D-Score: An Expert-Based Method for Assessing the Detectability of IoT-Related Cyber-Attacks

Yair Meidan, Daniel Benatar, Ron Bitton, Dan Avraham, and Asaf Shabtai

Abstract—**IoT devices are known to be vulnerable to various cyber-attacks, such as data exfiltration and the execution of flooding attacks as part of a DDoS attack. When it comes to detecting such attacks using network traffic analysis, it has been shown that some attack scenarios are not always equally easy to detect if they involve different IoT models. That is, when targeted at some IoT models, a given attack can be detected rather accurately, while when targeted at others the same attack may result in too many false alarms. In this research, we attempt to explain this variability of IoT attack detectability and devise a risk assessment method capable of addressing a key question:** *how easy is it for an anomaly-based network intrusion detection system to detect a given cyber-attack involving a specific IoT model?* **In the process of addressing this question we** *(a)* **investigate the predictability of IoT network traffic,** *(b)* **present a novel taxonomy for IoT attack detection which also encapsulates traffic predictability aspects,** *(c)* **propose an expert-based attack detectability estimation method which uses this taxonomy to derive a detectability score (termed** *'D-Score'***) for a given combination of IoT model and attack scenario, and** *(d)* **empirically evaluate our method while comparing it with a data-driven method.**

Index Terms—**Internet of Things (IoT) Security, Attack Detection, Network Traffic Predictability, Multi-Criteria Decision Making, Analytical Hierarchical Process (AHP).**

I. INTRODUCTION

The *Internet of Things* (IoT) is a rapidly evolving trend in wireless communication [1], where various objects are connected to the Internet and can cooperate with one another to reach common goals. In home or enterprise environments, such objects may include *smart* webcams, light bulbs, and motion detectors. Although designed to improve various aspects of modern life, IoT devices have also become known as easy targets for various cyber-attacks [2], [3] which continuously increase [4], [5] along with the worldwide proliferation of the IoT [6], [7].

A fundamental problem in IoT security is that usually, [4] attacks on IoT devices are stealthy. On top of that, anomaly-based network intrusion detection systems (AIDSs) [9] might under-perform if the normal traffic behavior is complex and thus challenging to characterize. That is, when a relatively complex IoT device is compromised by, e.g., an IoT botnet, the malicious traffic might be camouflaged by the already-irregular benign traffic, such that the malicious traffic would not be deemed anomalous by the AIDS. This could happen due

to high inherent variability in the device's *dynamic* features (e.g., diversification of source and destination IP addresses, unsteadiness of packet sizes and inter-arrival times, etc.), which are possibly influenced by the device's *static* features (e.g., quantity and diversity of sensors and actuators, memory size and CPU speed).

Overall, due to either the attack's stealthiness or the normal traffic's irregularity, users may not notice that their devices are being exploited, so these devices might remain connected to networks and pose an ongoing threat. Therefore, home users and an enterprise network security administrators are advised to be cautious when deploying IoT devices which, once exploited, the attacks they might execute are hard to detect.

In many studies (surveyed in [3], [10], [11]), machine learning methods have been proposed as means for AIDSs to monitor ongoing network traffic data and differentiate benign IoT activity and malicious events. The empirical results of such research papers typically demonstrate promising *overall* classification performance. However, as noted in [12], [13], [14], the classification performance of an AIDS may vary among different IoT models even when they are compromised in the same way. This phenomenon of variability of IoT attack detection performance leads to the following two groups of interesting questions which we attempt to answer throughout this research:

1) **Network traffic predictability:** Is there a difference between the regularity/predictability of IoT network traffic and non-IoT traffic? How variable is the traffic predictability among various IoT models? Can we correlate between IoT model complexity (i.e., static features) and traffic predictability (dynamic features)?
2) **IoT attack detectability:** For IoT devices, can we explain why a given attack is relatively easy to detect on one IoT model, yet difficult to detect on another by the same AIDS? Can the ability to detect a given attack scenario involving a given IoT model be quantified in advance, and standardized into an IoT security index or label?

Past related studies have already attempted at developing standard IoT security indices [15] and designing related labels [16]. However, none of them measures attacks detectability. In this study, inspired by the notion of *IoT traffic predictability* [12], [17], [18], [19], [20], we address the above questions and coin a new complementary term, namely, the *IoT Attack Detectability Score* (abbreviated as *"D-Score"*). The D-Score is an assessment as to the ease of detecting IoT-related cyber-attacks by an AIDS. It is calculated for combinations of an attack scenario (A) and IoT model (M), as outlined in
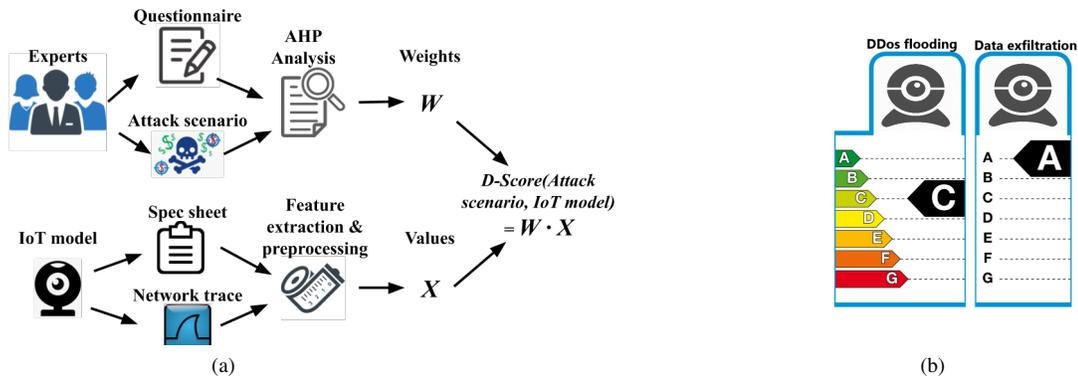
Fig. 1: Overview of *(a)* the proposed method for calculating D-Score and *(b)* the derived D-Score label, inspired by [8].

Eq. (1):

$$D - Score : A \times M \Rightarrow \mathbb{R} \tag{1}$$

To calculate the D-Score, a data-driven approach can be considered, where machine learning-based anomaly detectors are first trained and then applied to various IoT models in order to capture the relationship between the AIDS's performance and the IoT model's static and/or dynamic features. However, in practice [21], this approach requires large amounts of labeled training data that are often expensive [22] and time-consuming to acquire (see Subsection VIII-C). As an alternative, in this paper we propose an expert-based method (outlined in Fig. 1a, elaborated in Section VI). Generally, for an attack scenario A, cyber-security experts fill in a questionnaire which is constructed and analyzed in accordance with the analytical hierarchical process (AHP) [23], to obtain feature weights W. Then, for an IoT model M, static and dynamic feature values X are obtained from the spec sheet and network traces, respectively. Eventually, the D-Score is calculated as a weighted sum $(W \cdot X)$. The resultant scalar ranges from zero ("impossible to detect an attack") to one ("easy to detect"), and it can be translated into a detectability *label* (see Fig. 1b). Once D-Score labels are associated with IoT devices, domestic and enterprise customers will be more able to make informed decisions.

We summarize our contributions as follows:

1) Numerous past studies have relied on the (reasonable) hypothesis that the network traffic of IoT devices is significantly more steady and predictable than the traffic of non-IoT devices (computers, smartphones, etc.). Compared to them, we are the first to quantitatively investigate this hypothesis, using two datasets. In addition, we examine the variability of traffic predictability among disparate IoT models, correlate this variability with static features of IoT models, and leverage the variability in traffic predictability to quantify in advance IoT attack detectability.

