Exploring Machine Learning Algorithms for User Activity Inference from IoT Network Traffic

Kuai Xu, Yinxin Wan, Xuanli Lin, Feng Wang, Guoliang Xue
Arizona State University
{kuai.xu, ywan28, xlin54, fwang25, xue}@asu.edu

***Abstract*—The availability of ubiquitous and heterogeneous Internet-of-Things (IoT) devices in smart homes and their interactions with users provide a unique opportunity to monitor, understand, recognize, learn, and infer user activities for safety monitoring, connected health, energy saving as well as other disruptive services. Our analysis on IoT network traffic from smart homes with a variety of IoT devices has discovered that user activities often trigger overlapping traffic waves from multiple IoT devices that are deployed near the activities. This insight leads us to adopt wavelet analysis to decompose IoT network traffic in smart homes into low, middle, and high frequency bands that distinguish IoT traffic waves triggered by user activities from background noises such as heartbeat signals between IoT devices and cloud servers. Subsequently, we extract a broad range of traffic features from these IoT traffic waves and explore supervised machine learning (ML) algorithms to classify various user activities with these features. Based on the labelled user activities and IoT network traffic data collected from real smart home environments, our experiments have demonstrated that the ML-based algorithms are able to use IoT network traffic to accurately infer various user activities in smart homes.**

I. Introduction

The last decade has witnessed the rapid and ubiquitous deployment of IoT devices in smart homes, smart buildings, and smart cities, which have created unique opportunities to improve the situation awareness in these environments for safety monitoring and connected health [1], [2]. For example, several recent studies [3], [4], [5] collect and analyze network traffic of IoT devices for recognizing and inferring a variety of user activities in smart homes. These studies on user activity inference (UAI) rely on the accurate extraction of IoT device events [6], [7] from the underlying network traffic, e.g., TCP/IP data packets, as well as the correct temporal sequences of such device events. However, device malfunctions, packet losses, and network latency dynamics [8], [9] often lead to missing or out-of-order IoT device events, which creates substantial challenges for these existing solutions.

To tackle this issue, this paper explores supervised machine learning (ML) algorithms to directly infer user activities in smart homes from IoT network traffic without relying on IoT device events extracted from network traffic. Figure 1 illustrates the overall architecture of our proposed approach. The first step is to collect IoT network traffic via programmable routers in smart homes while recording and labelling user activities [10] with timestamps that trigger heterogeneous IoT devices. Subsequently, we analyze and characterize network traffic of these IoT devices before, during, and after smart home user activities and discover substantial and overlapping *traffic waves* of IoT devices that are strongly correlated with user activities in time and space, i.e., approximate locations in the homes.

Based on our observations on diverse patterns of IoT background traffic and IoT device event traffic triggered by user activities, we adopt a wavelet analysis approach [11], [12], inspired by prior work [13], [14] on signal analysis for traffic anomalies, to learn and extract IoT network traffic with low, middle, and high frequency bands. As IoT network traffic waves triggered by user activities are often captured by the synthesized low frequency signals, we are able to identify the start and end timestamps of traffic waves due to user activities, and extract and enigeer 19 features from these traffic waves. The availability of these IoT network traffic features and the corresponding labelled user activities as ground-truth leads us to explore a suite of supervised ML-based algorithms such as Naive Bayes, random forest, k-nearest neighbors (k-NN), gradient boosting, and support vector machine (SVM), to infer a wide range of user activities in smart homes.

To evaluate the performance of our proposed ML-based solution for inferring user activities, we set up two real smart home environments and collect network traffic from a number of heterogeneous IoT devices. In addition, we identify 23 different user activities which could trigger at least one IoT device deployed in these two homes. Our experiment results show that our proposed ML-based algorithms are able to accurately infer user activities. For example, best overall performing algorithm, Naive Bayes, achieves an average accuracy, precision, recall, F_1 score, and AUC (area under curve) of 0.8783, 0.8801, 0.8783, 0.8779, 0.9498, respectively. The ability of automatically and accurately inferring user activities in smart homes is a crucial step to leverage today's ubiquitous and heterogeneous IoT devices for home safety and security, remote patient monitoring, and independent senior living.

The contributions of this paper are summarized as follows:

- This paper introduces the problem of inferring user activities via directly characterizing and learning network traffic of heterogeneous IoT devices in smart homes.
- This paper explores a suite of supervised ML-based algorithms to infer user activities with a set of features

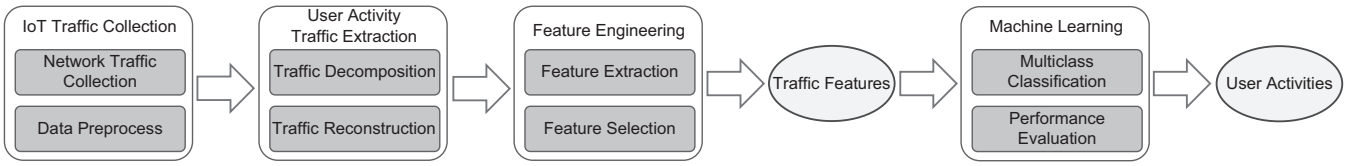

Fig. 1. The overall system architecture of our supervised ML-based approach for inferring user activities in smart homes.

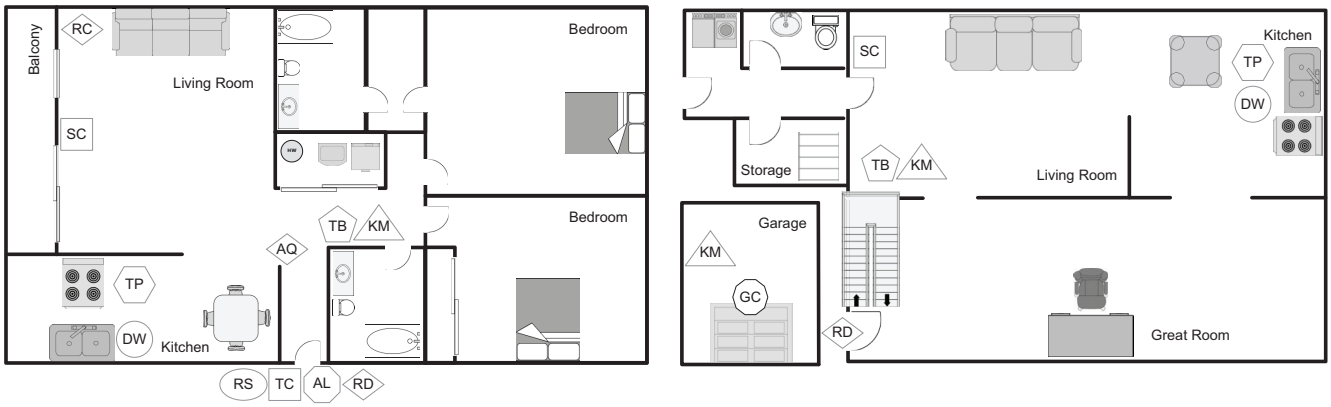

(a) Smart home A (b) Smart home B

Fig. 2. The layouts of IoT device deployment in two real-world smart home experimental environments. The mappings between the abbreviation names and the actual IoT devices are summarized in Table I.

generated from IoT network traffic.

