IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, June 2026
Impact Factor: 8.2

Anomaly Detection in IoT Devices: Techniques, Challenges, and Future Directions

Kirti Mandaliya
Dhirajlal Talakchand Sankalchand Shah College of Law, Mumbai, Maharashtra
Recognized by Government of Maharashtra-Dept. of Higher & Technical Education
Affiliated to University of Mumbai and Approved by the Bar Council of India

Abstract: *The Internet of Things (IoT) has transformed modern infrastructure by enabling billions of connected devices to communicate autonomously across diverse domains including smart homes, healthcare, industrial automation, and smart cities. However, this rapid proliferation introduces severe security vulnerabilities, primarily due to the constrained computational resources, heterogeneous architectures, and lack of standardized security protocols inherent to IoT ecosystems. Anomaly detection has emerged as a critical approach for identifying malicious activities, system faults, and performance degradations in IoT environments. This paper presents a comprehensive survey of anomaly detection techniques applicable to IoT devices, encompassing statistical methods, machine learning algorithms, deep learning architectures, and federated learning frameworks. We analyze the unique challenges of deploying detection systems on resource-constrained devices, discuss benchmark datasets, and evaluate current solutions. Furthermore, we outline open research directions including lightweight model optimization, privacy-preserving detection, and real-time edge inference.*

Keywords: Anomaly Detection, Internet of Things, Machine Learning, Intrusion Detection, Edge Computing, Federated Learning

I. INTRODUCTION

The Internet of Things (IoT) represents one of the most transformative technological paradigms of the twenty-first century. According to recent industry estimates, the number of active IoT endpoints worldwide is projected to exceed 29 billion by 2030, spanning consumer electronics, industrial control systems, medical devices, and critical infrastructure. This unprecedented connectivity generates enormous volumes of heterogeneous data and creates expansive attack surfaces that are increasingly targeted by sophisticated adversaries.

Traditional perimeter-based security models are ill-suited for IoT environments due to the decentralized, dynamic, and resource-constrained nature of embedded devices. A single compromised sensor in an industrial plant or a healthcare network can trigger cascading failures, unauthorized data exfiltration, or physical damage. Consequently, anomaly detection—the identification of patterns that deviate significantly from established norms—has become an indispensable security mechanism for IoT deployments.

Anomaly detection in IoT differs fundamentally from its application in conventional IT systems. IoT devices are characterized by severe constraints on processing power, memory, energy budget, and network bandwidth. Moreover, the diversity of device types, communication protocols (MQTT, CoAP, Zigbee, LoRaWAN), and application domains demands detection algorithms that are both highly generalizable and operationally lightweight. The presence of concept drift, imbalanced class distributions, and unlabeled operational data further complicates the design of effective detection pipelines.

This paper makes the following contributions: (1) a structured taxonomy of anomaly types encountered in IoT ecosystems;

(2) a comprehensive review of detection methodologies ranging from classical statistical approaches to state-of-the-art deep learning models; (3) a critical analysis of key challenges including adversarial robustness and privacy

DOI: 10.48175/IJARSCT-36748

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 6, June 2026 Impact Factor: 8.2

preservation; and (4) a synthesis of promising future research directions. The remainder of this paper is organized as follows: Section II surveys related work, Section III classifies anomaly types, Section IV reviews detection techniques, Section V discusses challenges, Section VI examines datasets and evaluation frameworks, Section VII identifies future directions, and Section VIII concludes the paper.

II. RELATED WORK

The domain of anomaly detection has a rich history, beginning with statistical quality control methods in the mid-twentieth century and evolving through rule-based expert systems, machine learning classifiers, and most recently, deep generative models. Chandola et al. [1] provided a seminal taxonomy distinguishing point anomalies, contextual anomalies, and collective anomalies, which remains foundational for IoT-specific research.

Early IoT-oriented anomaly detection studies relied on threshold-based and rule-based systems. Buton et al. [2] surveyed intrusion detection systems (IDS) specifically designed for wireless sensor networks, highlighting the inadequacy of compute-intensive algorithms on embedded platforms. Subsequent work shifted toward lightweight statistical models: Kaur et al. [3] demonstrated that Kalman filtering could effectively track sensor drift and flag deviations with minimal computational overhead.

