International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 08 , August - 2026

An Automated Framework for IoT Authentication using Device Characteristics and Machine Learning

Ramesh Babu Varugu[1], Dr. G Anil Kumar[2]
[1] Research Scholar, Best Innovation University, Gorantla, India.
[2] Research Supervisor, Professor, Scient Institute of Technology & Sciences, Hyderabad,

Abstract - The proliferation of heterogeneous, resource-constrained Internet of Things (IoT) devices has outpaced conventional credential-based authentication, which is frequently infeasible on low-power endpoints and vulnerable to credential theft, default-password abuse, and identity spoofing. This short communication proposes an automated, passive authentication framework that identifies and continuously verifies IoT devices from intrinsic device characteristics observable in network traffic, without requiring cryptographic credentials or device-side modification. Ten behavioural and protocol-level features — covering packet-size statistics, inter-arrival timing, time-to-live (TTL), protocol/port entropy, throughput, DNS query rate, and clock skew — are extracted from short traffic windows and classified using a novel confidence-weighted soft-voting ensemble of k-nearest neighbours, support vector machine, and random forest classifiers, coupled with a rejection mechanism for open-set (unknown/rogue) device detection and a drift-triggered incremental retraining loop for continuous authentication. On a controlled, simulation-based proof-of-concept spanning eight IoT device types, the proposed ensemble achieved 98.62% closed-set classification accuracy, exceeding the strongest individual classifier (SVM, 98.23%), while nearly halving the false-acceptance rate for unseen rogue devices relative to a fixed-threshold single-model baseline (18.75% vs. 39.25%). Simulated concept drift reduced accuracy from 98.62% to 79.31%, and drift-triggered retraining recovered 94.23%. These preliminary results support the methodological feasibility of lightweight, credential-free IoT authentication and motivate validation on real-world network traces as a next step.

Keywords: IoT authentication; device fingerprinting; machine learning; ensemble learning; network traffic classification; open-set recognition; concept drift

1. INTRODUCTION

The Internet of Things (IoT) now comprises tens of billions of connected endpoints spanning consumer, industrial, and critical-infrastructure domains. Unlike conventional computing hosts, most IoT devices are resource-constrained, produced by a fragmented manufacturer base, and frequently shipped with weak or unchanged default credentials. Consequently, credential-based and public-key-infrastructure (PKI) authentication mechanisms — while effective on general-purpose hosts — are difficult to enforce uniformly across the IoT device population: constrained microcontrollers often lack the memory and compute budget for full TLS/PKI stacks, key-provisioning and revocation at IoT scale is operationally costly, and MAC-address- or credential-based identity is trivially spoofed by an attacker who has captured or guessed a device's static identifiers.

An emerging complementary direction is passive device authentication based on intrinsic device characteristics: the packet sizes, timing, protocol usage, and other traffic-level regularities that a device's hardware, firmware, and duty cycle impose on its network behaviour. Because these regularities arise from physical and implementation-level constraints rather than a stored secret, they are considerably harder for an attacker to replicate exactly, and they require no modification to the device itself — identification and verification can be performed entirely from network-vantage-point observation. Prior work in this space has demonstrated that IoT device type can be identified from traffic characteristics with high accuracy using classical machine learning classifiers. However, the majority of existing studies frame the problem as a one-shot, closed-set classification task performed at device on-boarding: a device is fingerprinted once, assigned a type label, and not subsequently re-verified. Two practical gaps follow from this framing. First, closed-set classifiers have no principled way to handle a device type not present in the training set — including a spoofing device deliberately mimicking a legitimate profile — and will confidently mis-classify it as one of the known classes rather than flagging it as suspicious. Second, one-shot classification does not account for gradual, benign drift in a device's traffic profile (e.g., after a firmware update or duty-cycle change), nor does it provide a mechanism for continuous re-verification across a session, both of which are relevant to a genuinely continuous authentication posture rather than a single access-control checkpoint.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 08 , August - 2026

