Resource thresholds	Rare	Energy, latency (mostly reported, not enforced)	Largely conceptual

4. Taxonomy of Machine Learning Techniques for IoT Security

The heterogeneity of Internet of Things (IoT) deployments, comprising immensely resource-constrained end devices, edge gateways, and cloud-fused infrastructure, has led to the deployment of a host of machine learning (ML)-based intrusion detection and security monitoring solutions. The available literature differs in learning algorithms, deployment architecture, computational overhead, and data representation, and comparison among studies is not always straightforward. To overcome this conceptual gap, a deployment-sensitive taxonomy of ML techniques for IoT security is presented, organized into architectural, algorithmic, compositional, and representational dimensions. The taxonomy provides a standard framework for comparing existing approaches in terms of detection capability and deployment viability.

4.1. Centralized and Distributed IoT Intrusion Detection Architectures

An ML-based intrusion detection system (IDS) for IoT can be categorized based on where data processing and learning occur. In a centralized IDS architecture, traffic information from multiple IoT devices is collected, and model training and inference are performed on a central or cloud server. This approach provides global visibility and supports computationally intensive models, but it introduces scalability bottlenecks, increased detection latency, and a higher risk of privacy leakage because raw data are aggregated [21]. Practical constraints may limit the suitability of centralized IDSs for latency- or privacy-sensitive IoT applications.

Distributed IDS architectures, on the other hand, perform learning and inference at edge devices or gateways, thereby minimizing communication overhead and enabling quicker responses. Federated and collaborative learning paradigms are examples of this pattern, as they facilitate decentralized model updating without raw-data sharing, which is more suited to large-scale, privacy-sensitive IoT deployments [22]. Nevertheless, distributed IDSs introduce challenges associated with system heterogeneity, synchronization, partial observability, and client drift, especially when traffic is non-IID [23]. This architectural distinction fundamentally influences the feasibility and performance of ML-based IoT security solutions.

4.2. Traditional Machine Learning vs. Deep Learning Approaches

From an algorithmic perspective, ML-based IoT security solutions can be classified into traditional machine learning and deep learning strategies. Traditional ML methods, including decision trees, support vector machines, k-nearest neighbors, and naive Bayes classifiers, rely on handcrafted features and relatively shallow models. These approaches are computationally efficient and explainable, and hence appealing for implementation on constrained IoT devices or edge nodes [24]. Deep learning (DL) methods, such as convolutional neural networks, recurrent neural networks, and autoencoders, automatically learn hierarchical representations from raw or minimally processed data. This capability has enabled improved detection performance across various IoT security settings, especially when dealing with large-scale or high-dimensional traffic data [25].

However, DL models are generally more computationally demanding and memory-intensive, which increases latency, energy consumption, and explainability challenges. As a result, the choice between traditional ML and DL methods introduces a trade-off between detection performance and deployment feasibility in IoT environments.

4.3. Hybrid and Ensemble Models

Hybrid and ensemble models aim to provide a compromise between expressiveness and robustness by integrating multiple learning methods. Hybrid models generally combine feature engineering with ML, or integrate conventional ML classifiers with DL-based feature extractors, to achieve improved generalization on heterogeneous IoT traffic while partially managing computational costs [26]. These methods are particularly useful in complex or variable traffic environments.

Ensemble learning techniques, including bagging, boosting, and stacking, combine predictions from multiple base learners to reduce variance and improve robustness to noise and class imbalance. Ensemble-based IDSs can be more resilient to evolving attacks than single-model IDSs in IoT security settings [27]. However, ensemble schemes increase inference latency, memory usage, and management complexity, which may restrict their use in resource-constrained IoT deployments. Their inclusion in this taxonomy highlights the need to consider both performance improvements and operational costs.

4.4. Lightweight Machine Learning for Resource-Constrained IoT Devices

A large percentage of IoT devices operate under stringent memory, processing power, and energy constraints, necessitating the development of lightweight ML methods. Lightweight IDS designs use simplified model architectures, reduced feature sets, or model compression to minimize resource consumption while maintaining acceptable detection performance [28].

Recent studies show that carefully designed lightweight DL models (e.g., shallow neural networks or optimized autoencoders) can achieve competitive performance while significantly reducing computational cost [29]. However, the effectiveness of these methods depends heavily on model configuration and workload characteristics. Resource-conscious learning is therefore essential for achieving operational sustainability in real-world low-power IoT environments, where architectural security considerations remain critical [30].

4.5. Graph-Based and Traffic Flow Modeling Techniques

In addition to flat feature-vector representations, graph-based and traffic-flow modeling methods have proven effective for detecting IoT intrusions. These approaches model network entities and their interactions as a graph, allowing structural and relational information to be captured that may be obscured in traditional representations. In particular, graph neural networks (GNNs) have shown potential for identifying coordinated and distributed attacks by exploiting topological and temporal dependencies [22].

In IoT environments, communication among devices, gateways, and services is analyzed using graph-based IDS approaches to identify anomalies at the system level rather than at isolated nodes. Although recent methods demonstrate improved detection of advanced attack scenarios, graph-based approaches introduce challenges, including graph construction overhead, scalability limitations, and real-time inference constraints [31]. Doctoral research also emphasizes that, despite the representational advantages of traffic graph modeling, its practical implementation in large-scale IoT systems remains challenging [32].

This taxonomy section indicates that no universal ML technique is optimal for IoT security. Detection effectiveness cannot be evaluated independently of deployment architecture, resource constraints, and data representation. Consequently, IoT intrusion detection should be treated as a system-level design challenge, where learning algorithms are selected not only for accuracy but also for scalability, adaptability, and robustness in real-world environments.

All references listed in Table 4, including static baseline comparators, are drawn exclusively from the 38 studies included in the final systematic review (see Figure 1). “Representative” indicates illustrative examples selected to exemplify each taxonomy category, not an exhaustive listing. Background surveys and non-included studies are cited only in the narrative text and are not used as evidentiary support in this table.