2) We are the first (to the best of our knowledge) to explore the quantification of IoT attack detectability in an explicit and systematic manner. Note that in contrast to past studies which addressed the ease of *compromising* an IoT device [24] and *executing* a cyber-attack, no past study attempted to quantify the ease of *detecting* such attacks if and when they are executed.

3) To quantify IoT attack detectability in advance, we propose a novel expert-based method, which could spare some technically-challenging and time-consuming actions, mostly attack implementation and precise ground-truth labeling. As a basis for our method, we present a novel designated taxonomy, and for each feature in the taxonomy we also propose a means of calculating its value for a given IoT model.

4) To facilitate the collection of responses from cyber-security experts, we developed an online questionnaire [1]. This questionnaire incorporates two novel elements, namely *(a)* preliminary filtering of categories and sub-categories, and *(b)* dynamic selection of the attack scenario to be addressed by the respondent. The former minimizes the number of comparisons and thus reduces the burden on respondents; the latter facilitates the addition of attack scenarios in future research, without the need to implement and maintain multiple versions of the questionnaire.

5) The questionnaire, which addresses 4 common IoT-related attack scenarios, was completed by 40 cyber-security experts from various organizations (both from academia and industry) in the European Union, the Middle East and the Far East. We share [2] the set of (anonymous) responses with the research community so it can be utilized in future research.

II. USE CASES, SCOPE AND ASSUMPTIONS

A. Use Cases

1) Looking backward, a D-Score can provide guidance for specialists of why the same attack was detected with varying levels of accuracy for diverse IoT models. In academic settings, this can help explaining the variability in attack detection performance among experimented IoT devices [12], [13], [14]. In enterprise settings, the D-Score can assist in analysing the functioning of an AIDS in IoT-related cyber-incidents.

2) Looking forward, a D-Score can be leveraged to support procurement decisions of home users and enterprise network administrators. For instance, when contemplating

---

[1]https://dscore.limequery.com/915153
[2]http://doi.org/10.5281/zenodo.4018614

which *smart* security camera model to deploy, a model with a higher D-Score is expected to demonstrate better detection performance, thus safer and more advisable. The D-Score can also supplement organizational processes of risk assessment and risk mitigation. Further details are provided in Subsection III-C and Section VI, respectively.

B. Scope

In this research we propose a taxonomy and a method to assess in advance the attack detectability (rather than device exploitability) of IoT models. To this end, similarly to [25], we define an *IoT model* as a combination of four elements: Type/manufacturer/model number/firmware version. For instance, Webcam/Provision/PT-838/v.1.3.3 [26] is an IoT model we use in this research for evaluation. We focus on the IoT model granularity rather than, e.g., the device's manufacturer or type, because in many cases [25], [27], [28] IoT malware rely on exploiting vulnerabilities that are strongly associated with specific firmware versions of specific models made by specific manufacturers. From among the various application domains in which the IoT is proliferating [29], we concentrate on the sectors of *smart* homes and enterprises, and relate mostly to high-end commercial IoT devices such as web-enabled cameras, TVs, light bulbs, sockets, baby monitors, thermostats, door bells, etc. Those devices typically connect to the Internet via Wi-Fi and transfer data using the TCP/IP stack. Our taxonomy (see Section V) is tailored to reflect those typical characteristics.

C. Assumptions

Considering the scope defined above, we make several assumptions in this research. First, we assume that there is concern that a stealthy IoT-related attack might occur and go undetected by an AIDS. More specifically, we address attacks that can potentially be detected by an AIDS (though with various detectability levels), such as data exfiltration [30], [31], botnet gathering [18], communication with a command and control (C&C) server [2], and execution of flooding attacks as part of a distributed denial of service (DDoS) attack [2]. We also assume that the IoT device is connected to the Internet directly (i.e., not via a communal gateway) and that its network traffic data is being monitored by either enterprise-level [32] or domestic [33], [34] intrusion detection systems.

D. Threat Model

In this research we assume a system in which an IoT device is connected to a local area network (LAN) and the Internet. We also assume that the LAN is monitored by an AIDS in order to identify a variety of cyber-attacks. In this research we consider two threat models and evaluate them quantitatively (see Section VII. In the first threat model, we consider an infected IoT device that is fully controlled by the attacker. In this threat model, we assume that the attacker can receive/transmit network packets from/to the LAN and the Internet. By utilizing these capabilities, the attacker can execute the following attacks (summarized in Table V): *(a)*

C&C communication where the infected IoT device tries to communicate with a C&C server controlled by the attacker and located on the Internet. To avoid detection of C&C communication, adversaries often mimic normal traffic. *(b) DDoS flooding attacks* where the infected IoT devices are used to perform network denial-of-service attacks. This is done by exhausting the network bandwidth of the target service (such as specific websites, email services, the DNS, and web-based applications). *(c) Data exfiltration attacks* where the attacker can scan the target network and send reconnaissance data to a server controlled by the attacker. To do so, adversaries often utilize compression and encryption, and they can also set size limits on the data transmission to avoid detection. In the second threat model, we consider a legitimate IoT device that is deployed on the target network. We further assume that the IoT device is reachable from the Internet. In this threat model, we consider a *bot scanning* scenario where an adversary scans the IoT device for recruitment as part of botnet propagation. Note that the threat models we assume and evaluate can easily be extended in the future, without the need to adjust the taxonomy presented in this paper (see Section V).

III. BACKGROUND AND RELATED WORK

A. IoT vs. Non-IoT Security

Compared with the cyber-security aspects of other, non-IoT, systems, IoT security is known for its large attack surface. This attack surface exists on billions of devices, many of which are unfortunately known to have increased vulnerability [24]. Broad heterogeneity is one of the common attributes used for the characterization of IoT devices, making it difficult for device manufacturers to institute a single standard of security, unlike traditional IT systems, which are less heterogeneous. Moreover, compared to traditional IT systems, an IoT deployment may involve numerous devices and countless lines of code created by a much larger pool of developers and have many types of hardware and operating systems. Also, traditional IT devices such as local servers and personal computers are typically located in a closed environment with access control measures, whereas IoT devices are often located in open and unattended environments, leaving an opening for intruders to intentionally gain physical access to them [35]. IoT devices are often resource-constrained in terms of processing power and memory, such that they might lack robust security protocols, and sufficient computational power for encryption [36]. Due to these hardware limitations, many smart devices are not capable of protecting themselves with host-based security solutions such as endpoint protection platforms, anti-malware, and endpoint detection and response. These solutions, while well suited for traditional IT systems, are too computationally intensive for the limited power, battery, and memory capacity of IoT devices. Unlike servers and PCs that constantly receive security updates and patches, IoT devices are updated infrequently and fail to address the risk of zero-day vulnerabilities. Moreover, in many cases, end users are not capable of adding security measures to the built-in operational system. Lastly, IoT devices depend on third-party libraries and components that act as a 'black box,' making them difficult to patch, control, and scan for vulnerabilities.

B. Recent Security Research in the IoT

IoT security is a broad area of research, in which a variety of subjects are being studied. These subjects include various IoT attack scenarios, defense mechanisms, risk mitigation approaches, and risk assessment methods. In multiple recent studies on IoT security, several attacks are repeatedly mentioned as open-ended key challenges. Among these these attacks are DDoS attacks [37] (addressed in our research as well) which are considered a convenient way for an attacker to abuse a system due to the large number of IoT devices that are continuously connected to the Internet [38]. In a recent review paper on IoT security challenges [39], DDoS attacks and data sniffing are mentioned as major threats to IoT systems, while in [40], the authors proposed a deep learning-based methodology for DDoS attack detection and zero-day attacks in IoT networks. Another IoT-related attack that has been addressed in several recent studies (including our research) is data exfiltration. This type of attack is commonly considered a big security challenge [41], [42]. Another prominent security issue in IoT systems is bot detection, a challenging task that has received a lot of attention from researchers, including [43], where deep learning techniques are used for bot detection in IoT network traffic.