- Our extensive experiments based on data-sets collected from two real smart homes have demonstrated that the ML-based algorithms are able to accurately infer different user activities from IoT network traffic.

II. Research Background

In this section, we first briefly describe the wide deployment of IoT devices in smart homes and present the layout of two smart home experimental environments for this study. Subsequently, we discuss user activities triggering IoT devices in these two smart homes and explain the temporal and spatial correlations between user activities and IoT network traffic.

A. IoT Devices in Smart Homes

TABLE I
THE LIST OF HETEROGENEOUS IOT DEVICES DEPLOYED IN TWO REAL SMART HOME EXPERIMENTAL ENVIRONMENTS FOR THIS STUDY.

Device Type	Device Name	Protocols	Home	Abbr.
Bulb	TP-Link Bulb	WiFi	A & B	TB
Camera	Arlo Q Camera	WiFi	A	AQ
	Reolink Camera	Ethernet	A	RC
Controller & Sensor	MyQ Garage Controller	WiFi	B	GC
Doorbell	Ring Doorbell	WiFi	A & B	RD
Lock	August Lock	WiFi	A	AL
Plug	TP-Link Plug	WiFi	A & B	TP
Sensor	D-Link Water Sensor	WiFi	A & B	DW
	Kangaroo Motion Sensor	WiFi	A & B	KM
	Smart Life Contact Sensor	WiFi	A & B	SC
	Tessan Contact Sensor	WiFi	A & B	TC
Spotlight	Ring Spotlight	WiFi	A	RS

Recent advances in embedded systems, cloud computing, and artificial intelligence have driven the rapid and wide adoption of IoT devices in smart homes for a broad range of innovative and disruptive applications such as energy saving, safety monitoring, and home automation. Figure 2 illustrates the deployment of 11 and 8 heterogeneous IoT devices in our two smart home experimental environments, respectively. These IoT devices provide many benefits and functions for smart home users. For example, an outdoor motion sensor near the front door could automatically notify residents on the presence of guests or delivery drivers.

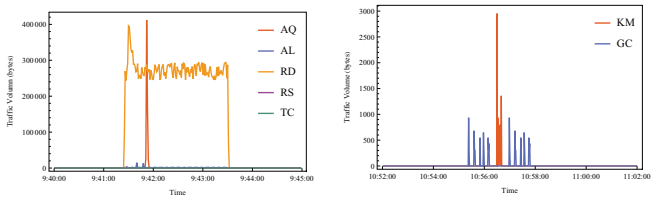

(a) Network traffic from Arlo Q Camera (AQ), August Lock (AL), Ring Doorbell (RD), Ring Spotlight (RS) and Tessan Contact Sensor (TC) which are triggered by a resident entering the front door of the home A

(b) Network traffic from Kangaroo Motion Sensor (KM) and MyQ Garage Controller (GC), which are triggered by a resident driving into the garage of the home B

Fig. 3. The temporal and spatial correlations between user activities and IoT network traffic.

B. User Activities in Smart Homes

Given the dense deployment of diverse IoT devices in smart homes, user activities often trigger multiple IoT devices in

parallel [15], [16], [3], [17]. For example, a resident entering the house through the front door of the smart home A would trigger various IoT devices including the outside IP cameras, the smart lock and the contact sensor. For another instance, a resident driving to the garage in the smart home B would trigger the smart garage controller and the motion sensor. Figure 3 illustrates network traffic of these IoT devices over a 5-minute time window during which these two user activities have happened.

As observed in Figure 3, user activities often simultaneously trigger multiple IoT devices which are deployed near the user activities based on the physical layouts of IoT device deployment in the smart homes. In other words, user activities in smart homes with dense IoT deployment have strong temporal and spatial correlations with network traffic of IoT devices. In addition, the IoT network traffic triggered by different user activities exhibits distinct patterns in volumes, shapes, and durations.

The observation of distinct IoT traffic patterns for different user activities inspires us to explore supervised machine learning algorithms to infer user activities from network traffic of heterogeneous IoT devices in smart homes. In the next three sections, we will describe how we i) collect and analyze IoT network traffic via programmable routers in smart homes, ii) explore a suite of machine learning algorithms to infer user activities based on a broad range of IoT network traffic features, and iii) evaluate the performance of our proposed ML-based approach for inferring user activities with data-sets collected from real smart homes.

III. IoT Network Traffic in Smart Homes

In this section, we first discuss how we collect and process network traffic of IoT devices from smart homes, and then shed light on distinct IoT network traffic waves triggered by different user activities.

A. Collecting and Processing IoT Network Traffic in Smart Homes

Similar to many prior research on IoT network traffic, this study also relies on programmable routers in smart homes to monitor, capture, and collect TCP/IP data packets which are sent from or received by IoT devices in home networks. Considering the privacy concerns of users in the experimental environments, network traffic data collected by the system does not carry any personal identifiable information (PII) such as the external and public IP address of home routers, the detailed packet payload, and the IP addresses of end systems on the Internet such as cloud servers, which communicate with IoT devices in smart homes. Figure 4 illustrates a simple yet effective IoT network traffic collection system we have developed for collecting, processing, and analyzing network traffic in smart homes.

To differentiate heterogeneous IoT devices in the same smart home network, we use the invariant physical media access control (MAC) address to uniquely identify and represent each IoT device. Subsequently, we use the Dynamic Host Configuration Protocol (DHCP) logs to identify the private IP addresses of different IoT devices based on their unique physical addresses. As a result, our data collection system has the ability to recognize and separate the network traffic of each IoT device for fine-grained and in-depth analysis. For example, Figure 5 shows the observations of separate network traffic for 11 individual IoT devices in smart home A during an 8-hour time span.