Table 4: Taxonomy of Machine Learning Techniques for IoT Security (with Representative Included Studies)

Category	**Learning Paradigm**	**Typical Algorithms / Models**	**Deployment Architecture**	**Key Advantages**	**Key Limitations**	**Representative References**
Architectural	Centralized IDS	SVM, Random Forest, Deep Neural Networks	Cloud / Central Server	Global visibility, high detection accuracy, simplified model management	High latency, scalability issues, privacy concerns	[21, 24]
Architectural	Distributed / Federated IDS	Federated averaging, collaborative ML, edge-based DL	Edge / Fog / Federated	Privacy preservation, reduced communication overhead, improved scalability	Client drift, system heterogeneity, synchronization complexity	[22, 23]
Algorithmic	Traditional ML	Decision Trees, k-NN, Naïve Bayes, SVM	Edge / Gateway	Low computational cost, interpretability, suitability for constrained devices	Limited representation power, manual feature engineering	[24]
Algorithmic	Deep Learning	CNN, RNN, LSTM, Autoencoders	Cloud / Edge	Automatic feature extraction, strong performance on complex traffic	High resource consumption, limited explainability	[25]
Model Composition	Hybrid Models	Feature engineering + DL, ML–DL fusion	Edge / Cloud	Balanced performance and complexity, adaptability to heterogeneous traffic	Increased system complexity, tuning overhead	[26]
Model Composition	Ensemble Models	Bagging, Boosting, Stacking	Edge / Cloud	Improved robustness, reduced variance, resilience to noise	Higher inference latency, model management overhead	[27]
Resource Awareness	Lightweight ML	Shallow NN, compressed DL, reduced-feature ML	IoT Device / Edge	Low latency, reduced energy consumption, real-time feasibility	Potential accuracy degradation, limited expressiveness	[28, 29]
Representation	Graph-Based Models	Graph Neural Networks (GNNs), Flow Graphs	Edge / Cloud	Captures relational and structural patterns, effective for coordinated attacks	Graph construction overhead, scalability challenges	[22, 31, 32]

5. Thematic Review of State-of-the-Art Studies

Recent studies on machine learning-based IoT security have increasingly shifted toward adaptive, privacy-aware, and distributed learning paradigms rather than static intrusion detection approaches. However, much of the current literature remains focused on algorithmic performance in controlled experimental settings, with limited critical discussion of real-time viability, system overhead, and deployment constraints. This section addresses these limitations by summarizing cutting-edge research across four prevailing themes: hybrid and ensemble learning, federated learning, feature selection and optimization, and deep learning-based intrusion detection, while explicitly considering their applicability to real-time IoT security implementation. Instead of focusing solely on detection accuracy, the analysis emphasizes latency, scalability, adaptability to evolving traffic, and operational cost.

5.1. Hybrid and Ensemble Machine Learning Approaches for DDoS Detection

Hybrid and ensemble learning methods are frequently recommended for detecting distributed denial-of-service (DDoS) attacks in IoT networks due to their robustness in monitoring heterogeneous traffic patterns. Hybrid models combine complementary techniques, such as feature engineering with machine learning classifiers, traditional ML with deep learning feature extractors, or ensemble methods that integrate multiple base learners to reduce variance and improve generalization.

Empirical research consistently reports higher detection rates compared to single-model baselines, especially for heterogeneous or noisy IoT traffic [33, 34]. Nevertheless, these performance gains involve significant deployment costs. Ensemble inference increases latency, memory overhead, and energy consumption, which are constrained in edge- or gateway-level deployments. Consequently, models that perform strongly in offline experiments may exceed real-time processing capabilities in operational IoT environments [35].

Importantly, most hybrid and ensemble studies evaluate performance using fixed datasets and offline testing pipelines, providing limited insight into long-term behavior under evolving traffic or adversarial adaptation. This reveals a fundamental trade-off: hybrid and ensemble models enhance detection robustness, but their structural complexity may hinder real-time execution. Without explicit analysis of latency, throughput, and energy consumption, it is difficult to translate reported performance gains into deployable IoT security solutions.

5.2. Federated Learning for Privacy-Preserving IoT Security

Federated learning (FL) has emerged as a prominent framework for privacy-preserving IoT intrusion detection, enabling collaborative model training without centralized raw data aggregation. FL-based systems retain traffic data on local devices or edge nodes, thereby supporting regulatory compliance and privacy guarantees while improving scalability in distributed environments.

Recent research indicates that FL can reduce data exposure and improve scalability [36]. However, real-time deployment challenges remain significant. Non-IID traffic distributions across IoT devices often lead to client drift and unstable model convergence, degrading detection performance over time. These effects are amplified in highly heterogeneous IoT environments where device behavior, workloads, and network conditions vary considerably.

Real-time operation is further constrained by communication overhead. Frequent model updates may overload limited network bandwidth and edge infrastructure, resulting in unacceptable latency. Adaptive aggregation and resource-aware scheduling have been proposed to mitigate these challenges [23], but they introduce additional trade-offs among convergence speed, detection accuracy, and system complexity. Furthermore, although FL is designed to preserve privacy, FL-based IDSs remain vulnerable to inference attacks, such as gradient leakage and model inversion, raising concerns about their suitability for sensitive IoT applications [37].

Federated learning represents a promising architectural direction for IoT security; however, its practical deployment in real-time environments remains constrained by communication cost, convergence instability, and unresolved privacy risks [38].

5.3. Feature Selection and Optimization-Based Approaches

Feature selection and optimization-based methods aim to improve IoT intrusion detection performance by reducing feature dimensionality while preserving discriminative power. These techniques are particularly relevant for real-time IoT implementations, where computational efficiency directly influences latency and energy consumption. Metaheuristic optimization methods have been shown to identify compact feature subsets that enable lower training and inference cost with minimal loss of accuracy [39, 40].

The Arithmetic Optimization Algorithm (AOA), for example, has been reported to provide balanced exploration and exploitation capabilities with competitive detection performance [41]. Despite these advantages, optimization-based IDSs may lack robustness under varying IoT traffic conditions. Their performance is often sensitive to dataset characteristics, parameter configuration, and attack distribution, raising concerns regarding generalizability. Moreover, most evaluations are conducted on benchmark datasets, offering limited evidence of robustness under shifting attack patterns. Foundational studies in optimization emphasize that efficiency improvements observed under controlled experimental conditions do not guarantee reliable real-time deployment without system-level validation [42].