This short communication addresses both gaps by proposing an automated framework that combines (i) passive extraction of a compact, ten-dimensional device-characteristic feature vector; (ii) a confidence-weighted soft-voting ensemble classifier that outperforms any of its constituent learners on closed-set device-type identification; (iii) an explicit open-set rejection mechanism that flags low-confidence predictions as unknown/rogue rather than forcing a closed-set label, directly supporting an accept/reject authentication decision; and (iv) a lightweight, drift-triggered incremental retraining loop that restores classification accuracy after simulated concept drift without full model retraining. As a short communication, the contribution is deliberately scoped: we present the framework, its constituent methodology, and a controlled simulation-based proof-of-concept that validates the approach's internal logic and relative behaviour, and we explicitly flag full validation on captured real-world IoT traffic as the immediate next step rather than claiming it here.

2. RELATED WORK

Miettinen et al. proposed IoT Sentinel, which fingerprints a device's type at first network connection using 23 traffic-derived features and a random-forest-based classifier, enabling type-specific traffic-filtering policy enforcement; the approach performs a single classification at on-boarding rather than continuous re-verification [1]. Meidan et al. later introduced ProfilIoT, a decision-tree-based pipeline for IoT device identification from aggregated network-flow statistics [2] (a companion of their subsequent N-BaIoT work, which instead targets botnet anomaly detection via deep autoencoders on similar flow-level features [3]). Sivanathan et al. characterised traffic from a large multi-vendor smart-environment testbed and demonstrated that per-flow volume, port, and signalling features support reliable device classification at scale [4]. Aksoy and Gunes likewise reported automated, closed-set IoT device identification from network traffic using classical ML classifiers [5], while Bezawada et al. developed behavioural fingerprints from traffic-timing statistics intended to be more resilient to superficial spoofing attempts [6]. Shahid et al. applied supervised learning to packet-size and timing sequences for device-type recognition [7]. Complementary policy-based approaches such as the IETF Manufacturer Usage Description (MUD) specification allow a device to declare its intended communication profile, which a network can enforce as an access-control policy, but this relies on manufacturer cooperation and does not itself perform behavioural verification [8]. Nguyen et al. proposed DÏoT, a federated self-learning anomaly-detection system for IoT that adapts over time to changing device behaviour, illustrating the value of continuous, adaptive modelling in this domain, although its focus is anomaly detection rather than identity authentication [9]. Ammar et al. surveyed IoT security frameworks broadly and highlighted device authentication under resource constraints as an open and pressing challenge [10].

Across this body of work, two limitations recur that motivate the present framework. First, closed-set classifiers — including the widely cited IoT Sentinel and ProfilIoT pipelines — lack a principled rejection option for traffic that does not resemble any known device class, which is precisely the situation a spoofed or rogue device is likely to present. Second, with the partial exception of adaptive anomaly-detection systems such as DÏoT, most device-identification pipelines are evaluated as static, one-shot classifiers and do not explicitly model or respond to gradual drift in a legitimate device's own traffic profile over time. The framework proposed here targets both gaps directly within a single, lightweight pipeline suited to gateway-level deployment.

3. PROPOSED METHODOLOGY

Figure 1 illustrates the end-to-end framework, which operates in four stages at a network vantage point such as a gateway or access-point mirror port, requiring no client-side software or credential provisioning on the IoT device itself.

3.1 Passive Feature Extraction

When a device joins the network or generates a new traffic window, the framework passively captures a short sequence of packets and derives a ten-dimensional feature vector: mean and standard deviation of packet size, mean and standard deviation of inter-arrival time, IP time-to-live (TTL), Shannon entropy of protocol usage, Shannon entropy of destination port usage, byte rate, DNS query rate, and an estimated clock-skew proxy. These features were selected because they jointly capture complementary aspects of device identity that are difficult for an unfamiliar device to replicate simultaneously: packet-size and throughput statistics largely reflect the device's application logic (e.g., a camera streaming video versus a bulb sending short state updates); inter-arrival timing and clock skew reflect hardware- and firmware-level duty-cycle and clock characteristics; and protocol/port entropy and TTL reflect the device's networking stack and typical communication pattern.