Similar to previous studies on IoT background network traffic and IoT device event traffic [18], our study also discovers the persistent background traffic of IoT devices, e.g., periodic heartbeat signals between IoT devices and remote servers in the cloud, as well as IoT device event traffic which is generated by the control commands on IoT devices, e.g., turning on or off smart lights via the smartphone companion app by a home user.

We adopt the same strategy used in [18] to identify the background traffic of each IoT device in smart homes via creating an "idle" time period in which there are no interactions between users and IoT devices. As illustrated in Figure 5, IoT devices AL, AQ, RC, RD, RS, TB, and TP exhibit consistent frequent background traffic, illustrated by the dotted lines, during the 8-hour "busy" time window when both background and non-background network traffic are observed, while IoT devices DW, KM, SC, and TC send or receive little background traffic during the same time period.

In this study, we are interested in studying non-background network traffic for IoT devices, which are often triggered by user activities in smart homes. Removing background traffic in Figure 5 leads to the observation in Figure 6, which shows the non-background network traffic for the same 11 IoT devices as in Figure 5. As shown in Figure 6, the non-background network traffic for the 5-device group DW, RC, RD, TB, and TP as well as the 2-device group AQ and SC exhibits very strong temporal correlations. Our in-depth analysis on these correlations leads to two observations: i) the IoT devices in the same groups are physically deployed very close in the smart homes, and ii) these devices are often triggered by the same user activities.

As the individual dots in Figure 5 and Figure 6 do not convey the actual traffic volumes and dynamics during user activities, we use Figure 7 with four typical IoT devices, i.e., `August Smart Lock`, `TP-Link Bulb`, `TP-Link Plug`, and `Kangaroo Motion Sensor` in smart home A, to present both background network traffic over a short 20-minute time window for these IoT devices as well as non-background traffic due to different user activities during the same 20-minute time window in smart home A. As illustrated in Figure 7, the repeated smaller spikes for `August Smart Lock`, `TP-Link Bulb`, and `TP-Link Plug` reflect background network traffic of these three devices, while the larger spikes for all four devices capture non-background traffic due to user activities in the smart home.

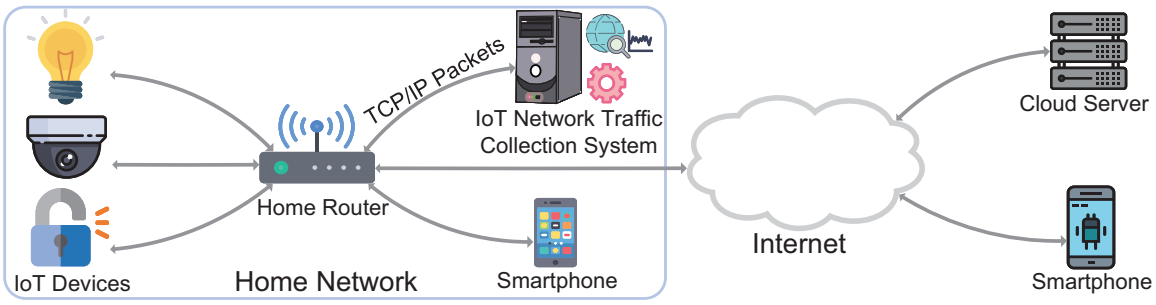

Fig. 4. An IoT network traffic collection system via programmable routers in smart homes.

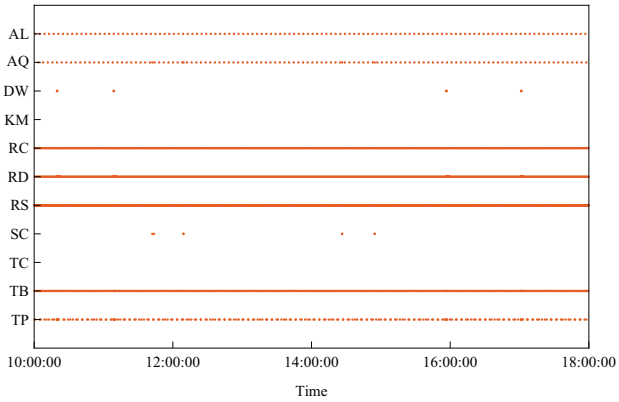

Fig. 5. The observations of *all* network traffic for 11 IoT devices in smart home A over an 8-hour time span. (cf. Table I for the mappings between the abbreviation names and the actual IoT devices)

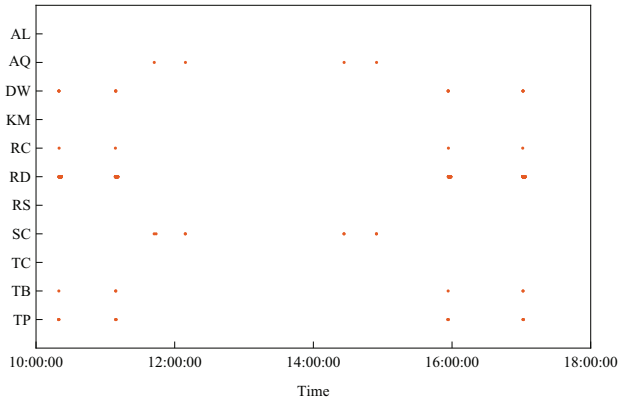

Fig. 6. The observations of *non-background* network traffic for the same 11 IoT devices Figure 5 during the same time period. (cf. Table I for the mappings between the abbreviation names and the actual IoT devices)

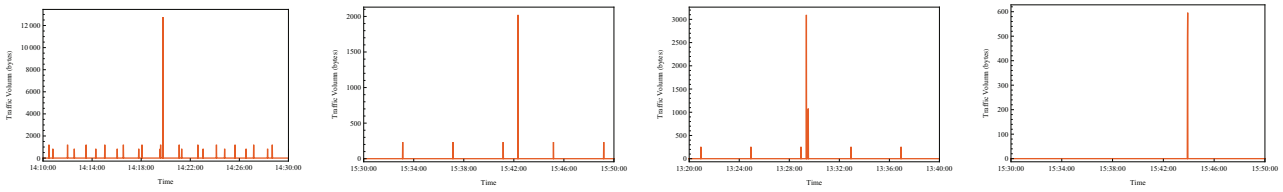

(a) August Smart Lock (b) TP-Link Bulb (c) TP-Link Plug (d) Kangaroo Motion Sensor

Fig. 7. The observations of IoT network traffic time series of four selected IoT devices, which are triggered by different user activities at different times in smart home A. Each device is triggered by exactly one user activity during the time window.