5.4. Deep Learning–Based IoT Intrusion Detection

Deep learning (DL) methods, such as convolutional neural networks (CNNs), recurrent neural networks (RNNs), and long short-term memory (LSTM) models, have been widely studied in IoT intrusion detection because of their ability to capture complex temporal and spatial traffic patterns. Temporal DL architectures have demonstrated strong performance in modeling sequential network traffic dependencies [43].

However, real-time deployment of DL-based IDSs remains highly constrained. Many models have substantial computational and memory requirements, resulting in high inference latency and energy consumption, which limits their suitability for single-board IoT devices or edge nodes. Moreover, deep models are often overfitted to benchmark datasets; their performance degrades when exposed to unknown or evolving traffic, reflecting limited robustness in operational settings [44].

Explainability further restricts real-world applicability. Opaque decision-making in DL complicates incident response and reduces trust among IoT operators. Although explainable DL methods have demonstrated potential to improve transparency without significantly degrading detection performance [45], their scalability and real-time feasibility remain insufficiently explored. Doctoral research emphasizes that unless latency, interpretability, and robustness are addressed jointly, DL-based IDSs may remain laboratory solutions rather than practically deployable security mechanisms [46].

6. Evaluation of Datasets and Performance Metrics in IoT Security Research

The quality of datasets used to train and evaluate machine learning-based IoT security solutions, along with the performance metrics used to report results, fundamentally constrains their reliability. Although methodological sophistication has increased, much of the literature still emphasizes high detection accuracy without critically examining dataset representativeness, temporal validity, and the operational implications of false alarms, latency, and resource consumption. This section provides a deployment-focused critical analysis of commonly used IoT security datasets and evaluation metrics, and discusses how dataset bias and metric selection may distort claims of effectiveness and real-time viability.

6.1. Commonly Used IoT Security Datasets

Empirical evaluation of IoT intrusion detection studies is dominated by several benchmark datasets, including IoT-23 and CIC-IoT-2023, as well as smaller, domain-specific datasets designed for healthcare IoT systems. Although these datasets facilitate reproducible experimentation, they are not always representative of real-world IoT environments.

The IoT-23 dataset contains labeled benign and botnet traffic collected from real IoT devices, providing insight into malware-driven attack behavior under controlled conditions [47]. However, its experimental configuration covers a limited set of device types, workloads, and attack evolutions, restricting its ability to model large-scale heterogeneous deployments. CIC-IoT-2023 introduces more recent attack scenarios and broader traffic diversity, but remains constrained by scripted attack execution and laboratory network conditions [48].

Specialized healthcare IoT datasets present additional challenges. Although realistic in terms of protocol usage and privacy constraints, they are often small, curated, and difficult to generalize beyond their specific application context [49]. As a result, models trained exclusively on such datasets may exhibit optimistic performance that does not generalize to uncontrolled, heterogeneous IoT environments.

Reliance on static benchmark datasets also obscures the effects of traffic evolution, device churn, and adversarial adaptation, leading to evaluation outcomes that may overstate real-world performance.

6.2. Dataset Quality Issues and Bias

A pervasive limitation across IoT security datasets is class imbalance, where attack traffic may dominate the dataset or be severely underrepresented relative to benign samples. In such cases, accuracy-based evaluation becomes misleading, as models may achieve high overall accuracy while performing poorly on minority attack classes that are operationally critical [50]. This issue is particularly significant in IoT settings, where infrequent but high-impact attacks must be reliably detected.

Another major concern is the reliance on synthetic or simulated traffic. Although synthetic datasets are easier to generate and label, they often lack the variability, noise, and unpredictability of real IoT networks. Empirical studies show that systems trained on artificial data frequently fail under live traffic conditions, revealing discrepancies between test and operational performance [51].

Dataset aging further undermines evaluation validity. Attack patterns, communication protocols, and device behaviors evolve rapidly, rendering many benchmark datasets obsolete. Aging datasets have been shown to degrade the performance of models when confronted with contemporary attack variants, highlighting the risks of drawing long-term conclusions from static benchmarks [52].

6.3. Implications for Performance Metrics Beyond Accuracy

Dataset limitations directly influence the reliability of reported performance metrics. Although accuracy remains the most commonly reported metric, it is insufficient for evaluating intrusion detection in imbalanced and evolving IoT environments. Operationally meaningful metrics, such as false positive rate (FPR) and false negative rate (FNR), provide clearer insight into system behavior, since excessive false alarms disrupt operations, whereas missed attacks compromise security.

Robustness-oriented measures, including the Matthews correlation coefficient and balanced accuracy, have been recommended for more reliable evaluation under class imbalance [53]. Additionally, effective assessment requires cross-dataset validation and evaluation across diverse attack scenarios rather than reliance on single-dataset benchmarking.

Beyond detection correctness, real-time IoT security demands evaluation of latency, throughput, and resource overhead. IDSs operating at the edge or fog layer must satisfy strict timing constraints; otherwise, inference delays may negate their practical value. Empirical studies indicate that models achieving high accuracy may still be impractical due to excessive inference latency or energy consumption [54]. Frameworks emphasizing comprehensive evaluation across detection quality, robustness, latency, and energy efficiency have been proposed [55], but remain underutilized in current IoT security research.

6.4. Descriptive Quantitative Summary of Evaluation Practices

To complement the qualitative analysis, a descriptive quantitative synthesis was conducted across the 38 included studies to summarize evaluation practices related to dataset quality, performance metrics, and robustness assessment. Table counts were derived from the standardized data-extraction protocol described in Section 2.6.

Table 5: Frequency of Evaluation Practices Across Included Studies ($n = 38$)

Evaluation Aspect	**Number of Studies**	**Percentage (%)**
Report accuracy / detection rate	38	100
Report precision / recall / F1	29	76
Report false positive rate (FPR)	16	42
Report false negative rate (FNR)	11	29
Report latency or inference time	9	24
Report energy or resource usage	6	16
Address class imbalance explicitly	14	37
Perform cross-dataset testing	5	13
Evaluate dataset aging / temporal validity	4	11

The results indicate a strong emphasis on accuracy-centric evaluation, with limited reporting of operational metrics such as latency, energy consumption, and robustness to dataset aging. Cross-dataset validation and temporal robustness remain underreported, reinforcing concerns regarding real-world generalizability.