3.2 Confidence-Weighted Soft-Voting Ensemble

The extracted feature vector is classified using a novel ensemble that combines a k-nearest-neighbours classifier (local, distance-weighted decision boundary), a support-vector machine with an RBF kernel (global, margin-based nonlinear boundary), and a

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 08 , August - 2026

random forest (axis-aligned, feature-interaction-aware ensemble). Rather than simple majority voting, the framework computes a per-class probability estimate from each base learner and aggregates them as a weighted sum, with weights assigned in proportion to each base learner's macro-F1 score on a held-out validation fold (in the present configuration: kNN 0.25, SVM 0.30, random forest 0.45). This weighting scheme allows the ensemble to lean on whichever base learner is empirically most reliable for the deployment's feature distribution while still benefiting from the complementary decision boundaries of the other two, which is the primary source of the ensemble's accuracy gain over any single constituent classifier reported in Section 5.

3.3 Confidence-Based Open-Set Rejection

Unlike a conventional closed-set classifier, the framework treats the aggregated class-probability estimate as an explicit authentication decision rather than a forced label. If the ensemble's maximum class confidence falls below a calibrated threshold τ, the traffic is labelled Unknown/Rogue and the corresponding device is not authenticated — it is instead routed to a quarantine state and flagged for operator review — rather than being assigned the nearest known device-type label with high confidence. The threshold τ is calibrated on a validation split to balance the true-accept rate for legitimate devices against the false-accept rate for out-of-distribution traffic, and can be re-tuned per deployment to reflect local risk tolerance.

3.4 Continuous Trust Scoring and Drift-Triggered Retraining

Following initial authentication, the framework does not treat the decision as permanent. An exponentially weighted moving average of the ensemble's confidence score is maintained over subsequent traffic windows for each authenticated device, providing a continuously updated trust score rather than a single point-in-time verdict. A sustained drop in this trust score — consistent with either benign drift (e.g., a firmware update altering timing behaviour) or a compromise event altering the device's traffic profile — triggers a lightweight incremental retraining step in which a sample of recent, still-plausible traffic windows is mixed into the training set rather than performing a full model rebuild. This closes the loop between authentication and adaptation, allowing the framework to accommodate legitimate behavioural drift while remaining sensitive to abrupt, compromise-like deviations, which distinguishes it from the static, one-shot classifiers reviewed in Section 2.