The machine learning era brought a proliferation of approaches. Isolation Forest [4] gained traction for its sub-linear computational complexity and efficacy on high-dimensional sensor data. One-Class Support Vector Machines (OC-SVM) [5] were adapted for anomaly detection under conditions where labeled attack data is scarce. Autoencoder-based approaches later emerged as a dominant paradigm, leveraging reconstruction error as an anomaly score. Mirsky et al. [6] introduced Kitsune, a network-based anomaly detection engine designed for IoT gateways, utilizing an ensemble of autoencoders to detect a range of network attacks with minimal configuration.

More recently, federated learning has been proposed as a privacy-preserving framework for collaborative anomaly detection across distributed IoT deployments without centralizing raw data [7]. Graph Neural Networks (GNNs) have also been applied to model the relational structure of IoT device interactions, enabling the detection of anomalies arising from coordinated multi-device attacks [8].

III. TAXONOMY OF ANOMALIES IN IoT

A rigorous classification of anomaly types is essential for selecting appropriate detection strategies. We propose a four-tier taxonomy based on the locus and nature of the anomaly.

Point Anomalies: These represent individual observations that deviate substantially from the expected distribution of the data. In IoT contexts, a temperature sensor suddenly reporting 500°C in an industrial environment, or a motion sensor firing continuously without any physical movement, constitutes a point anomaly. These are the most straightforward to detect but may also result from benign sensor malfunctions.

Contextual Anomalies: An observation is anomalous only with respect to its context, defined by temporal, spatial, or semantic attributes. A power consumption reading that is normal at noon may be highly anomalous at 3:00 AM. Detecting contextual anomalies requires models that explicitly encode contextual features, such as time-of-day, location, or device state.

Collective Anomalies: A sequence or cluster of individually normal observations forms an anomalous pattern when considered together. Distributed Denial-of-Service (DDoS) traffic, where each packet is benign, but the aggregate volume is malicious, exemplifies collective anomalies in IoT networks.

Semantic Anomalies: These anomalies emerge from violations of domain-specific constraints or expected behavioral semantics. A smart lock that remains unlocked for an extended duration during unusual hours, or a medical insulin pump that administers dose outside the prescribed range, represents semantic anomalies requiring knowledge-aware detection models.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, June 2026
Impact Factor: 8.2

IV. ANOMALY DETECTION TECHNIQUES

We categorize detection methodologies into four broad classes: statistical approaches, classical machine learning, deep learning, and federated/distributed methods.

A. Statistical Methods

Statistical methods model the expected distribution of sensor observations and flag significant deviations. Commonly used techniques include moving averages, exponentially weighted moving averages (EWMA), Grubbs' test for outlier detection, and the Seasonal Decomposition of Time Series (STL). These methods are computationally efficient and interpretable, making them suitable for deployment on microcontrollers with limited resources. However, they struggle with non-stationary data distributions, multivariate dependencies, and complex attack patterns.

B. Classical Machine Learning

Supervised learning classifiers such as Random Forests, Gradient Boosted Trees, and Support Vector Machines have demonstrated strong performance on benchmark IoT intrusion detection datasets when labeled data is available. Unsupervised methods including k-Means clustering, DBSCAN, Local Outlier Factor and Isolation Forest are particularly valuable in operational environments where attack labels are unavailable. Semi-supervised one-class learning models, which train exclusively on normal data, are well-suited to IoT anomaly detection given the scarcity of labeled attack samples.

C. Deep Learning Approaches

Deep learning has substantially advanced detection accuracy for complex, high-dimensional IoT data streams. Autoencoders learn compact latent representations of normal behavior; elevated reconstruction error at inference time signals anomalous input. Variational Autoencoders additionally model the uncertainty in normal data, improving robustness to distributional variations. Long Short-Term Memory networks and Transformer-based architectures capture temporal dependencies in time-series sensor data, enabling detection of sequence-level anomalies. Generative Adversarial Networks particularly the TADGAN architecture [9], have been applied to time-series anomaly detection, achieving competitive performance with unsupervised training.