7. Real-Time Deployment Challenges in IoT Environments

Despite the high performance of reported intrusion detection systems (IDSs) using machine learning models in experimental studies, their extrapolation to real-time deployments in Internet of Things (IoT) environments remains limited. A major limitation in the literature is the emphasis on detection accuracy at the expense of system-level considerations, including inference latency, energy usage, retraining cost, and deployment architecture. These constraints ultimately determine whether an IDS can operate effectively in real IoT networks rather than in controlled laboratory environments [56]. This section critically reviews the main deployment issues that hinder real-time adoption of ML-based IoT security solutions.

7.1. Edge vs. Cloud vs. Fog-Based Deployment

The deployment layer fundamentally determines real-time performance of intrusion detection. Cloud-based IDS deployments enable centralized management and support computationally intensive models, but introduce communication latency and dependence on continuous connectivity. Such latency may render detection ineffective in time-sensitive IoT applications, particularly in geographically dispersed or bandwidth-limited environments [57].

Edge-based IDS deployment brings traffic processing closer to IoT devices, minimizing inference latency and network overhead and enabling faster responses. However, edge nodes have limited computational and energy resources, which significantly constrain the complexity of deployable models [54]. As a result, many high-accuracy models proposed in the literature cannot be executed at the edge without substantial simplification.

Fog computing aims to balance these trade-offs by distributing processing across intermediate nodes. Although fog architectures can reduce latency compared to cloud-only deployments, they introduce additional coordination, synchronization, and management complexity [58]. Notably, many IDS studies do not clearly specify the assumed deployment layer, making it difficult to determine whether reported performance can be achieved in real-time operational contexts.

7.2. Resource Constraints of IoT Devices

Resource scarcity represents one of the most immediate barriers to real-time IoT security implementation. IoT devices and gateways are tightly constrained by memory capacity, processing power, and energy supply. Even moderately complex deep learning models may exceed these limits, resulting in unacceptable inference delays or rapid battery depletion [59].

Empirical studies on energy-efficient ML inference consistently show that improvements in detection accuracy obtained through increased model complexity come at the cost of higher energy consumption and reduced system lifespan [54]. Therefore, resource efficiency should not be treated as a secondary optimization objective but rather as a primary design requirement, especially for adaptive IDSs that require continuous monitoring and periodic model updates.

7.3. Retraining Cost and Model Update Frequency

Periodic model updates are necessary to maintain detection performance in dynamic IoT environments where traffic patterns and attack behaviors evolve. However, retraining imposes significant computational and communication costs, particularly in large-scale or distributed deployments. Excessive retraining may overload system resources, whereas infrequent updates may increase vulnerability to concept drift.

Federated and distributed learning systems reduce raw data transfer but introduce challenges related to synchronization, device heterogeneity, and convergence stability [60]. Despite these trade-offs, many studies assume fixed retraining schedules and do not evaluate their operational cost or impact on real-time performance [55]. This gap between adaptive learning theory and deployment realities represents a persistent weakness in current IoT security research.

7.4. Edge–Cloud Orchestration Challenges

Effective real-time IoT security increasingly depends on coordinated edge–cloud orchestration, where decisions regarding model placement, inference execution, aggregation, and retraining are dynamically managed across system layers. Poor orchestration may lead to redundant computation, delayed response, and inconsistent detection performance across distributed nodes [56].

In practice, orchestration must account for device heterogeneity, variable workloads, intermittent connectivity, and evolving threat conditions. However, many ML-based IDS architectures treat edge and cloud components as independent entities rather than as elements of a unified security pipeline. Systems research emphasizes the importance of deployment-aware orchestration to achieve scalable and reliable real-time analytics in IoT environments [61]. The limited focus on orchestration strategies in current IDS proposals constitutes a significant barrier to operational adoption.

7.5. Deployment Evaluation Framework and Evidence Mapping

To systematically assess deployment realism, a deployment evaluation framework was applied to the 38 included studies. Each study was annotated against four measurable criteria derived from edge-IoT operational constraints.

Deployment Criteria Defined

- D1: Deployment layer specified (Edge / Fog / Cloud / Not specified)
- D2: Latency budget discussed or measured
- D3: Resource constraints evaluated (CPU, memory, or energy)
- D4: Real-time feasibility explicitly claimed or validated

Table 6: Deployment Criteria Coverage Across Included Studies ($n = 38$)

Deployment Criterion	**Studies Meeting Criterion**	**Percentage (%)**
D1: Deployment layer specified	17	45
D2: Latency budget analyzed	9	24
D3: Resource constraints evaluated	8	21
D4: Real-time feasibility validated	6	16

Although many studies acknowledge deployment constraints conceptually, fewer than one-quarter provide measurable latency or resource evaluations. Explicit mapping of IDS models to edge- or device-level deployment remains limited, supporting concerns regarding the gap between experimental performance and operational feasibility.

8. Adversarial and Privacy Risks: Background vs. Evidence in Reviewed Studies

With machine learning-based intrusion detection systems (IDSs) increasingly integrated into IoT security architectures, the security of the learning models themselves has become an important and often underexamined consideration. Much of the literature assumes benign training and inference conditions and does not account for the reality that ML models deployed in IoT environments are directly vulnerable to adversarial manipulation, data poisoning, evasion, or privacy breaches. These threats undermine reported detection performance and introduce significant risks in real-world deployments.

8.1. Adversarial Attacks on ML-Based IoT IDS

Adversarial attacks aim to compromise learning systems by degrading detection performance or inducing systematic misclassification. In IoT environments, such attacks are particularly effective due to distributed data sources, limited supervision, and automated decision-making processes.

Poisoning attacks target the training or update phase by injecting malicious or mislabeled data into the learning pipeline. In adaptive IoT IDSs, where models are periodically retrained or updated using live traffic, poisoning can progressively bias decision boundaries, resulting in sustained misclassification of malicious traffic. This threat is amplified in distributed and federated learning environments, where compromised devices may contribute poisoned updates without centralized validation [62]. In contrast, evasion attacks manipulate input traffic during inference to avoid detection. Attackers exploit weaknesses in feature representations or decision thresholds by crafting traffic patterns that appear benign. Empirical evidence indicates that IDS models trained on outdated datasets are particularly susceptible to evasion strategies, raising concerns about the reliability of accuracy-centric evaluations [4].