B. Extracting IoT Network Traffic Waves Trigged by User Activities

Our experiments on real smart home environments have observed that user activities often trigger multiple IoT devices in parallel due to their proximity. For example, Figure 8 shows three separate spikes in the overall IoT network traffic in smart homes, which exhibit strong temporal correlations with three separate user activities during a 10-minute time window in smart home A. As shown in Figure 8, the distinct patterns of sudden IoT network traffic jumps for three different user activities suggest the possibility of learning and recognizing different user activities in smart homes via supervised ML algorithms.

Our in-depth investigation on IoT network traffic during this 30-minute time window discovers that each of these traffic

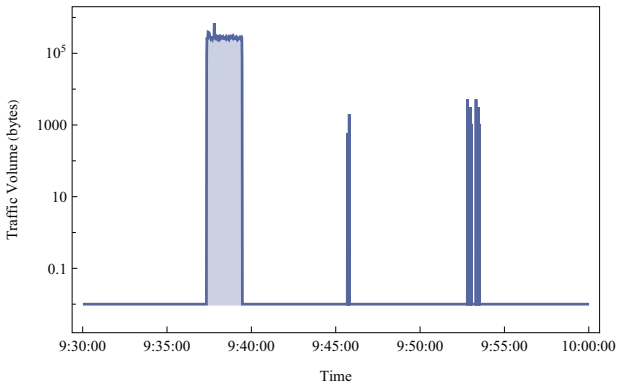

Fig. 8. The aggregated IoT network traffic waves due to three separate user activities during a 30-minute time window in smart home A.

spikes is contributed by two or more IoT devices exhibiting multiple traffic waves in parallel. More importantly, these network traffic spikes contributed by user activities in Figure 8 trigger multiple IoT devices in parallel due to their close proximity. Figure 9 breaks down network traffic in Figure 8 for each IoT device triggered by these three user activities. As shown in Figure 9, AL, AQ, RD, RS, and TC are triggered by the first user activity, KM and TB are triggered by the second activity, and TP and DW are triggered by the third activity.

These observations in Figure 8 and Figure 9 confirm that each user activity in smart homes leads to district network traffic patterns and waves of IoT devices, and lead us to explore multiclass classification algorithms for inferring different user activities with each activity forming a unique class or label.

IV. Inferring User Activities from IoT Network Traffic with Supervised ML Algorithms

In this section, we first discuss how we extract and develop a set of features from IoT network traffic waves for inferring user activities. Subsequently, we describe how we explore a suite of supervised ML algorithms for inferring user activities based on our labeled data-sets from two real smart home environments.

A. Feature Extraction from IoT Network Traffic Waves

In order to extract features from IoT network traffic for inferring user activities, we need to first identify and extract the start and end timestamps of network traffic spikes [19] for IoT devices. Specifically, we adopt a change point detection algorithm based on wavelet analysis [13] to locate the boundaries of individual traffic waves for each IoT device in smart homes.

Unlike smartphones and end hosts such as desktops and laptops in home networks, IoT devices typically have much less noisy network traffic [5], [20] due to their simple and specific functionalities. Prior studies on IoT devices network traffic [20], [6], [5], [7] have revealed that smart home IoT devices have periodical low-volume background traffic with cloud servers, and occasionally these devices exhibit with high-volumes of network traffic spikes due to user activities. For example, our traffic analysis on `August Smart Lock` device shows that the smart lock synchronizes with two of the manufacturer's cloud servers every 90 seconds and 105 seconds, and the lock periodically checks for the latest firmware updates every 6 hours. However, compared with such background traffic, the network traffic of the same smart lock triggered by user activities is rarely observed due to the infrequent nature of user activities in smart homes. Our experiments with other IoT devices lead to similar observations.

The above findings of IoT network traffic with varying frequencies lead us to explore signal processing techniques [12], [11], [14], [13] on IoT network traffic time-series for decomposing different types of traffic, and most importantly for extracting the boundaries of IoT network traffic waves caused by user activities. Specifically, we adopt the wavelet analysis approach, similar to [13], for learning and extracting network traffic with low, middle, and high frequency bands due to the strong similarity between volume-based network traffic anomalies in [13] and IoT network traffic waves in this study.

Performing wavelet analysis on IoT network traffic collected from smart homes has two complementary steps: i) decomposition (also referred to as analysis) and ii) reconstruction (also referred to as synthesis) [13]. The decomposition step extracts a hierarchy of derived *strata* with varying frequencies from the original signal of IoT network traffic s over a time period t of length n. The multi-level wavelet decomposition splits up the original signal with one low-pass sub band $L(s)$ (also referred to as the approximation level), and one or more high-pass sub bands $H_1(s), \ldots, H_m(s)$, $(m \geq 1)$ (also referred to as the fined-grained detail levels). In our experiments, we choose m as 1 thanks to simple traffic patterns of IoT devices in smart homes, which are mostly contributed by management and maintenance traffic and device events. In addition, our empirical experiments with m set in the range of [1, 5] obtain similar results. The low-pass sub band can be further split into $L^2(s)$ and $H_1L(s)$ with the same decomposition process. For simplicity, we use $HL(s)$ to denote $H_1L(s)$ throughout the rest of this paper as we choose m as 1 in our experiments. As such iteration proceeds, we obtain the derived signals, i.e., wavelet coefficients, $L^i(s)$ and $HL^{i-1}(s)(i = 1, \ldots, j)$, where j denotes the number of iterations. The iteration process stops at j when $L^j(s)$ has few signals to be further split.

The reconstruction step runs the inverse of the multi-level wavelet decomposition process with the same number of iterations and assembles new and aggregated signals with various wavelet coefficients generated from the decomposition. Similar to the discrete wavelet transform process in [13], our study also obtains the high-band, mid-band, and low-band aggregated signals via synthesizing the wavelet coefficients from the frequency level groups of [1], [2, 3], and $[4, \ldots, j]$, respectively.