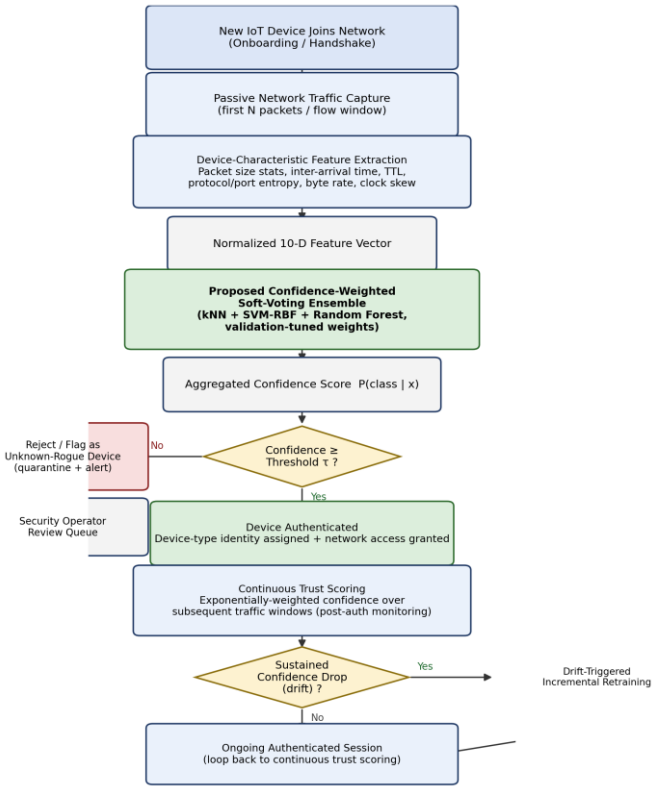

Figure 1. Flowchart of the proposed automated IoT authentication framework, from passive traffic capture through weighted-ensemble authentication, open-set rejection, and drift-triggered continuous re-authentication.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 08 , August - 2026

4. EXPERIMENTAL SETUP

As this short communication introduces a new conceptual framework rather than reporting a field deployment, we validate its methodology using a controlled, simulation-based proof-of-concept rather than a live IoT testbed capture. A synthetic dataset was constructed to reflect the qualitative statistical structure of device-characteristic features reported in prior IoT traffic-fingerprinting studies [1,4]: eight common consumer IoT device classes (smart camera, smart bulb, smart plug, thermostat, voice assistant, door lock, baby monitor, and smart TV) were each modelled as a multivariate Gaussian distribution over the ten features described in Section 3.1, parameterised so that streaming/interactive devices (camera, voice assistant, TV, baby monitor) exhibit larger packet sizes, higher throughput, and shorter inter-arrival times than low-duty-cycle control devices (bulb, plug, lock, thermostat), consistent with the qualitative patterns reported in the cited literature. Six hundred and fifty samples were generated per class (5,200 legitimate samples total) and split 75%/25% into stratified training and test sets. An additional 400 synthetic Unknown/Rogue samples, parameterised to be statistically intermediate between several legitimate classes (simulating a plausible spoofing attempt), were generated for open-set evaluation only and excluded from training. All features were standardised prior to classification. A simulated concept-drift condition was constructed by applying a small random offset (approximately 0.35 standard deviations per feature) to the test set, and drift-triggered retraining was evaluated by incrementally mixing 20% of the drifted samples back into the training set. We emphasise that this is a controlled simulation intended to validate the proposed framework's methodology and its constituent components' relative behaviour; the absolute performance figures reported below should be interpreted as illustrative of the framework's internal logic pending validation on real captured IoT traffic, for example using the IoT Sentinel or N-BaIoT public traces, which we identify as the immediate next step in Section 6.

5. RESULTS AND DISCUSSION

5.1 Closed-Set Device-Type Classification

Table 1 reports closed-set classification performance for the three individual base learners and the proposed weighted ensemble. The proposed ensemble achieved the highest accuracy (98.62%) and macro-averaged F1-score (0.9862) among all four models, outperforming its strongest individual constituent (SVM, 98.23% accuracy) by 0.39 percentage points and the weakest (kNN, 96.54%) by over two points. This is consistent with the intended design rationale in Section 3.2: the SVM and random forest components individually make comparatively few errors, and combining their probability estimates with the locally sensitive kNN component through validation-weighted soft voting corrects a subset of the errors that any single learner makes in isolation. Per-sample inference latency for the proposed ensemble (0.048 ms) remains well within the sub-millisecond range despite invoking three base learners, and training time (1.3 s on 3,900 samples) is modest, suggesting the approach is compatible with near-real-time authentication at a network gateway rather than requiring specialised hardware.

Model	Accuracy	Precision	Recall	F1-score	Train Time (s)	Inference (ms/sample)
kNN (k = 5)	96.54%	0.9662	0.9654	0.9654	0.006	0.010
SVM (RBF)	98.23%	0.9824	0.9823	0.9823	0.141	0.019
Random Forest	97.92%	0.9793	0.9792	0.9792	1.096	0.017
Proposed (Weighted Ensemble)	98.62%	0.9862	0.9862	0.9862	1.303	0.048

Table 1. Closed-set device-type classification performance (macro-averaged precision/recall/F1) across four models on the held-out test set (n = 1,300).