C. Risk Assessment in the IoT

Along with the IoT expansion, the demand for appropriate risk assessment methods has increased. However, due to the heterogeneity, scalability and connectivity of IoT devices, IoT risk assessment presents significant challenges. Over the years, various methods have been developed as a means of evaluating the security risks associated with the IoT. Among those methods, IoT attack graphs [44], [45], [46] provide a comprehensive view of attacks, and can assist in identifying multi-hop high-risk attack paths in a large-scale environment. A different objective was addressed in [47], where the authors aimed at quantifying the ease of IoT-related attack execution. In more recent research studies [48], [49], the authors attempted at developing a method for IoT device risk assessment, i.e., for quantifying in advance the influence of an IoT device on the security level of the environment within which it operates. Unlike these and other studies, our goal is not to identify IoT attack paths, determine the ease of *compromising* an IoT device or estimate an IoT device's influence on the security level of its environment. In contrast, our goal is to quantify in advance the ease of *detecting* an IoT attack, assuming that an attack of this kind is likely to occur.

D. Existing Taxonomies in IoT Security

For the quantification of IoT attack detectability, the methodology we propose relies on two elements: *(a)* a novel taxonomy, and *(b)* an expert-based mechanism that assigns a weight to each feature in the taxonomy. Regarding the first element, several taxonomies have already been proposed, none of which sufficiently supports our goal. For example, none of the existing taxonomies addresses user interaction factors, which are strongly correlated [50], [51] with traffic generation.

Most existing IoT security-related taxonomies can be assigned to one of the following two categories: *(a)* IoT architecture taxonomies [52], or *(b)* IoT attack mechanism taxonomies [53], [54].

IoT architecture taxonomies are challenging to define, mainly due to the heterogeneous nature of IoT devices and the lack of standardization in this domain [55]. In addition, IoT architectures may need to be more adaptive than traditional computing architectures, in order to cope with the rapid advancement in this field as well as the typical real-time interaction of a given IoT device within its environment. In [52], a four-layered IoT architecture was suggested, and each layer was associated with various attack types. The authors of [52] showed how previously proposed IoT architectures rely on the OSI model [56] using varying granularity levels. These architectures are relatively general (rather than attack-oriented), as they mostly discuss the structure of the layers, their functionality, and their main components.

IoT attack mechanism taxonomies are also used to characterize IoT devices, however with an emphasis on related attacks. In [57], the authors focused on applying traditional Internet standards to smart devices, in order to simplify the integration of cyber-attacks in the IoT context. In contrast, [53] and [54] studied the limitations of traditional security when applied to smart devices, and suggested IoT-specific taxonomies of security attacks. The authors of [53] presented six elements to consider when characterizing IoT attack mechanisms: device properties, adversary location, attack strategy, access level, information damage level, and host compromise. In [54] two additional elements were suggested: the attack protocol and the communication stack protocol.

Of the above two categories, our study is more closely related to the latter, as we present a security-oriented IoT taxonomy. In comparison to existing taxonomies, we add behavioral and IoT-specific features in a measurable manner, thus facilitate the quantification of IoT attack detectability. In contrast to existing taxonomies, ours also proposes a hierarchical structure that serves to characterize a given IoT model from a broad range of viewpoints, i.e., hardware, software, user interaction (not addressed before in the literature), and networking. Moreover, we propose an expert-based means of weighing the features in this novel taxonomy for various attack scenarios.

IV. PREDICTABILITY OF IOT NETWORK TRAFFIC

The notion of *IoT traffic predictability* was mentioned several times in the literature [12], [17], [18], [19], [20], however mostly as an (educated) assumption, i.e., without validating it. Throughout this section we quantitatively investigate this assumption and review related literature. For the quantitative investigation we utilize the IoT-deNAT dataset [58], which contains (in the form of NetFlows) the outbound traffic generated by commercial IoT models as well as non-IoT devices (see Table I). Note that for simplicity (and also since the firmware of the IoT devices did not change during the data collection period [25]), throughout this paper we omit the *firmware version* component from the IoT model names.
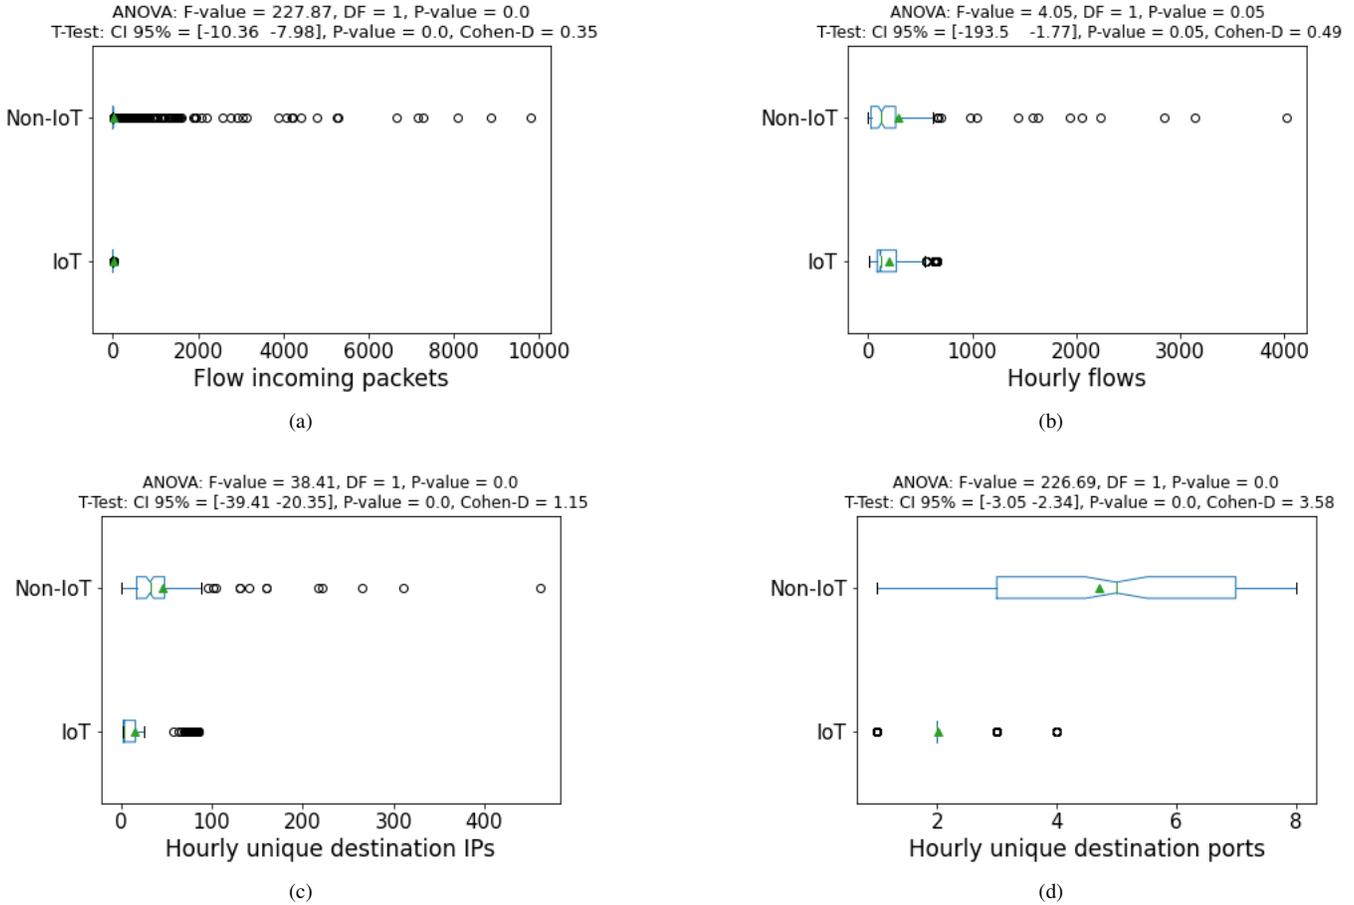
Fig. 2: Comparison of network traffic predictability between IoT and non-IoT devices.