8.2. Privacy Leakage and Inference Risks

Model inversion and membership inference attacks enable adversaries to determine whether specific data samples were included in training or to infer sensitive attributes from model outputs. Even federated learning, commonly described as privacy-preserving, remains vulnerable to such attacks unless appropriate safeguards are implemented [37]. These risks are further amplified in heterogeneous IoT environments, where partially trusted participants constrain the practical privacy guarantees of collaborative learning frameworks [38].

8.3. Defense Strategies and Practical Limitations

Various defense mechanisms have been proposed to mitigate adversarial and privacy risks, including adversarial training, robust optimization, secure aggregation, and privacy-preserving update mechanisms. However, many of these defenses introduce substantial computational and communication overhead, conflicting with the resource limitations of IoT devices.

Adversarial training enhances robustness against evasion but requires repeated retraining with more complex models, which may be impractical in real-time IoT systems [62]. Similarly, cryptographic approaches for securing federated learning updates reduce privacy leakage but significantly increase communication and computation costs, limiting scalability in large IoT deployments [23]. Consequently, many proposed defenses remain challenging to implement outside controlled experimental environments.

8.4. Implications for Adaptive IoT Security

The presence of adversarial and privacy risks fundamentally alters the design considerations of adaptive IoT security systems. As adaptive learning mechanisms become more responsive to evolving threats, they also expand the attack surface to include poisoning and inference attacks. At the same time, real-time deployment constraints restrict the feasibility of computationally intensive defense strategies. These observations indicate that adaptive ML systems for IoT security should be designed with robustness, privacy, and efficiency addressed jointly rather than treated as competing objectives. Without integrating adversarial resilience and privacy protection into the core design of adaptive IDSs, machine learning may become a vulnerability rather than a solution within IoT security architectures.

8.5. Evidence of Adversarial Evaluation in Included Studies

Although adversarial machine learning threats such as poisoning, evasion, and privacy inference are well established in the broader literature, their empirical evaluation within the included IoT security studies remains limited.

Table 7: Adversarial Evaluation Coverage in Included Studies ($n = 38$)

Adversarial Aspect Evaluated	**Number of Studies**	**Percentage (%)**
Evasion attacks tested	7	18
Poisoning attacks tested	3	8
Privacy leakage / inference evaluated	2	5
No adversarial evaluation	30	79

The majority of included studies discuss adversarial risks only at a conceptual level. Explicit adversarial testing—particularly poisoning resistance and privacy leakage assessment—remains uncommon, indicating a significant evidence gap between theoretical threat models and experimental validation.

9. Comparative Analysis of Existing Studies

Although the primary focus of this review is on adaptive machine learning approaches, several static (non-adaptive) intrusion detection studies were intentionally retained in Table 8 as baseline comparators. These studies represent widely used, high-performing static ML and DL models that are frequently referenced or extended by adaptive approaches. Their inclusion enables clearer comparison of adaptation benefits, deployment trade-offs, and performance limitations relative to non-adaptive baselines.

Static-only studies are therefore not treated as evidence of adaptation mechanisms, but as contextual benchmarks against which adaptive methods are evaluated. This section presents a structured comparative analysis of representative machine learning-based IoT security studies. Instead of reporting accuracy values in isolation, the comparison emphasizes adaptation mechanisms, dataset quality, deployment feasibility, resource consumption, and documented limitations. The revised table supports a comprehensive evaluation of current techniques and highlights research gaps that inform future directions.

Table 8: Comparative Analysis of Adaptive and Static Baseline ML-Based IoT Security Studies

Study	**Target Threat / Domain**	**ML Technique**	**Adaptation Type**	**Dataset(s)**	**Dataset Quality Issues**	**Reported Metrics**	**Resource / Deployment Considerations**	**Key Limitations**
Mahdi et al. (2024)	DDoS in IoT networks	Traditional ML (RF, SVM)	Static	CIC-DDoS2019	Synthetic traffic; class imbalance	Accuracy, DR	Evaluated offline; no latency analysis	No real-time validation; static model
Bhayo et al. (2022)	IoT intrusion detection	Hybrid ML (FS + classifier)	Incremental retraining	CIC-IDS-based IoT data	Limited device diversity	Accuracy, F1-score	Moderate complexity; edge feasibility not analyzed	Retraining cost not evaluated
Abdallah et al. (2022)	DDoS detection	Hybrid ML	Static	Custom IoT traffic	Synthetic environment	Accuracy, precision	No deployment discussion	Scalability unclear
Verma & Ranga (2023)	IoT IDS	Ensemble learning	Static	Benchmark IoT datasets	Dataset aging ignored	Accuracy, recall	Increased inference overhead	High complexity for constrained devices
Kim et al. (2022)	IoT intrusion detection	Federated learning	Federated adaptation	Distributed IoT traffic	Non-IID data	Accuracy, convergence	Communication cost analyzed	Client drift affects stability
Ioannou et al. (2024)	Medical IoT security	Federated ML	Federated + adaptive	Healthcare IoT dataset	Domain-specific bias	Accuracy, energy usage	Edge–cloud deployment considered	Privacy leakage not fully addressed
Hassanien et al. (2024)	IoT IDS	AOA-based feature selection + ML	Static	Benchmark IDS datasets	Imbalance partially addressed	Accuracy, latency	Reduced feature cost	Generalization not tested
Roy et al. (2022)	IoT IDS	Deep learning (LSTM)	Static	IoT traffic traces	Overfitting risk	Accuracy, loss	High training cost	Poor explainability
Zhang et al. (2023)	IoT IDS	Explainable DL	Static	Benchmark IoT datasets	Limited attack diversity	Accuracy, explanation fidelity	Not evaluated under real-time constraints	Scalability unclear
Arisdakessian et al. (2024)	IoT intrusion detection	Deep learning	Drift-aware retraining	IoT datasets	Dataset aging discussed	Accuracy, robustness	Training overhead high	Resource constraints ignored

10. Research Challenges and Future Directions

Despite substantial progress in machine learning for IoT security, the preceding analysis reveals persistent research challenges that limit long-term robustness and real-world applicability. Addressing these issues requires shifting from accuracy-centric evaluation toward resilient, adaptive, and deployment-aware security architectures.