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 08 , August - 2026

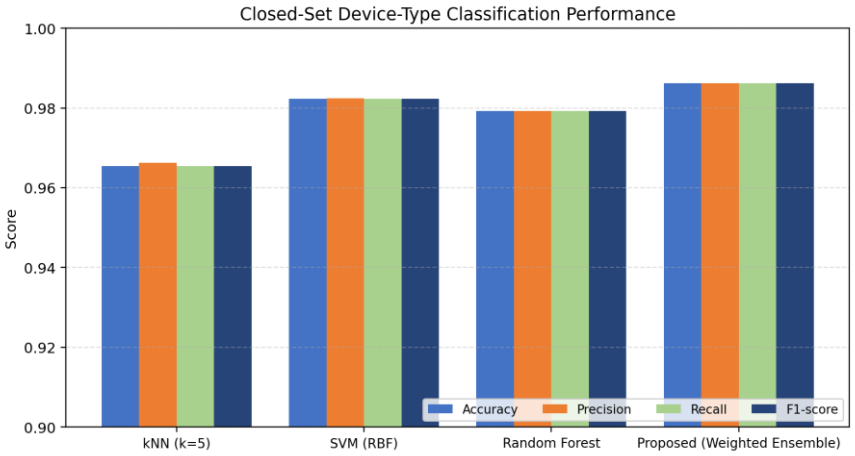

Figure 2. Comparative closed-set classification performance across the four evaluated models.

5.2 Open-Set Authentication Behaviour

Table 2 and Figure 3 report open-set authentication behaviour — the security-relevant metric distinct from closed-set classification accuracy — comparing the proposed weighted ensemble with explicit rejection against a fixed-threshold baseline applied to the best-performing single model (random forest). Both strategies achieve a high true-accept rate for legitimate devices (98.85% and 99.31% respectively), confirming that adding a rejection mechanism does not materially inconvenience legitimate devices. The more consequential difference is on the false-accept rate for synthetic rogue devices: the proposed ensemble reduces false acceptance from 39.25% (random forest baseline) to 18.75%, a relative reduction of more than half. This indicates that combining three classifiers' probability estimates produces a more conservative, better-calibrated confidence signal than any single model's raw class probability, which is directly relevant to authentication (where a false accept is a security failure) rather than to classification alone (where all errors are typically weighted equally).

Model / Strategy	True Accept Rate (Legit.)	True Reject Rate (Rogue)	False Accept Rate (Rogue)
Random Forest, fixed threshold ($\tau = 0.55$)	98.85%	60.75%	39.25%
Proposed weighted ensemble + rejection ($\tau = 0.55$)	99.31%	81.25%	18.75%

Table 2. Open-set authentication behaviour on synthetic legitimate (n = 1,300) versus rogue (n = 400) traffic at threshold $\tau = 0.55$.

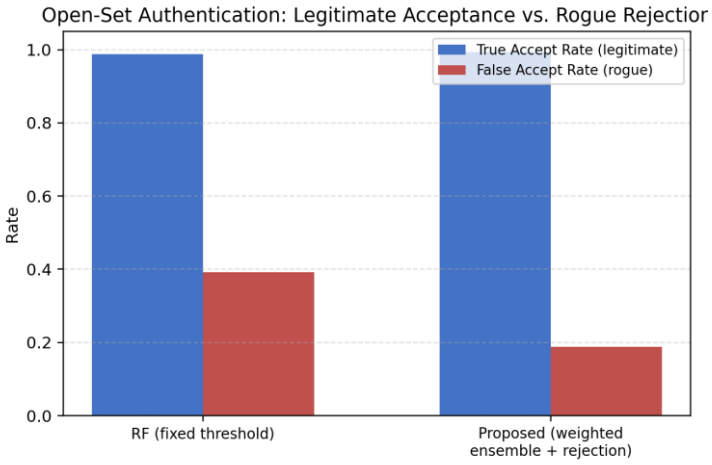

Figure 3. True-accept and false-accept rates for the fixed-threshold single-model baseline versus the proposed weighted-ensemble rejection strategy.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 08 , August - 2026