Further details about this dataset, which was captured during a substantial period of 37 days, are provided in [25].

In our experiments we explored the IoT network traffic predictability from three different perspectives: *(a)* Comparison between the traffic predictability of IoT vs. non-IoT devices as two distinct groups, *(b)* comparison of traffic predictability among different models of IoT devices, and *(c)* association of IoT model complexity with traffic predictability, later to be translated into attack detectability. Following are the results of those experiments, in terms of four key performance indicators (KPIs) of network traffic:

- *Flow incoming packets*: Number of packets in an inbound flow
- *Hourly flows*: Number of outbound flows per hour
- *Hourly unique destination IPs*: Number of unique destination IP addresses per hour, corresponding with the number of destinations communicated
- *Hourly unique destination ports*: Number of unique destination ports per hour, corresponding with the number of protocols used

A. Traffic Predictability of IoT vs. Non-IoT Devices

The basic reasoning which guides the notion of IoT traffic predictability is that IoT devices tend to have specialized

TABLE I: IoT and non-IoT models used to quantitatively evaluate the traffic predictability, based on the IoT-deNAT dataset [58]

Category	Model	Number of outbound flows	Number of sensors	Number of actuators	CPU speed (MHz)	Supports adding apps
IoT	webcam/Amcrest/IPM-721W	479,492	5	1	1,000	True
	webcam/Samsung/SNH-1011N	217,899	5	1	216	False
	light_bulb/TP_Link/LB130	102,767	0	0	200	True
	socket/Wemo/Insight	96,426	1	1	360	True
	speaker/Sonos/One	78,480	0	1	400	True
	doorbell/Amazon/Ring	17,179	4	2	80	True
Non-IoT	laptop/Dell/Latitude_E6430	23,254				
	smartphone/Samsung/Galaxy_Note_5	15,001				
	laptop/Dell/Latitude_7400	4,864				
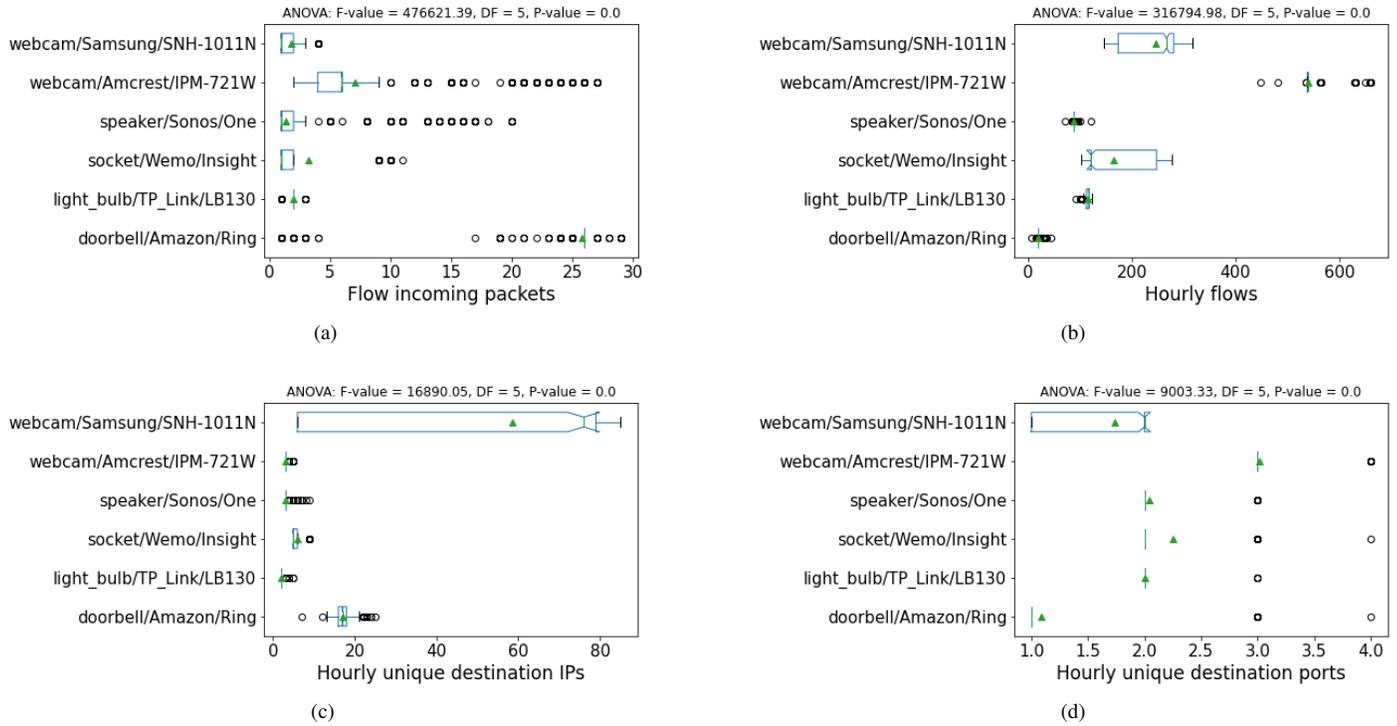
Fig. 3: Comparison of network traffic predictability among IoT models.

functions with limited input, and they also normally have only a limited variety of functional states [59], such as "On" and "Off", "Connecting to Wi-Fi", "idle", "status check", etc., where each state has distinctive network traffic patterns. Thus, as opposed to typical non-IoT devices (e.g., servers, desktop computers, laptops, tablets and mobile phones [60], [61]), IoT devices are hypothesized to behave in a relatively predictable manner in terms of traffic patterns. For instance, laptops and smartphones typically access a large number of web endpoints for browsing, while IoT devices normally send only automated pings (or messages with a predefined structure) to a finite number of endpoints.

In order to test the (reasonable) hypothesis regrading the relative predictability of IoT network traffic, compared with non-IoT devices, we *(a)* conducted one-way ANOVA tests [62] to ascertain a statistical difference of each KPI between the two groups (IoT vs. non-IoT), *(b)* conducted two-sided t-tests [63] which provide 95% confidence intervals (CIs) of the difference in the means of each KPI, and *(c)* visualized the distributions of these KPIs using boxplots. As can be seen in Fig. 2, the ANOVA and t-tests reveal significant differences ($P-value \leq 0.05$) between IoT and non-IoT devices. That is, IoT devices tend to have fewer packets in an inbound flow as well as fewer total outbound flows per hour; IoT devices also communicate with fewer destinations than non-IoT devices while using fewer ports. As visualized by the boxplots in Fig. 2, not only the means of the KPIs differ between IoT and non-IoT, but rather also the range and the variability of each KPI are smaller for the evaluated devices. For example, in most cases (see plot and CI in Fig 2c), IoT devices communicate with [20.35, 39.41] fewer destination IP addresses per hour,

and in any case this number reaches only a few dozens at most (depicted as outliers). In comparison, non-IoT devices have demonstrated several hours, during which hundreds of destination IPs were communicated per hour.