10.1. Lightweight and Resource-Aware Adaptive Intrusion Detection

A major challenge is the design of adaptive intrusion detection systems (IDSs) that operate effectively under strict resource constraints of IoT devices and edge nodes. Although adaptive and deep learning models can enhance detection capability, they often incur high computational and energy cost. Future research should prioritize lightweight adaptation strategies, including selective retraining, feature-efficient modeling, and hardware-aware optimization, to support sustainable real-time deployment across heterogeneous IoT platforms.

10.2. Continual, Self-Supervised, and Drift-Aware Learning

IoT environments are characterized by evolving traffic patterns and concept drift, rendering static training paradigms insufficient. While incremental and federated learning have received attention, continual and self-supervised learning remain underexplored in IoT security. Future work should focus on drift-aware mechanisms capable of autonomously detecting distributional changes and triggering adaptation with minimal reliance on labeled data and limited retraining overhead.

10.3. Robustness Against Adversarial and Poisoning Attacks

The vulnerability of ML models themselves constitutes a critical weakness in IoT security architectures. As discussed in Section 8, adaptive and federated systems expand the attack surface for poisoning, evasion, and inference attacks. Future research should aim to develop robustness-by-design adaptive models, integrating adversarial resilience during training, updating, and deployment phases rather than treating it as an afterthought. Achieving robustness without incurring prohibitive computational cost remains an open challenge.

11. Conclusions

This review provides a systematic and critical analysis of adaptive machine learning methods for securing IoT environments, addressing limitations identified in previous surveys. Rather than cataloging algorithms or presenting detection accuracy as an isolated metric, this study emphasizes adaptation mechanisms, dataset quality, evaluation metrics, real-time feasibility, and adversarial robustness.

Using a unified taxonomy and thematic analysis, the review demonstrates that while machine learning offers powerful capabilities for IoT intrusion detection, high experimental accuracy does not guarantee real-world performance. Operational viability is strongly influenced by dataset bias, insufficient performance metrics, resource limitations, and deployment architecture constraints. Moreover, emerging adaptive and federated learning paradigms introduce new privacy and adversarial risks that must be addressed comprehensively.

Adaptive machine learning is essential for protecting dynamic IoT environments; however, effective implementation requires joint consideration of robustness, efficiency, interpretability, and system-level deployment constraints. This review synthesizes current evidence and outlines key research challenges and future directions to guide the development of scalable, reliable, and practically deployable IoT security solutions.

Declaration of Competing Interests

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.

Funding Declaration

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Ethics Approval and Consent to Participate

This article is a review study based exclusively on previously published literature. It does not involve human participants, animals, or confidential data requiring ethical approval or informed consent.

Data Availability Statement

Data sharing is not applicable to this article as no new data were created or analyzed. All information synthesized in this review is derived from publicly available published studies cited in the reference list.

AI Usage Disclosure

The authors used an AI-based language assistance tool for minor grammatical refinement and formatting support. The scientific content, analysis, interpretation, and conclusions were independently developed, verified, and approved by the authors.

Author Contributions

Jayashri Jayesh Patil: Conceptualization, Methodology, Literature Search, Data Curation, Formal Analysis, Writing – Original Draft, Visualization; **Dr. Ramkumar Solanki**: Supervision, Validation, Writing – Review and Editing, Project Administration.

References

[1] M. Gelgi and M. Çelik, "A systematic literature review of IoT botnet DDoS attacks," *Sensors*, vol. 24, no. 3, pp. 1–29, 2024.

[2] F. Alwahedi, A. Aldhaheri, M. A. Ferrag, A. Battah, and N. Tihanyi, "Machine learning techniques for IoT security: Current research and future vision with generative AI and large language models," *Internet of Things and Cyber-Physical Systems*, vol. 4, pp. 167–185, 2024.

[3] R. Liu, J. Shi, X. Chen, and C. Lu, "Network anomaly detection and security defense technology based on machine learning: A review," *Computers & Electrical Engineering*, vol. 119, p. 109581, 2024.

[4] M. M. Rahman, M. S. Hossain, and M. M. Hassan, "Intrusion detection systems for IoT networks: Recent advances and challenges," *Journal of Network and Computer Applications*, vol. 226, p. 103927, 2024.

[5] C. Ni and S. Li, "Machine learning-enabled industrial IoT security: Challenges, trends, and solutions," *Journal of Industrial Information Integration*, vol. 38, p. 100549, 2024.

[6] E. C. Pinto Neto, S. Dadkhah, S. Sadeghi, H. Molyneaux, and A. A. Ghorbani, "A review of machine learning–based IoT security in healthcare: A dataset perspective," *Computer Communications*, vol. 213, pp. 61–77, 2024.

[7] F. Hinder, A. Artelt, B. Hammer, and M. Biehl, "One or two things we know about concept drift," *Frontiers in Artificial Intelligence*, vol. 7, p. 1296484, 2024.

[8] A. L. Suárez-Cetrulo, A. Bifet, and J. Calvo-Zaragoza, "A survey on machine learning for recurring concept drifting data streams," *Applied Soft Computing*, vol. 132, p. 109890, 2023.

[9] Z. Mahdi, N. Abdalhussien, N. Mahmood, and R. Zaki, "Detection of real-time distributed denial-of-service (DDoS) attacks on internet of things (IoT) networks using machine learning algorithms," *Computers, Materials & Continua*, vol. 80, no. 2, pp. 2139–2159, 2024.

[10] I. Ioannou, P. Nagaradjane, P. Angin, P. Balasubramanian, K. J. Kavitha, P. Murugan, and V. Vassiliou, "GEMLIDS-MIOT: A green effective machine learning intrusion detection system based on federated learning for medical IoT network security hardening," *Computer Communications*, vol. 218, pp. 209–239, 2024.

[11] I. Khamassi, M. Sayed-Mouchaweh, M. Hammami, and K. Ghédira, "Concept drift detection and adaptation with incremental learning: A survey," *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, no. 4, pp. 1739–1756, 2022.

[12] J. Lu, A. Liu, F. Dong, F. Gu, J. Gama, and G. Zhang, "Learning under concept drift: A review," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 2, pp. 1242–1260, 2023.