The top plot of Figure 10 illustrates the original signal, i.e., network traffic volumes in bytes measured every one second, of `August Smart Lock` in smart home A during a 25-minute time window, while the bottom three plots of Figure 10 show the decomposed high-band, mid-band, and low-band signals during the same time period, respectively. Our in-depth examination confirms that the synthesized high-frequency signal captures very frequent IoT network traffic, e.g., background device state synchronization. The synthesized middle frequency signal represents network traffic which periodically occurs but with larger time intervals, e.g., new firmware checks and DHCP renewals, while the low frequency part corresponds to IoT network traffic that has a low recurrence rate, e.g., unpredictable and infrequent IoT device events triggered by user activities. More importantly, IoT network traffic waves triggered by user activities are always captured by the synthesized low frequency signal. In other words, identifying the *start* and *end* timestamps of network traffic spikes for IoT devices has naturally become a simpler problem of extracting the timestamps of the synthesized low-frequency signals with significant traffic volume *increases* and *decreases*,

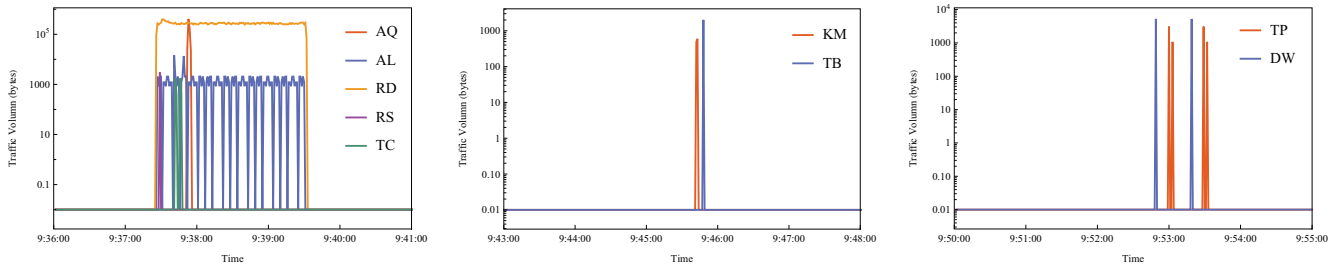

(a) The first user activity (b) The second user activity (c) The third user activity

Fig. 9. The overlapping traffic waves from multiple IoT devices for the same user activities.

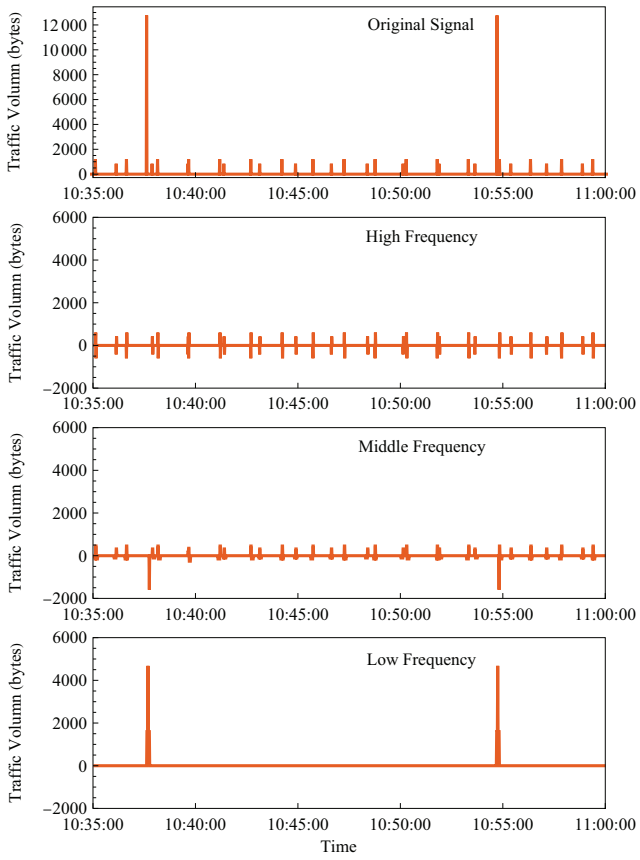

Fig. 10. The original IoT network traffic signal of August smart lock over a 25-minute time window as well as the decomposed high-band, mid-band, and low-band signals via wavelet analysis.

as evidenced in Figure 10.

After identifying the overlapping network traffic waves from two or more IoT devices, we extract and derive 19 features from the aggregated network traffic from these IoT devices for each manually-labeled user activity. These 19 IoT network traffic features include 1) number of IoT devices with the observed traffic waves, 2) number of total packets in the waves, 3) number of total bytes in the waves, 4) average bytes per packet in the waves, 5) overall time duration of the user activity, 6) number of packets per second, 7) number of bytes per second, 8) number of unique application protocols, 9) number of unique application protocols over TCP, 10) number of unique application protocols over UDP, 11) number of network flow sessions, 12) number of network flow sessions over TCP, 13) number of network flow sessions over UDP, 14) maximum duration of all network flow sessions, 15) minimum duration of all network flow sessions, 16) average duration of all network flow sessions, 17) maximum interval of inter-packet arrival times among all sessions, 18) minimum interval of inter-packet arrival times among all sessions, and 19) average interval of inter-packet arrival times among all sessions.

These features are similar to prior research [21], [22], [23] on IoT network traffic analysis. The availability of rich IoT network traffic features and labeled user activities allows us to explore supervised ML algorithms for inferring a wide range of user activities in smart homes.

B. Supervised ML Algorithms for Inferring User Activities

Given the diverse set of user activities, we use multiclass classification algorithms, rather than binary classification algorithms, for inferring user activities. In this study, we focus on several widely-used multiclass classification algorithms including Naive Bayes, random forest, k-NN, gradient boosting, and SVM, for classifying 23 different user activities in our experiments. The first three algorithms inherently support multi-class classification, while the last two algorithms are fundamentally binary classifiers. Thus, we transform the problem of multiclass classifications into multiple binary classification problems for gradient boosting and SVM via developing 23 binary classifiers with the "one-vs-the-rest" strategy for performing multiclass classification.

To realize the best performance of each supervised ML algorithm for inferring user activities, we exploit the parameter tuning process to identify the optimal hyperparameters, e.g., k in the k-NN algorithm, which are not found from the training data but must be specified before running the algorithm. We repeatedly run the ML algorithm with varying values for each hyperparameter, e.g., trying the values from 1 to 10 for the hyperparameter k in the k-NN algorithm.