B. Variability in Traffic Predictability among IoT Models

In many cases, the relative predictability of IoT network traffic makes it feasible to establish a solid baseline profile of normal activities. As such, it has been leveraged for several security-related applications such as device type identification [27], [64], user authentication [65], [66], human activity recognition [67], [68] and anomaly detection for intrusion detection [12], [17]. Nevertheless, in the context of anomaly detection, the authors of [12], [17], [20] suspected that not all of the IoT models have equally predictable network traffic patterns. As a result, certain IoT models are more amenable than others to network anomaly detection, even if the anomaly is due to the same reason (i.e., the same cyber-attack).

To quantitatively investigate the above suspicion regarding unequality of traffic predictability among disparate IoT models, we conducted one-way ANOVA tests. As evident in Fig 3, significant differences in means ($P-value = 0.0$) were found, suggesting that dissimilar sub-populations (i.e., groups of IoT models) are present for each of the four KPIs. For instance, as illustrated in Fig. 3c, webcam/Amcrest/IPM-721W [69] communicates with fewer destination IP addresses per hour (and also with lower variability) than doorbell/Amazon/Ring [70]. In terms of our use-case, this means that eventually, if these two IoT models are recruited to the same botnet and take part in executing a DDoS attack pointed at the same victim (identified by a destination IP address), this attack would likely
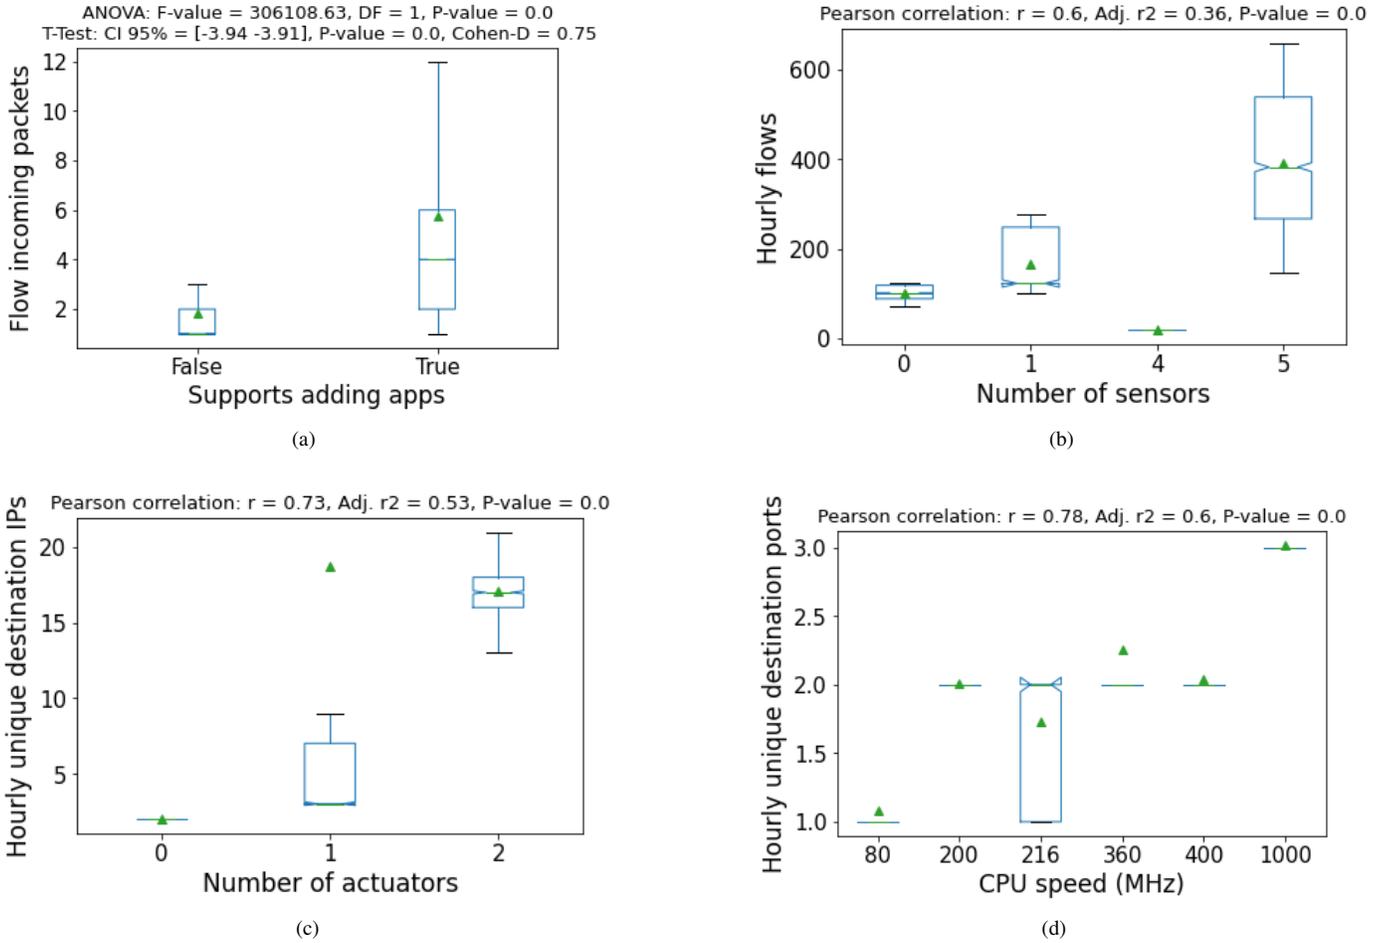
Fig. 4: Correlation between IoT model complexity and network traffic predictability.

pop up more clearly as an anomaly for webcam/Amcrest/IPM-721W (i.e., higher D-Score).

For quantitative investigation of traffic predictability, in addition to comparing KPI means using ANOVA tests (as described above), we also compare the KPIs' tendency to deviate from a random walk using the Hurst exponent [71], denoted as H. This technique has been widely used in a variety of domains, including finance [72], medicine [73] and IoT security [74]. In practice, we treat the consecutive values of each KPI as a time series and then we calculate H and use it to classify the KPI to one of three categories: *Brownian motion* or *geometric random walk* ($H = 0.5$); *anti-persistent behavior* ($H < 0.5$), where the time series reverts to the mean; or *persistent trending behavior* ($H > 0.5$). Generally, the farther H is from 0.5, the stronger non-random (i.e., the more predictable) the KPI's behavior is. To calculate H we used a designated Python package named *hurst* [75] and applied it to the IoT-deNAT dataset [58], which we previously used for analyzing traffic predictability. The H values calculated for this dataset are presented in Table II, where the rows correspond to the same devices in Table I, and the four right hand side columns correspond to the same KPIs in Figs. 2, 3 and 4. In each column, the H values that are the closest or farthest from 0.5 (the point which

indicates the least predictability) are highlighted in bold. With regard to the comparison of traffic predictability between IoT and non-IoT devices, as can be seen in Table II, in most cases, IoT devices are the ones to have H values farthest from 0.5, meaning that they are more predictable than non-IoT devices. This difference in H is most prevalent in the 'Flow incoming packets' KPI, where for IoT devices H is very close to one (indicating that for this KPI a high value is often followed by an even higher one, and vice versa), while for the non-IoT devices H is closer to 0.5 (indicating weaker connections between consecutive values). Regarding the comparison of traffic predictability among IoT models, Table II presents non-negligible ranges of H. This finding led us to further investigate this variability and examine whether an IoT device's network traffic predictability is associated with its complexity, as described next.