[13] J. Gama, I. Żliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Computing Surveys*, vol. 46, p. 44, Apr. 2014.

[14] Y. Ren, *Interactive Causality Enabled Adaptive Machine Learning*. PhD thesis, UC Irvine, 2023. ProQuest ID: Ren_uci_0030D_18546. Merritt ID: ark:/13030/m5kn0cfk. Retrieved from https://escholarship.org/uc/item/3dj43270.

[15] A. Bifet, R. Gavalda, and G. Holmes, "Machine learning for evolving data streams," *ACM Computing Surveys*, vol. 54, no. 8, pp. 1–36, 2022.

[16] P. Kairouz, H. B. McMahan, *et al.*, "Advances and open problems in federated learning," *Foundations and Trends® in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2023.

[17] J. Zhang, B. Chen, and Y. Wang, "Adaptive federated learning for edge-enabled IoT systems," *IEEE Transactions on Mobile Computing*, 2024.

[18] L. Jing and Y. Tian, "Self-supervised learning for representation learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 4, pp. 4143–4166, 2023.

[19] Z. Zhang and M. R. Sabuncu, "Generalized cross entropy loss for training deep neural networks with noisy labels," *IEEE Transactions on Neural Networks and Learning Systems*, 2023.

[20] V. Losing, B. Hammer, and H. Wersing, "Choosing the best algorithm for an evolving problem," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 33, no. 1, pp. 1–15, 2022.

[21] V. Mothukuri, S. Pouriyeh, A. Dehghantanha, R. M. Parizi, G. Srivastava, and M. Shafiq, "A survey on security and privacy of federated learning," *Future Generation Computer Systems*, vol. 115, pp. 619–640, 2022.

[22] T. T. Nguyen, V. J. Reddi, and K. R. Chowdhury, "Graph neural networks for network intrusion detection: A survey," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 2, pp. 1241–1267, 2022.

[23] T. Li, S. Hu, A. Beirami, and V. Smith, "Fair resource allocation in federated learning," *Proceedings of the IEEE*, vol. 111, no. 3, pp. 352–370, 2023.

[24] I. Ullah, Q. H. Mahmoud, and W. Alasmary, "A comparative analysis of machine learning and deep learning approaches for intrusion detection in internet of things networks," *IEEE Access*, vol. 10, pp. 100456–100474, 2022.

[25] M. A. Ferrag, L. Shu, H. Djallel, and L. Maglaras, "Deep learning-based intrusion detection for internet of things: A comprehensive survey," *Computer Networks*, vol. 230, p. 109806, 2023.

[26] J. Bhayo, I. A. Hameed, and A. Ahmed, "Hybrid machine learning models for intrusion detection in internet of things networks," *Future Generation Computer Systems*, vol. 129, pp. 63–78, 2022.

[27] A. Verma and V. Ranga, "Ensemble learning based intrusion detection systems for internet of things," *Journal of Network and Computer Applications*, vol. 209, p. 103531, 2023.

[28] S. M. Alqahtani, A. Alzahrani, and H. Aljuaid, "Lightweight intrusion detection for resource-constrained internet of things devices," *IEEE Internet of Things Journal*, vol. 10, no. 9, pp. 7862–7875, 2023.

[29] N. Shone, T. N. Ngoc, V. D. Phai, and Q. Shi, "A lightweight deep learning approach for intrusion detection in internet of things environments," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 6, no. 3, pp. 552–565, 2022.

[30] S. Raza, L. Wallgren, and T. Voigt, *Security and Privacy in Low-Power Internet of Things Systems*. Cambridge University Press, 2023.

[31] Y. Zhou, G. Cheng, and X. Xu, "Graph-based intrusion detection for internet of things networks using graph neural networks," *IEEE Internet of Things Journal*, vol. 10, no. 15, pp. 13489–13502, 2023.

[32] Y. Chen, *Graph-based traffic modeling for intrusion detection in Internet of Things networks*. PhD thesis, Tsinghua University, 2024. (Doctoral dissertation).

[33] M. Abdallah, H. Hijazi, and A. Awad, "Hybrid machine learning models for DDoS detection in internet of things networks," *IEEE Systems Journal*, vol. 16, no. 4, pp. 6432–6443, 2022.

[34] A. S. Alqahtani and M. A. Babar, "Ensemble-based intrusion detection for DDoS attacks in IoT environments," *Computers & Security*, vol. 123, p. 102951, 2023.

[35] M. Hassan, *Hybrid and ensemble learning techniques for distributed denial-of-service detection in IoT networks*. PhD thesis, University of Manchester, 2023. (Doctoral dissertation).

[36] Y. Kim, J. Park, and M. Bennis, "Federated intrusion detection in IoT networks under non-IID data," *IEEE Transactions on Network and Service Management*, vol. 19, no. 3, pp. 3194–3208, 2022.

[37] M. Nasr, R. Shokri, and A. Houmansadr, "Comprehensive privacy analysis of federated learning," in *Proceedings of the IEEE Symposium on Security and Privacy (SP)*, IEEE, 2022.

[38] M. Alazab, *Privacy-preserving federated intrusion detection for large-scale IoT systems*. PhD thesis, Deakin University, 2024. (Doctoral dissertation). Deakin University, Australia.

[39] W. Elmasry, A. Akbulut, and A. H. Zaim, "Feature selection for intrusion detection systems using metaheuristic optimization," *Expert Systems with Applications*, vol. 195, p. 116567, 2022.

[40] G. Kaur and P. Singh, "Energy-efficient intrusion detection in internet of things using feature optimization techniques," *Sustainable Computing: Informatics and Systems*, vol. 39, p. 100855, 2023.

[41] A. E. Hassanien, M. Kilany, M. Abd Elaziz, and A. A. Ewees, "Arithmetic optimization algorithm for feature selection in cybersecurity applications," *Applied Soft Computing*, vol. 148, p. 110863, 2024.

[42] X.-S. Yang, *Nature-inspired Optimization Algorithms*. Elsevier, 2 ed., 2023.

[43] S. S. Roy, A. Mallik, and S. Das, "Deep learning–based intrusion detection in internet of things networks: A temporal analysis," *IEEE Transactions on Industrial Informatics*, vol. 18, no. 6, pp. 4098–4107, 2022.