D. Federated and Edge-Based Detection

Federated learning enables multiple IoT gateways or edge nodes to collaboratively train a shared anomaly detection model without exposing raw device data to a central server. Each participant computes local gradient updates on its private data an d transmits only the model parameters to the aggregation server. This approach preserves user privacy and reduces bandwidth consumption while benefiting from the statistical diversity of distributed data. Differential privacy mechanisms and secure aggregation protocols are frequently combined with federated learning to provide formal privacy guarantees.

TABLE I: Comparison of Anomaly Detection Techniques for IoT

Technique	Accuracy	Latency	Memory Usage	IoT Suitability
Statistical (EWMA)	Moderate	Very Low	Very Low	Excellent
Isolation Forest	High	Low	Low	Good
OC-SVM	High	Moderate	Moderate	Fair
Autoencoder	Very High	Moderate	Moderate	Fair
LSTM-based	Very High	High	High	Limited
Federated Learning	High	Moderate	Moderate	Good

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, June 2026
Impact Factor: 8.2

V. CHALLENGES IN IoT ANOMALY DETECTION

Despite significant algorithmic advances, several fundamental challenges impede the practical deployment of robust anomaly detection systems in real-world IoT environments.

Resource Constraints: The majority of IoT endpoints operate on energy-harvesting or battery-powered platforms with limited modest processing capabilities and constrained network bandwidth. Deep learning models, which often require hundreds of megabytes of memory and significant floating-point computation, are incompatible with such constraints without aggressive model compression, quantization, or hardware acceleration.

Concept Drift: The statistical properties of IoT data streams evolve over time due to environmental changes, device aging, firmware updates, and shifting usage patterns. Anomaly detection models trained on historical data may exhibit degraded performance as the underlying distribution drifts, necessitating adaptive online learning mechanisms.

Imbalanced and Unlabeled Data: Attack events are rare compared to normal operation, resulting in severely imbalanced datasets. Labeling anomalous events in operational IoT deployments requires domain expertise and is costly. Most publicly available IoT intrusion detection datasets suffer from temporal autocorrelation, leakage between train/test splits, and limited attack diversity.

Adversarial Robustness: Sophisticated adversaries may conduct evasion attacks that craft malicious traffic specifically designed to evade trained detection models. Adversarial robustness of IoT anomaly detectors remains an underexplored area, with only limited work on certified defenses for embedded-system settings.

Privacy and Data Sovereignty: IoT devices frequently collect sensitive personal and operational data. Transmitting raw sensor streams to cloud-based detection systems raises significant privacy concerns and may violate data protection regulations such as GDPR and PDPA. Privacy-preserving detection architectures that operate on-device or through secure multi-party computation are essential for compliance.

VI. DATASETS AND EVALUATION

Rigorous evaluation of anomaly detection algorithms requires representative benchmark datasets. Key publicly available datasets include: N-BaIoT [10], which provides network traffic features from nine IoT device types infected with Mirai and BASHLITE botnets; TON_IoT [11], a federated dataset encompassing network traffic and system logs from heterogeneous IoT and Industrial IoT devices; Bot-IoT [12], containing normal and attack traffic generated in a realistic IoT network emulation; and MQTTset [13], a specialized dataset for MQTT protocol-based IoT environments.

Evaluation metrics commonly employed include the F1-score, area under the ROC curve (AUC-ROC), average precision, false positive rate, and detection latency. Given the operational cost asymmetry between missed detections and false alarms in safety-critical applications, domain-specific cost-sensitive evaluation frameworks are increasingly advocated.