C. Association of IoT model complexity with network traffic predictability

In this research, we hypothesize that the variability of traffic predictability among disparate IoT models (demonstrated above) is correlated with their difference in basic complexity, and can also be translated into varying levels of attack detectability. That is, the more complicated an IoT model is (in

TABLE II: Hurst exponent (H) values calculated for devices in the IoT-deNAT dataset [58]

Category	Model	Flow incoming packets	Hourly unique destination IPs	Hourly unique destination ports	Hourly flows
IoT	webcam/Amcrest/IPM-721W	1.001	**1.256**	1.264	0.144
	webcam/Samsung/SNH-1011N	0.988	**0.399**	0.984	0.640
	light_bulb/TP_Link/LB130	**1.139**	1.219	**1.359**	0.963
	socket/Wemo/Insight	0.965	0.842	0.949	**0.277**
	speaker/Sonos/One	0.932	0.962	1.063	0.785
	doorbell/Amazon/Ring	0.623	0.975	1.056	0.724
	Average	0.941	0.942	1.113	0.589
	Range	0.516	0.857	0.410	0.819
Non-IoT	laptop/Dell/Latitude_E6430	0.665	0.980	1.064	**1.004**
	smartphone/Samsung/Galaxy_Note_5	**0.602**	0.951	**0.855**	0.944
	laptop/Dell/Latitude_7400	0.837	0.972	0.913	0.857
	Average	0.701	0.968	0.944	0.935
	Range	0.235	0.029	0.209	0.147

terms of hardware, software, user interaction, and networking), the less predictable its network traffic would be, such that generating a solid baseline profile of normal traffic behavior is less feasible. As a result, any attack would be less predictable by an AIDS, because anomalies would be harder to detect within noisy and non-deterministic traffic data. To test the hypothesis regarding the association between the complexity of an IoT model and its network traffic predictability, we characterized the evaluated IoT models based on features from our proposed taxonomy (introduced in Section V) and correlated them with the four KPIs. As illustrated in Fig. 4, although not linear or entirely monotonous, trends do seem to exist between the static features (i.e., complexity characteristics) and the dynamic features (i.e., network traffic KPIs) of the evaluated IoT devices. For example, higher values of hourly flows are associated with a larger number of sensors (Fig. 4b), logically because more data is captured from the IoT device's surroundings by multiple and various sensors, and the data is delivered to the designated servers via the Internet using larger quantities of outbound network traffic.

In this section we investigated the predictability of IoT network traffic from three complementary perspectives (IoT vs. non-IoT devices, differences among IoT models, and correlation with static IoT features). As part of this investigation, we conducted statistical hypothesis testing and visualization using data from [12]. As can be seen in Appendix A, we obtained similar findings from another freely-available dataset [3], collected by UNSW researchers [76] from various IoT and non-IoT devices during 20 days. Some of our experimental findings, obtained using the IoT-deNAT and the UNSW datasets, conform with past research, as follows: Regarding the traffic predictability of IoT vs. non-IoT devices (Subsection IV-A), similar to us the authors of [60] noted that overall, IoT devices generate much less traffic compared to traditional non-IoT devices. Also, according to their findings, an IoT device typically transfers only less than 1 KB per session. However, since their goal was to facilitate IoT model fingerprinting, their quantitative evaluation focused on finding distinct levels of traffic attributes which are characteristic of certain IoT models, rather than on statistical validation of the *differences* between IoT and non-IoT (as in our case). In a later paper [76] which extended the above-mentioned fingerprinting goal and method-

ology, the results also highlighted the most prominent values of traffic features (e.g., distinctive destination ports) instead of comparing the number of unique values per hour between IoT and non-IoT (like we do). Regarding the variability in traffic predictability among IoT models (Subsection IV-B), the authors of [77] found statistically significant differences in the distribution of packet inter-arrival time among disparate IoT models, and the authors of [59] visualized the distribution of packet size among several IoT devices. However, each of these studies focused on one feature only, whereas we statistically compared among IoT models using four other (dynamic) features, while also correlating them with (static) device features. The authors of [20] classified the traffic predictability of smart home IoT devices into four levels, using the coefficient of variation of the ratio of the received and sent volume of traffic. However, as opposed to our study, they did not compare the IoT predictability with non-IoT, nor did they associate traffic predictability with device complexity or attack detectability. Regarding the association of IoT model complexity with network traffic predictability (Subsection IV-C), we were unable to find any existing study which performed a similar analysis.
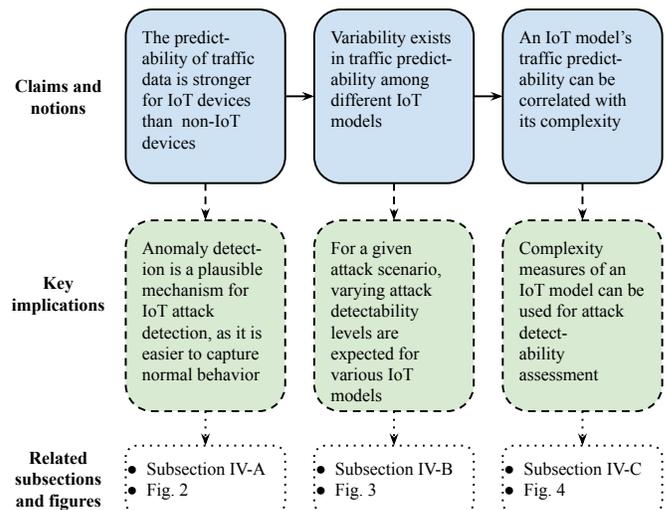
Fig. 5: The logic flow behind the statistical analysis covered in Section IV, which motivated the development of our IoT attack detectability assessment method (D-Score).

---

[3]https://iotanalytics.unsw.edu.au/iottraces
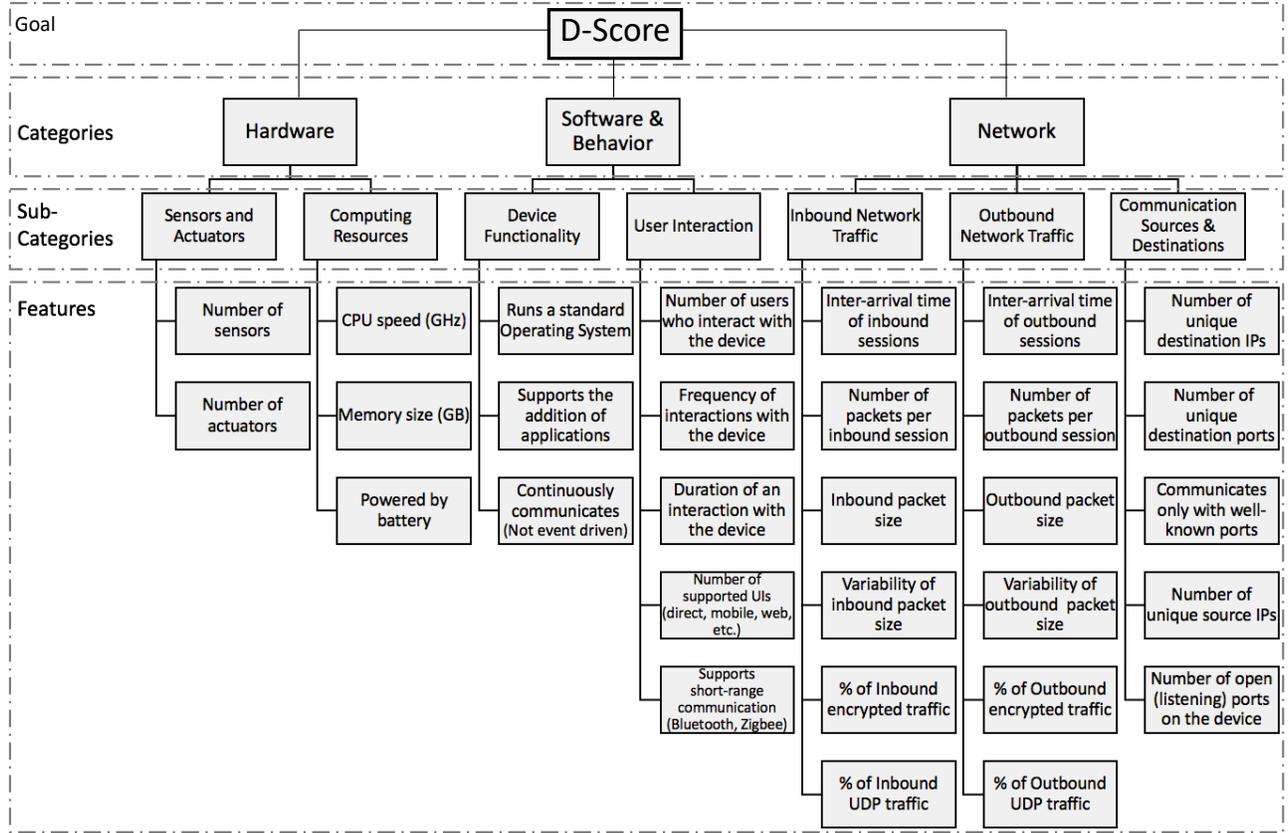
Fig. 6: Proposed taxonomy for IoT attack detectability.