[44] C. Arisdakessian, O. A. Wahab, and A. Mourad, "Overfitting-aware deep intrusion detection for IoT networks," *Computer Networks*, vol. 242, p. 110181, 2024.

[45] Y. Zhang, X. Chen, and J. Li, "Explainable deep learning for intrusion detection in internet of things systems," *IEEE Internet of Things Journal*, vol. 10, no. 11, pp. 9634–9646, 2023.

[46] R. Kumar, *Deep learning architectures for scalable and explainable Internet of Things intrusion detection*. PhD thesis, Indian Institute of Technology Delhi, 2023. (Doctoral dissertation). Indian Institute of Technology Delhi, India.

[47] S. García, M. Grill, J. Stiborek, and A. Zunino, "An empirical comparison of botnet detection methods," *Computers & Security*, vol. 45, pp. 100–123, 2020.

[48] A. H. Lashkari, G. Draper-Gil, M. S. I. Mamun, and A. A. Ghorbani, "CIC-IoT-2023: A realistic dataset for IoT intrusion detection," *IEEE Access*, vol. 11, pp. 341–356, 2023.

[49] F. Ullah, F. Al-Turjman, and L. Mostarda, "Security datasets and evaluation challenges in healthcare internet of things systems," *IEEE Journal of Biomedical and Health Informatics*, vol. 27, no. 9, pp. 4521–4535, 2023.

[50] B. Krawczyk, "Learning from imbalanced data: Open challenges and future directions," *Progress in Artificial Intelligence*, vol. 11, pp. 221–232, 2022.

[51] M. Ring, S. Wunderlich, D. Grüdl, D. Landes, and A. Hotho, "Flow-based benchmark datasets for intrusion detection," *Computers & Security*, vol. 114, p. 102597, 2022.

[52] F. Pacheco and R. Fernandes, "Realistic traffic generation and dataset aging effects in network intrusion detection," *Computer Networks*, vol. 228, p. 109790, 2023.

[53] D. Chicco and G. Jurman, "The advantages of the Matthews correlation coefficient over F1 score and accuracy in binary classification evaluation," *BMC Genomics*, vol. 24, p. 173, 2023.

[54] S. Wang, X. Zhang, Y. Zhang, and L. Wang, "Energy-efficient machine learning inference at the edge for internet of things applications," *IEEE Transactions on Green Communications and Networking*, vol. 7, no. 2, pp. 873–885, 2023.

[55] A. Alshamrani, *Evaluation frameworks for real-time intrusion detection in Internet of Things systems*. PhD thesis, King Saud University, 2024. (Doctoral dissertation). King Saud University, Saudi Arabia.

[56] M. Satyanarayanan, "The emergence of edge computing," *Computer*, vol. 55, no. 1, pp. 30–39, 2022.

[57] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," *IEEE Internet of Things Journal*, vol. 9, no. 2, pp. 1353–1370, 2022.

[58] J. Hong, Y. Diao, and Z. Zhang, "Fog computing for real-time IoT analytics: Architecture and challenges," *Future Generation Computer Systems*, vol. 139, pp. 1–14, 2023.

[59] X. Xu, W. Zhang, X. Liu, and R. Buyya, "A taxonomy of resource management in edge computing," *ACM Computing Surveys*, vol. 55, no. 3, p. 63, 2022.

[60] C. Zhang, Y. Xie, H. Bai, B. Yu, W. Gao, and P. S. Yu, "A survey on federated learning," *Knowledge and Information Systems*, vol. 63, no. 1, pp. 1–47, 2022.

[61] B. Varghese and R. Buyya, "Next-generation cloud computing: New trends and research directions," *Future Generation Computer Systems*, vol. 134, pp. 289–301, 2023.

[62] S. Alkadi, N. Moustafa, B. Turnbull, and K.-K. R. Choo, "Better safe than never: A survey on adversarial machine learning," *Applied Sciences*, vol. 13, no. 2, pp. 1–36, 2023.

Appendix A: Database-Specific Search Queries and Filters

This appendix documents the complete database-specific search strategy used to identify relevant literature for this systematic review. All searches were conducted on 12 January 2025 and covered publications from 1 January 2019 to 31 December 2024.

Appendix A.1 IEEE Xplore

Search fields:

- Metadata (Title, Abstract, Index Terms)

Query string:

```
("Internet of Things" OR IoT)
AND
("intrusion detection" OR "network security" OR "IoT security")
AND
("machine learning" OR "adaptive learning" OR "online learning"
 OR "incremental learning" OR "federated learning")
```

Filters applied:

- Content type: Journals and Conferences
- Language: English
- Publication years: 2019–2024

Records retrieved: 96

Appendix A.2 Scopus

Search fields:

- Title, Abstract, Keywords

Query string:

```
TITLE-ABS-KEY(
  ("Internet of Things" OR IoT)
  AND
  ("intrusion detection" OR "IoT security" OR "network attack")
  AND
  ("machine learning" OR "adaptive machine learning"
   OR "concept drift" OR "federated learning")
)
```

Filters applied:

- Document type: Article, Conference Paper
- Language: English
- Publication years: 2019–2024
- Subject areas: Computer Science, Engineering

Records retrieved: 88

Appendix A.3 Web of Science (Core Collection)

Search fields:

- Topic (Title, Abstract, Author Keywords, Keywords Plus)

Query string:

```
TS=(
  ("Internet of Things" OR IoT)
  AND
  ("intrusion detection" OR "IoT security")
  AND
  ("machine learning" OR "adaptive learning"
   OR "federated learning" OR "drift-aware")
)
```

Indexes searched:

- SCI-EXPANDED, SSCI

Filters applied:

- Language: English
- Publication years: 2019–2024

Records retrieved: 71

Appendix A.4 ScienceDirect

Search fields:

- Title, Abstract, Keywords

Query string:

```
("Internet of Things" OR IoT)
AND
("intrusion detection" OR "IoT security")
AND
("machine learning" OR "adaptive learning" OR "federated learning")
```

Filters applied:

- Article type: Research articles
- Subject areas: Computer Science, Engineering
- Language: English
- Publication years: 2019–2024

Records retrieved: 57