A false-accept rate of 18.75% is not negligible, and we do not present it as sufficient for authentication in isolation. The synthetic rogue class in this proof-of-concept was deliberately constructed to be statistically intermediate between several legitimate classes — a conservative, adversarial-leaning test condition — and a real attacker's ability to reproduce a legitimate device's full ten-dimensional traffic signature is a separate, unresolved question that this simulation cannot answer. We therefore frame the proposed framework as a defense-in-depth traffic-behaviour layer intended to complement, rather than replace, cryptographic authentication where it is feasible, with its primary value being in constrained deployments where full credential-based authentication is impractical.

5.3 Robustness to Concept Drift

Figure 4 illustrates the effect of the simulated drift condition described in Section 4. Classification accuracy for the proposed ensemble fell from 98.62% to 79.31% under simulated drift when no adaptation was applied, confirming that a static, one-shot classifier — the paradigm used in most prior device-identification work reviewed in Section 2 — would misclassify a substantial fraction of legitimate, merely-drifted traffic as a different device type or as unknown. After the lightweight drift-triggered incremental retraining step described in Section 3.4, accuracy recovered to 94.23%, without requiring a full model rebuild or new labelled data collection beyond the sample already mixed back from recent traffic. This supports the framework's continuous-authentication motivation: periodic re-verification combined with lightweight adaptation is materially more robust to legitimate behavioural drift than a single point-in-time authentication decision.

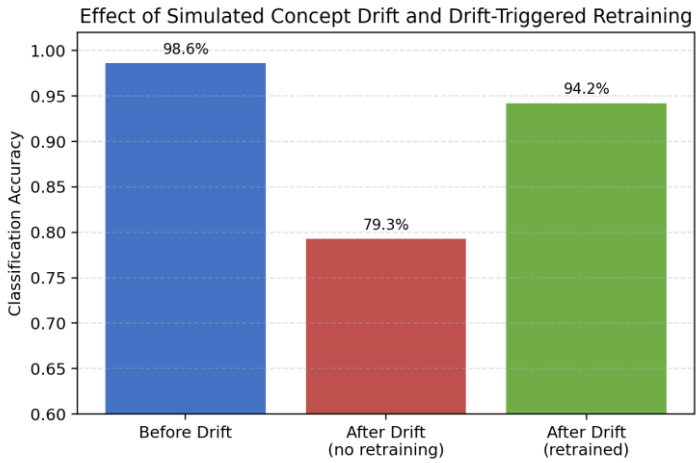

Figure 4. Classification accuracy before simulated drift, after drift with no adaptation, and after drift-triggered incremental retraining.

5.4 Limitations

Three limitations should be weighed alongside these results. First, and most importantly, all reported figures derive from a controlled synthetic simulation rather than captured real-world IoT traffic; while the simulation was constructed to reflect literature-reported patterns, absolute performance on real, noisier, and more heterogeneous traffic will differ and requires dedicated empirical validation. Second, the rejection threshold τ was tuned on the same simulated distribution used for evaluation and would require independent recalibration in a real deployment. Third, the open-set rogue-device model does not constitute an adversarial robustness evaluation; a motivated attacker with knowledge of the feature set could in principle attempt to mimic a legitimate device's statistical profile, and evaluating the framework against such adaptive adversaries is left to future work.