The logic flow (illustrated in Fig. 5) behind the statistical analysis covered in this section motivated the development of our IoT attack detectability assessment method (D-Score). That is, if IoT network traffic is more predictable than non-IoT traffic, variability exists in traffic predictability among IoT models, and this variability can be correlated with the IoT model's complexity, then anomaly-based attack detection is a plausible mechanism for attack detection in the IoT, where features that capture the complexity of an IoT model can be utilized to assess its attack detectability, i.e., its D-Score. Given this motivation, in the sections that follow we *(a)* present a novel taxonomy which is designed to characterize the basic complexity of a given IoT model (namely, "static features") as well as the predictability of its network traffic (namely, "dynamic features"), and *(b)* propose and empirically evaluate an expert based method which assigns weights to those features in the context of a given attack scenario, in order to ultimately produce a risk assessment in the form of D-Score.

V. TAXONOMY FOR IOT ATTACK DETECTABILITY

The authors of [78] propose a concise list they refer to as *detectable attack features*. It consists of twelve network traffic features to be analyzed by AIDSs, since they are the ones most likely to change upon the occurrence of a cyber-attack. In contrast, we propose a considerably richer taxonomy (illustrated in Fig. 6), which consists of thirty *features*, hierarchically distributed among three *categories* and seven *sub-categories*. Of the three categories in our taxonomy, the *Network* is largely congruent with the above list of detectable attack features. We refer to them as *direct features*, since they are extracted from network traffic and directly analyzed by an AIDS. The direct features in our taxonomy have been carefully designed by domain experts using a systematic process, following a thorough review of past related research, e.g., [12], [17], [78], [79], [80], [81].

In addition to the direct features, for the quantification of IoT attack detectability, we also propose the novel use of *indirect features*, which are those features that are not directly analyzed by an AIDS, but are likely to affect the direct features. The indirect features mainly fall in the *Hardware* and *Software and Behavior* categories, (see Table III). The indirect features are based mostly on domain knowledge as well as reasoning and experimentation concerning the correlation between device complexity and traffic predictability (see Subsection IV-C). For instance, an indirect feature that we find to be potentially indicative of IoT cyber-attacks is the *Number of sensors*. This feature is not analyzed directly by an AIDS, however, it is likely to affect direct features such as the *Variability of outbound packet size* or the number of *Hourly flows* (see Fig. 4b). For example, an IoT device with just one sensor (like a smart motion detector) is likely to have much more stable and predictable patterns of outbound traffic than another device which also includes a webcam,

TABLE III: Categories, sub-categories and features in our proposed taxonomy, as well as the means to evaluate them.

Category	Sub-category	Feature	Data Source / Calculation / Description	Unit
HW: Hardware	SNA: Sensors & actuators	NSNS: Number of sensors	Spec sheet	Count
		NACT: Number of actuators	Spec sheet	Count
	RSR: Computing resources	CPUS: CPU speed	Spec sheet	MHz
		MEMS: Memory size	Spec sheet	MB
		BATT: Powered by battery	Spec sheet	1/0
SB: Software & behavior	FNC: Device functionality	STOS: Runs a standard operating system	nmap -O <IP address>	1/0
		ADAP: Supports the addition of applications	Spec sheet	1/0
		CCOM: Continuously communicates (not event driven)	(Mean hourly packets at night) / (max hourly packets during the day)	%
	INT: User interaction	NUSR: Number of users who interact with the device	Expected number of unique users	Count
		FINT: Frequency of interactions with the device	Expected mean number of hourly user interactions	Count/hour
		DINT: Duration of an interactions with the device	Expected mean duration	Seconds
		NUIS: Number of supported UIs (direct, mobile, web, etc.)	Number of unique user interfaces	Count
		SRNG: Supports short-range communication (Bluetooth, Zigbee)	Spec sheet	1/0
NT: Network	INB: Inbound network traffic	IATI: Inter-arrival time of inbound sessions	Median	Seconds
		PCKI: Number of packets per inbound session	Median	Count
		PCSI: Inbound packet size	Sum(inbound packet sizes) / count(inbound sessions)	Bytes
		PCVI: Variability of inbound packet size	St.Dev(inbound packet size)	Bytes
		ENCI: % of inbound encrypted traffic	Count(inbound packets to ports 443 or 8443) / count(inbound packets)	%
		UDPI: % of inbound UDP traffic	Count(inbound UDP packets) / count(inbound packets)	%
	OUT: Outbound network traffic	IATO: Inter-arrival time of outbound sessions	Median	Seconds
		PCKO: Number of packets per outbound session	Median	Count
		PCSO: Outbound packet size	Sum(outbound packet sizes) / count(outbound sessions)	Bytes
		PCVO: Variability of outbound packet size	St.Dev(outbound packet size)	Bytes
		ENCO: % of outbound encrypted traffic	Count(outbound packets to ports 443 or 8443) / count(outbound packets)	%
		UDPO: % of outbound UDP traffic	Count(outbound UDP packets) / count(outbound packets)	%
	SRD: Communication sources & destinations	DSIP: Number of unique destination IP addresses	Hourly median	Count
		DSPR: Number of unique destination port numbers	Hourly median	Count
		WLPR: Communicates only with well-known ports	Count(outbound packets to ports<=1024) / count(outbound packets)	1/0
		SRIP: Number of unique source IP addresses	Hourly median	Count
		OPPR: Number of open (listening) ports on the device	nmap <IP address>	Count

microphone, temperature sensor, humidity sensor, and ambient light sensor (like a smart baby monitor). This is based on the logic that more sensors mean more communication protocols and therefore, greater diversity of packet sizes. The number of sensors can also affect the number of destination ports, destination IP addresses, variability of session lengths and inter-arrival times. Altogether, when the complexity of an IoT device is relatively low (e.g., it uses only a few sensors), the resultant traffic patterns are expected to be rather predictable and facilitate the creation of a normal behavior profile. In turn, traffic anomalies due to cyber-attacks are more likely to be prominent, and the IoT attack detectability (i.e., D-Score) is expected to be higher.