To prevent the potential overfitting during the process of tuning hyperparameters for these supervised ML-models, we adopt the standard k-fold cross-validation precedure with

$k = 10$, which partitions the labelled data-sets into 10 folds. For each of these 10 folds (referred to as the *holdout fold*), we iteratively train the ML models with the remaining 9 (i.e., $k - 1$) folds as the training data set, while using the selected *holdout fold* as the test data set. After the 10 iterations, we calculate the average ML model performance and select the optimal hyperparameters with the best performance for the ML models. For ML-models with two or more hyperparameters, we use the common grid search strategy for selecting the optimal hyperparameter set.

V. Experimental Evaluations

In this section, we first describe the experimental setup of two real smart home environments for evaluating the performance of inferring user activities from IoT network traffic. Subsequently, we present the performance results of running the algorithms on data-sets collected from these two homes over a 2-month time period.

A. Experimental Setup

To evaluate our proposed approach of inferring user activities in smart homes via supervised machine learning algorithms, we setup two real smart home environments, as illustrated in Figure 2, with heterogeneous IoT devices deployed across both homes. We identify 23 common daily activities of home users in these two homes. Table II lists these 23 user activities and the associated IoT devices, which are deployed near the locations where user activities happen.

For each user activity, we repeat 100 user experiments in the smart home over multiple weeks, and manually record the approximate timestamps of each user activity, which serve as the *ground truth* of user activities for model training, testing, and evaluations. The IoT network traffic collection system, deployed at programmable routers in two smart homes, are configured to continuously monitor and collect network traffic of all smart home IoT devices.

B. Performance Evaluations

Based on the data-sets collected from these two smart home environments, we evaluate the performance of our proposed ML-based algorithms for inferring user activities with the widely-used confusion matrix, i.e, true positive, true negative, false positive, false negative, as well as the accuracy, precision, recall, F_1 score, and AUC metrics.

In our study, a true positive (TP) suggests that our ML-based algorithm correctly infers one user activity that has happened, i.e., the inferred activity matches the ground truth, while a true negative (TN) indicates the algorithm does not infer a user activity that has not happened. A false positive (FP) suggests that our ML-based algorithm incorrectly infers one user activity that has not happened, while a false negative (FN) indicates the algorithm fails to infer one user activity that has happened. The accuracy metric, $\mathcal{A}$, is calculated as $\mathcal{A} = \frac{TP+TN}{TP+TN+FP+FN}$, while the precision ($\mathcal{P}$) and recall metrics ($\mathcal{R}$) are calculated as $\mathcal{P} = \frac{TP}{TP+FP}$ and $\mathcal{R} = \frac{TP}{TP+FN}$, respectively. The F-1 score (F_1), the harmonic mean of precision and recall, is derived as: $F_1 = 2 \times \frac{\mathcal{P} \times \mathcal{R}}{\mathcal{P}+\mathcal{R}}$. Lastly, the AUC metric quantifies the area under the receiver characteristic operator curve (ROC), which plots true positive rate vs. false positive rate of classification algorithms along all classification thresholds. In our study, the AUC measure essentially captures the performance of supervised learning algorithms in distinguishing different user activities.

Table III summarizes the average accuracy, precision, recall, F_1, and AUC metrics for running supervised machine learning algorithms for inferring all 23 user activities from two smart home environments. As shown in Table III, our proposed ML-based algorithms are very effective in inferring a broad range of user activities in smart homes based on IoT network traffic. For example, the best overall performing algorithm, Naive Bayes, achieves an average accuracy of 0.8783, an average precision of 0.8801, an average recall of 0.8783, an average F_1 score of 0.8779, and an average AUC of 0.9498.

VI. Related Work

As the usage of IoT devices continues to grow in smart homes, studying their network traffic has become increasingly important in understanding the trends and dynamics of the Internet ecosystem [24], [25], [26], [27], [28], [29]. A number of research studies have been devoted to collecting, analyzing, and characterizing network traffic patterns of IoT devices [20], [18], [30], [31]. For example, the studies in [20], [18] introduce an IoT traffic measurement framework to capture and analyze incoming, outgoing, and internal smart home network traffic to study traffic behavioral profiles of IoT devices, while HoMonit [30] monitors smart home apps via collecting and analyzing encrypted wireless network traffic of IoT devices, and adopts a deterministic finite automaton (DFA) matching algorithm to represent the interactions between the smart home apps and IoT devices.

Recognizing IoT device types and models, e.g., smart plugs and smart locks, has recently attracted significant interests [32], [33], [34] from the research community for anomaly detection and security enforcement [35]. For example, the study [32] uses the features from broadcast messages to identify the types and models of IoT devices over public WiFi networks, while IoT Sentinel [33] uses various traffic features from TCP/IP data packets as IoT device fingerprints for identifying the makes, models, and software versions of IoT devices. Similarly, [34] collects network traffic from different IoT devices over a 6-month time span, and uses a multi-stage ML-based classifier with network traffic features to identify these IoT devices.

A natural extension of understanding the IoT device types is to detect the events of IoT devices, e.g., smart plugs `ON`, via characterizing and modeling network traffic [7], [6], [36]. For example, IoTAthena [6] generates the signatures of IoT device events with the ordered sequences of TCP/IP data packets, and designs polynomial-time algorithms to unveil the sequence of IoT device events from IoT network traffic in smart homes.

Several recent studies have explored IoT network traffic to extract the knowledge of the deployed environments, in

TABLE II
USER ACTIVITIES AND THE ASSOCIATIONS WITH IOT DEVICES IN SMART HOMES.