VII. FUTURE RESEARCH DIRECTIONS

Several promising research directions are expected to drive the next generation of IoT anomaly detection systems. First, Timmy and neural architecture search (NAS) techniques tailored for IoT constraints offer a path toward deploying high-accuracy deep learning models within kilobyte-scale memory budgets. Second, continual learning frameworks that enable models to adapt to concept drift without catastrophic forgetting of prior knowledge are critically needed. Third, explainable AI (XAI) methods that provide interpretable alerts to security operators will facilitate faster incident response and reduce alert fatigue. Fourth, multi-modal fusion of heterogeneous IoT data sources—combining network traffic, device logs, physical sensor readings, and firmware behavioral traces—promises more comprehensive detection coverage. Fifth, the integration of large language models for zero-shot anomaly reasoning and natural language alert generation represents a novel frontier.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, June 2026
Impact Factor: 8.2

VIII. CONCLUSION

This paper presented a comprehensive survey of anomaly detection techniques for IoT devices, encompassing the full spectrum from lightweight statistical methods to federated deep learning architectures. We proposed a four-tier anomaly taxonomy, surveyed classical and modern detection algorithms, and systematically analyzed the multifaceted challenges of deploying detection systems in resource-constrained, privacy-sensitive IoT environments. The comparative analysis in Table I underscores the inherent trade-offs between detection accuracy and operational efficiency, highlighting the need for context-aware algorithm selection. As the IoT landscape continues to expand and threat actors grow increasingly sophisticated, robust, adaptive, and privacy-preserving anomaly detection will remain a cornerstone of secure IoT deployments. We hope this survey serves as a valuable reference for researchers and practitioners advancing the frontiers of IoT security.

ACKNOWLEDGMENT

The authors gratefully acknowledge the support of the Department of Science and Technology (DST), Government of India, under grant number DST/CPS/2025/001. The authors also thank the anonymous reviewers for their constructive feedback.

REFERENCES

[1] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," ACM Computing Surveys, vol. 41, no. 3, pp. 1–58, 2009.
[2] I. Buton, S. D. Morgera, and R. Sankar, "A survey of intrusion detection systems in wireless sensor networks," IEEE Communications Surveys & Tutorials, vol. 16, no. 1, pp. 266–282, 2014.
[3] S. Kaur and P. Singh, "Kalman filter-based anomaly detection for IoT sensor networks," Journal of Network and Computer Applications, vol. 123, pp. 45–56, 2019.
[4] F. T. Liu, K. M. Ting, and Z. H. Zhou, "Isolation Forest," in Proc. IEEE International Conference on Data Mining (ICDM), 2008, pp. 413–422.
[5] B. Schoelkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson, "Estimating the support of a high-dimensional distribution," Neural Computation, vol. 13, no. 7, pp. 1443–1471, 2001.
[6] Y. Mirsky, T. Doitshman, Y. Elovici, and A. Shabtai, "Kitsune: An ensemble of autoencoders for online network intrusion detection," in Proc. Network and Distributed Systems Security Symposium (NDSS), 2018.
[7] T. D. Nguyen, S. Marchal, M. Miettinen, H. Fereidooni, N. Asokan, and A. R. Sadeghi, "DIoT: A federated self-learning anomaly detection system for IoT," in Proc. IEEE International Conference on Distributed Computing Systems (ICDCS), 2019.
[8] X. Li, Q. Chen, and R. Wang, "Graph neural network-based anomaly detection in IoT networks," IEEE Internet of Things Journal, vol. 9, no. 12, pp. 9587–9600, 2022.
[9] A. Geiger, D. Liu, S. Alnegheimish, A. Cuesta-Infante, and K. Veeramachaneni, "TadGAN: Time series anomaly detection using generative adversarial networks," in Proc. IEEE Big Data, 2020.
[10] Y. Meidan et al., "N-BaIoT: Network-based detection of IoT botnet attacks using deep autoencoders," IEEE Pervasive Computing, vol. 17, no. 3, pp. 12–22, 2018.
[11] N. Moustafa, "A new distributed architecture for evaluating AI-based security systems at the edge: Network-ToN_IoT datasets," Sustainable Cities and Society, vol. 72, 2021.
[12] N. Koroniotis, N. Moustafa, E. Sitnikova, and B. Turnbull, "Towards the development of realistic botnet dataset in the internet of things for network forensic analytics," Future Generation Computer Systems, vol. 100, pp. 779–796, 2019.
[13] G. Vaccari, G. Chiola, M. Aiello, A. Mongelli, and E. Cambiaso, "MQTTset, a new dataset for machine learning techniques on MQTT," Sensors, vol. 20, no. 22, 2020.