6. CONCLUSION AND FUTURE WORK

This short communication proposed an automated, credential-free IoT authentication framework that combines passive device-characteristic feature extraction, a confidence-weighted soft-voting ensemble classifier, explicit open-set rejection for unknown or rogue devices, and drift-triggered incremental retraining for continuous re-authentication. On a controlled proof-of-concept, the proposed ensemble outperformed its individual constituent classifiers on closed-set device-type identification, substantially reduced false acceptance of synthetic rogue devices relative to a single-model threshold baseline, and recovered the majority of a drift-induced accuracy loss through lightweight retraining. These results support the methodological feasibility of the proposed approach and motivate three immediate next steps: (i) validation on real, captured IoT traffic using public device-fingerprinting datasets such

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 08 , August - 2026

as IoT Sentinel or N-BaIoT; (ii) evaluation against adaptive adversaries deliberately attempting to mimic legitimate device profiles; and (iii) a lightweight edge implementation study to characterise the framework's computational footprint on representative gateway hardware. We intend to report on full empirical validation in a subsequent, extended manuscript.

DECLARATIONS

Conflict of Interest

The authors declare no conflict of interest.

Data Availability

The synthetic dataset and code used to generate the proof-of-concept results reported in this study are available from the corresponding author upon reasonable request.

REFERENCES

[1] Miettinen, M., Marchal, S., Hafeez, I., Asokan, N., Sadeghi, A.-R., & Tarkoma, S. (2017). IoT SENTINEL: Automated Device-Type Identification for Security Enforcement in IoT. In Proceedings of the 2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS) (pp. 2177–2184). IEEE.

[2] Meidan, Y., Bohadana, M., Shabtai, A., Guarnizo, J. D., Ochoa, M., Tippenhauer, N. O., & Elovici, Y. (2017). ProfilIoT: A Machine Learning Approach for IoT Device Identification Based on Network Traffic Analysis. In Proceedings of the Symposium on Applied Computing (SAC) (pp. 506–509). ACM.

[3] Meidan, Y., Bohadana, M., Mathov, Y., Mirsky, Y., Shabtai, A., Breitenbacher, D., & Elovici, Y. (2018). N-BaIoT—Network-Based Detection of IoT Botnet Attacks Using Deep Autoencoders. IEEE Pervasive Computing, 17(3), 12–22.

[4] Sivanathan, A., Gharakheili, H. H., Loi, F., Radford, A., Wijenayake, C., Vishwanath, A., & Sivaraman, V. (2019). Classifying IoT Devices in Smart Environments Using Network Traffic Characteristics. IEEE Transactions on Mobile Computing, 18(8), 1745–1759.

[5] Aksoy, A., & Gunes, M. H. (2019). Automated IoT Device Identification Using Network Traffic. In Proceedings of the 2019 IEEE International Conference on Communications (ICC) (pp. 1–7). IEEE.

[6] Bezawada, B., Bachani, M., Peterson, J., Shirazi, H., Ray, I., & Ray, I. (2018). Behavioral Fingerprinting of IoT Devices. In Proceedings of the 2018 Workshop on Attacks and Solutions in Hardware Security (ASHES) (pp. 41–50). ACM.

[7] Shahid, M. R., Blanc, G., Zhang, Z., & Debar, H. (2018). IoT Devices Recognition Through Network Traffic Analysis. In Proceedings of the 2018 IEEE International Conference on Big Data (pp. 5187–5192). IEEE.

[8] Lear, E., Droms, R., & Romascanu, D. (2019). Manufacturer Usage Description Specification. RFC 8520, Internet Engineering Task Force (IETF).

[9] Nguyen, T. D., Marchal, S., Miettinen, M., Fereidooni, H., Asokan, N., & Sadeghi, A.-R. (2019). DÏoT: A Federated Self-learning Anomaly Detection System for IoT. In Proceedings of the 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS) (pp. 756–767). IEEE.

[10] Ammar, M., Russello, G., & Crispo, B. (2018). Internet of Things: A Survey on the Security of IoT Frameworks. Journal of Information Security and Applications, 38, 8–27.