No.	User Activity	IoT Devices Triggered
1	A person without key entering the home from the front door (day)	AL, AQ, RD, RS, TC
2	A person without key entering the home from the front door (night)	AL, AQ, RD, RS, TC
3	A person with app access entering the home from the front door (day)	AL, AQ, RD, RS, TC
4	A person with app access entering the home from the front door (night)	AL, AQ, RD, RS, TC
5	A person with key entering the home from thefront door (day)	AL, AQ, RD, RS, TC
6	A person with key entering the home from thefront door (night)	AL, AQ, RD, RS, TC
7	A person ring the doorbell and leave (day)	RD, RS
8	A person ring the doorbell and leave (night)	RD, RS
9	A person checking the front door of the home (day)	RD, RS
10	A person checking the front door of the home (night)	RD, RS
11	A person with key leaving the home from the front door (lock the door) (day)	AL, AQ, RD, RS, TC
12	A person with key leaving the home from the front door (lock the door) (night)	AL, AQ, RD, RS, TC
13	A person with key leaving the home from the front door (do not lock the door) (day)	AL, AQ, RD, RS, TC
14	A person with key leaving the home from the front door (do not lock the door) (night)	AL, AQ, RD, RS, TC
15	A person appearing in the hallway of the home	KM, TB
16	A person leaving the hallway of the home	KM, TB
17	A person checking the living room's camera streaming	RC
18	A person entering the balcony from living room or entering the living from balcony (contact sensor alarm off)	RC
19	An person entering the home from the balcony (contact sensor alarm on)	RC, SC
20	An person leaving the home to enter the balcony ((contact sensor alarm on)	RC, SC
21	A person checking water leakage	DW, TP
22	A person getting out of garage to the outside.	GC, KM
23	A person coming into the garage from the outside.	GC, KM

TABLE III
PERFORMANCE EVALUATIONS OF INFERRING USER ACTIVITIES WITH SUPERVISED LEARNING ALGORITHMS BASED ON DATA-SETS COLLECTED TWO REAL SMART HOME ENVIRONMENTS.

Algorithm	Accuracy	Precision	Recall	F_1 Score	AUC
Naive Bayes	0.8783	0.8801	0.8783	0.8779	0.9498
Random Forest	0.8630	0.8639	0.8630	0.8619	0.9394
k-NN	0.7304	0.7343	0.7304	0.7295	0.9033
Gradient Boosting	0.8522	0.8545	0.8522	0.8519	0.9319
SVM	0.8261	0.8204	0.8261	0.8184	0.9198

particular, inferring user activities in smart homes for safety monitoring, anomaly detection, and connected health [4], [3], [17]. For example, Peek-a-Boo [4] uses the passive traffic sniffing strategy to capture encrypted IoT network traffic for detecting IoT device types and device events, and ultimately infers user activities with traffic and device features extracted from network traffic as well as device locations and timestamps. IoTMosaic [3] generates the signatures of diverse user activities with sequences of IoT device events, and subsequently develops an approximate matching algorithm to infer user activities from IoT device events extracted from network traffic. The recent study in [17] formulates the problems of Events to Activities (E2A) and Events to Activity Patterns (E2AP) for discovering a sequence of user activities from an ordered sequence of IoT device events in smart homes. For solving the E2AP and E2A problems, [17] designs a two-phase solution that combines a deterministic algorithm in Phase 1 for computing a small number of representative matches of activity patterns and an unsupervised ML algorithm in Phase 2 for matching a compatible set of user activity patterns with an optimal set of weights in a finite number of iterations.

Different from these prior studies, this paper explores IoT network traffic directly to infer user activities without the intermediate step of extracting IoT device events which could be missing or out-of-order due to device malfunctions [35] and varying network latency [8], [9], [3]. In addition, certain IoT device events could be difficult to extract if their traffic signatures are unknown or frequently updated.

VII. CONCLUSIONS AND FUTURE WORK

The last decade has witnessed the rapid and ubiquitous deployment of heterogeneous IoT devices in smart homes, which creates new opportunities and challenges in improving environmental awareness and intelligence for a wide range of services such as home safety, remote health monitoring, and energy saving. In this paper, we explore supervised ML algorithms to infer user activities in smart homes based on IoT network traffic collected from programmable home routers.

Inspired by our experimental observations on distinct IoT traffic waves due to different user activities, we first adopt a wavelet analysis approach to extract the beginning and end timestamps of IoT network traffic waves from the frequent background and noisy traffic. Based on the labeled datasets of user activities in smart homes and the corresponding IoT network traffic, we extract and design a wide range of features from IoT network traffic waves for each activity, and explore five widely-used supervised ML algorithms to develop multiclass classifiers to infer various user activities. Our experimental results have demonstrated our proposed ML-based algorithms are able to accurately infer a wide range of user activities in smart homes.

REFERENCES

[1] S. Hui, H. Wang, D. Xu, J. Wu, Y. Li, and D. Jin, "Distinguishing Between Smartphones and IoT Devices via Network Traffic," *IEEE Internet of Things Journal*, vol. 9, no. 2, 2022.

[2] A. Mudgerikar and E. Bertino, "Jarvis: Moving Towards a Smarter Internet of Things," in *Proceedings IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2020.

[3] Y. Wan, K. Xu, F. Wang, and G. Xue, "IoTMosaic: Inferring User Activities from IoT Network Traffic in Smart Homes," in *Proceedings of IEEE Conference on Computer Communications (INFOCOM)*, 2022.

[4] A. Acar, H. Fereidooni, T. Abera, A. K. Sikder, M. Miettinen, H. Aksu, M. Conti, A.-R. Sadeghi, and S. Uluagac, "Peek-a-Boo: I See Your Smart Home Activities, Even Encrypted!" in *Proceedings of ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec)*, 2020.

[5] J. Ren, D. J. Dubois, D. Choffnes, A. M. Mandalari, R. Kolcun, and H. Haddadi, "Information Exposure From Consumer IoT Devices: A Multidimensional, Network-Informed Measurement Approach," in *Proceedings of ACM Internet Measurement Conference (IMC)*, 2019.

[6] Y. Wan, K. Xu, F. Wang, and G. Xue, "IoTAthena: Unveiling IoT Device Activities from Network Traffic," *IEEE Transactions on Wireless Communications*, vol. 21, no. 1, January 2022.

[7] R. Trimananda, J. Varmarken, A. Markopoulou, and B. Demsky, "Ping-Pong: Packet-Level Signatures for Smart Home Device Events," in *Proceedings of Network and Distributed System Security Symposium (NDSS)*, 2019.

[8] T. Høiland-Jørgensen, B. Ahlgren, P. Hurtig, and A. Brunstrom, "Measuring Latency Variation in the Internet," in *Proceedings of ACM International Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, 2016.

[9] A. Dhamdhere, D. D. Clark, A. Gamero-Garrido, M. Luckie, R. K. P. Mok, G. Akiwate, K. Gogia, V. Bajpai, A. C. Snoeren, and K. Claffy, "Inferring Persistent Interdomain Congestion," in *Proceedings of ACM SIGCOMM*, 2018.

[10] D. Campos and T. OConnor, "Towards Labeling On-Demand IoT Traffic," in *Proceedings of Cyber Security Experimentation and Test Workshop*, 2021.

[11] P. Abry, R. Baraniuk, P. Flandrin, R. Riedi, and D. Veitch, "Multiscale Nature of Network Traffic," *IEEE Signal Processing Magazine*, vol. 19, no. 3, May 2002.

[12] P. Abry and D. Veitch, "Wavelet Analysis of Long Range Dependent Traffic," *IEEE Transactions on Information Theory*, vol. 44, no. 1, January 1998.

[13] P. Barford, J. Kline, D. Plonka, and A. Ron, "A Signal Analysis of Network Traffic Anomalies," in *Proceedings of ACM SIGCOMM Workshop on Internet Measurement (IMW)*, 2002.

[14] P. Barford and D. Plonka, "Characteristics of Network Traffic Flow Anomalies," in *Proceedings of ACM SIGCOMM Workshop on Internet Measurement (IMW)*, November 2001.

[15] L. Chen and C. D. Nugent, *Human Activity Recognition and Behaviour Analysis*. Springer, 2019.

[16] S. Birnbach, S. Eberz, and I. Martinovic, "Peeves: Physical Event Verification in Smart Homes," in *Proceedings of ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2019.

[17] G. Xue, Y. Wan, X. Lin, K. Xu, and F. Wang, "An Effective Machine Learning based Algorithm for Inferring User Activities from IoT Device Events," *IEEE Journal on Selected Areas in Communications (JSAC) Series on Machine Learning in Communications and Networks*, 2022.

[18] Y. Wan, K. Xu, F. Wang, and G. Xue, "Characterizing and Mining Traffic Patterns of IoT Devices in Edge Networks," *IEEE Transactions on Network Science and Engineering (TNSE)*, vol. 8, no. 1, pp. 89–101, 2021.

[19] H. Wang, D. Zhang, and K. Shin, "Change-point Monitoring for the Detection of DoS Attacks," *IEEE Transactions on Dependable and Secure Computing*, vol. 1, no. 4, 2004.

[20] K. Xu, Y. Wan, G. Xue, and F. Wang, "Multidimensional Behavioral Profiling of Internet-of-Things in Edge Networks," in *Proceedings of IEEE/ACM International Symposium on Quality of Service (IWQoS)*, 2019.

[21] H. Jafari, O. Omotere, D. Adesina, H.-H. Wu, and L. Qian, "IoT Devices Fingerprinting Using Deep Learning," in *Proceedings of IEEE Military Communications Conference (MILCOM)*, 2018.

[22] Y. Meidan, M. Bohadana, A. Shabtai, J. D. Guarnizo, M. Ochoa, N. O. Tippenhauer, and Y. Elovici, "ProfilIoT: A Machine Learning Approach for IoT Device Identification Based on Network Traffic Analysis," in *Proceedings of ACM Symposium on Applied Computing*, 2017.

[23] Y. Wan, K. Xu, G. Xue, and F. Wang, "IoTArgos: A Multi-Layer Security Monitoring System for Internet-of-Things in Smart Homes," in *Proceedings of IEEE Conference on Computer Communications (INFOCOM)*, 2020.

[24] D. Kumar, K. Shen, B. Case, D. Garg, G. Alperovich, D. Kuznetsov, R. Gupta, and Z. Durumeric, "All Things Considered: An Analysis of IoT Devices on Home Networks," in *Proceedings of USENIX Security Symposium*, 2019.

[25] O. Alrawi, C. Lever, M. Antonakakis, and F. Monrose, "SoK: Security Evaluation of Home-Based IoT Deployments," in *Proceedings of IEEE Symposium on Security and Privacy (S&P)*, 2019.

[26] A. Wang, A. Mohaisen, and S. Chen, "XLF: A Cross-layer Framework to Secure the Internet of Things (IoT)," in *Proceedings IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2019.

[27] Q. Qin, K. Poularakis, and L. Tassiulas, "A Learning Approach with Programmable Data Plane towards IoT Security," in *Proceedings IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2020.

[28] X. Yang, Q. Qi, J. Wang, S. Guo, and J. Liao, "Towards Efficient Inference: Adaptively Cooperate in Heterogeneous IoT Edge Cluster," in *Proceedings IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2021.

[29] H. Chi, C. Fu, Q. Zeng, and X. Du, "Delay Wreaks Havoc on Your Smart Home: Delay-based: Automation Interference Attacks," in *Proceedings of IEEE Symposium on Security and Privacy (S&P)*, 2022.

[30] W. Zhang, Y. Meng, Y. Liu, X. Zhang, Y. Zhang, and H. Zhu, "HoMonit: Monitoring Smart Home Apps from Encrypted Traffic," in *Proceedings of ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2018.

[31] X. Ma, J. Qu, J. Li, J. C. Lui, Z. Li, and X. Guan, "Pinpointing Hidden IoT Devices via Spatial-temporal Traffic Fingerprinting," in *Proceedings of IEEE Conference on Computer Communications (INFOCOM)*, 2020.

[32] L. Yu, B. Luo, J. Ma, Z. Zhou, and Q. Liu, "You Are What You Broadcast: Identification of Mobile and IoT Devices from (Public) WiFi," in *Proceedings of USENIX Security Symposium*, 2020.

[33] M. Miettinen, S. Marchal, I. Hafeez, N. Asokan, A.-R. Sadeghi, and S. Tarkoma, "IoT Sentinel: Automated Device-Type Identification for Security Enforcement in IoT," in *Proceedings of IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2017.

[34] A. Sivanathan, H. H. Gharakheili, F. Loi, A. Radford, C. Wijenayake, A. Vishwanath, and V. Sivaraman, "Classifying IoT Devices in Smart Environments Using Network Traffic Characteristics," *IEEE Transactions on Mobile Computing*, vol. 18, no. 8, pp. 1745–1759, 2018.

[35] C. Fu, Q. Zeng, and X. Du, "HAWatcher: Semantics-Aware Anomaly Detection for Appified Smart Homes," in *Proceedings of USENIX Security Symposium*, 2021.

[36] V. Bhosale, L. De Carli, and I. Ray, "Detection of Anomalous User Activity for Home IoT Devices," in *Proceedings of International Conference on Internet of Things, Big Data and Security (IoTBDS)*, 2021.