When designing the taxonomy we balanced a number of considerations, such that the taxonomy would be:

- **IoT orientated**: As opposed to the general purpose OSI models discussed in Subsection III-D, our taxonomy is more specific to the IoT. For example, the features under the *Sensors and actuators* sub-category might not be relevant for PCs. Additionally, features such as the *Number of unique destination IPs* are more informative for IoT devices than for PCs, because unlike PCs, IoT devices typically communicate with only a small set of predefined servers.
- **Extensive**: Capable of (hierarchically) describing any IoT device from a broad range of perspectives, without leaving out aspects which are potentially informative for attack detection. These include hardware, software, user interaction, and networking viewpoints for any IoT model.
- **Informative for attack detection**: This includes direct features which are likely to behave abnormally in case of an attack as well as indirect features which affect the values and stability of the direct features.
- **Flexible**: Able to facilitate D-Score computation for *(a)* a large variety of known and unknown attack scenarios, and *(b)* unlimited IoT models (within the scope defined in

Subsection II-B), since it does not depend on predefined hardware components or specific communication protocols.
- **Coherent / self-explanatory**: Easily understood by security professionals with differing expertise areas, so they could thoughtfully respond to the questionnaire that uses this taxonomy.
- **Feasible**: Capable of easily capturing the features' values in practice, such that the D-Score computation for any given IoT model is achievable within a reasonable investment of time and effort.

Following is a description of our taxonomy's key components:

A. HW: Hardware

1) SNA: Sensors and Actuators: Sensors and actuators are key characteristics of an IoT device, which typically interacts with the physical world by collecting information from the device's surrounding environment and performing required actions [54], as opposed to typical servers, PCs, and mobile devices. Some common sensors in IoT devices are motion detectors and temperature sensors, while lights and speakers are examples of IoT actuators. The *total* number of sensors and actuators mainly affects the frequency of communication, and thus affects the inter-arrival time (IAT) of packets as well. Alternatively, the number of *unique* sensors or actuators mainly affects the variability of packet sizes and the diversity of communication protocols, sources, and destinations (IP addresses and ports). However, for consumer IoT devices (unlike smart cars, for example), it is anticipated that the number of unique sensors and actuators should be almost identical to the total number of sensors and actuators, so for simplicity, we did not allocate distinct (and probably redundant) features for both.

2) RSR: Computing Resources: The CPU, memory and power source are fundamental resources that can have an im-

pact on packet traffic throughput [82]. For instance, an insufficient CPU would not allow a highly intensive flooding attack, and might lead to crashes and significant traffic abnormalities under such circumstances. Also, additional memory may better support the computation needed to encrypt network traffic. Battery-powered IoT devices (e.g., smart smoke detectors), for the most part, are associated with low computational requirements. This implies simple functionality with minimal network traffic volume, so that attack-related traffic abnormalities are likely to emerge.

B. SB: Software and Behavior

1) FNC: Device Functionality: Non-standard IoT operation systems (OSs), defined here as any OS other than the general-purpose Linux OS, traditionally used in high-end IoT devices [83], are considerably more task-oriented. Therefore, in view of attack detectability, the normal traffic patterns of IoT devices which run these task-oriented lightweight OSs should be more predictable and thus, abnormalities should stand out more clearly. In addition, if an IoT device supports the addition of apps (as in smart TVs and smartwatches), a larger variety of communication servers, protocols, packet sizes, and IATs is foreseen, thus traffic predictability is likely to decrease. In contrast to IoT devices which produce network traffic very rarely, sometimes triggered only upon predefined scenarios such as the detection of motion or smoke, other devices communicate more frequently, e.g., for activity monitoring, tracking temperature or humidity, and so forth. Rationally, the less continuous normal traffic is, the more prominent any attack-related anomaly is expected to be.

2) INT: User Interaction: The user interaction (UA) features are "softer" and less technical. Nevertheless, as UA is strongly associated with traffic generation [50], [51], we find that it is a fundamental component of the taxonomy. In this regard, more users logically means more usage and traffic patterns and thus more variability in the detectable attack features [78]. For instance, the typically short [17] packet IAT in flooding attacks as part of a DDoS is less likely to come through as a traffic anomaly when more variability is present in the packet IAT, due to multiple users and UA patterns. Additionally, for IoT devices which can be interacted with not only via an app but rather also directly (e.g., a smartwatch) and/or via a web application (e.g., a smart webcam), more traffic variability is expected to be found and obscure attack-related traffic anomalies.

C. NT: Network

These *direct features* resemble the previously mentioned *detectable attack features* [78], and they are calculated using captured traffic data.

1) INB and OUT: Inbound and Outbound Network Traffic: Among the direct/detectable features, inbound network traffic features describe the traffic whose destination is the IoT device in terms of packet count and size, encryption, etc. The outbound network traffic features are similar to the inbound ones, however in this case, the IoT device is the source of the traffic for the features rather than the destination. In

past research, e.g., [12], [78], [79], [84], these features were found informative for IoT attack detection. For instance, most normal IoT packets are sent at regular time intervals [77] (e.g., automated network activities). In contrast, these intervals are dramatically shorter [17] for most DoS attack traffic and thus informative. The authors of [17] also observed that as opposed to the size variations of normal traffic, TCP SYN flood packets are quite small (typically under 100 bytes), largely in order to maximize the number of connection requests per second for victim exhaustion.

2) SRD: Communication Sources & Destinations: As opposed to PCs and smartphones, IoT devices normally communicate with a limited number of endpoints (IP destination) [85], e.g., for (de)activation from the cloud, retrieving firmware updates, and logging their status; thus, an increased number of destination endpoints might be indicative of attack traffic (e.g., sending messages to the botnet victims), while extraneous source IP addresses might indicate logging attempts into the device, as in Mirai [28]. Moreover, the set of destination IPs rarely changes over time [17]. Also, the traffic associated with DDoS attacks has different protocol distributions (roughly equivalent to the number of destination ports), and they also include fewer total protocols [17], thus highly likely to be informative for IoT attack detection.

VI. PROPOSED METHOD FOR COMPUTING THE D-SCORE OF AN IoT MODEL FOR AN ATTACK SCENARIO

Our taxonomy includes various features that are potentially relevant, either directly or indirectly, to the ability of an AIDS to detect attacks which involve IoT devices. Understanding the importance of those features to the D-Score of various attack scenarios is a crucial task, as the contribution of features may not be the same for each attack scenario. Two main approaches can be considered for this kind of problem; one is data-driven and the other is expert-based. In order to reduce the considerable investment associated with the former approach (see Subsection VIII-C) and still obtain a robust and highly generalizable detectability assessment mechanism, we suggest using an expert-based approach. As part of the suggested method, cyber-security experts rank the taxonomy's features with respect to attack scenarios of interest, with the aim of producing feature weights. However, ranking multiple criteria simultaneously is a difficult task for an individual to accomplish, often resulting in an inaccurate ranking [86]. The *analytical hierarchical process (AHP)* method was introduced by Saaty [23] to moderate this drawback. AHP represents the problem space as a hierarchical structure in which the experts only compare pairs of elements at the same level of the hierarchy, thus reducing the amount of comparison needed to a more reasonable amount and resulting in a more precise ranking. The number of pairwise comparisons among n elements in a given level using AHP is calculated as $\frac{n \cdot (n-1)}{2}$. Based on the pairwise comparisons done by each expert, the AHP computes a vector of feature weights [87], which is eventually averaged over all the experts.

Essentially, our method for quantifying IoT attack detectability includes two key phases which conform with AHP's requirements:

Phase A Deriving a detectability model for an IoT attack scenario		Phase B Computing the detectability score of a given IoT model for an attack scenario
Performed only once	Performed for each A	Performed for each M
Input: none	Input: IoT attack scenario (A)	Input: IoT model (M)
Step 1: Develop a taxonomy for IoT attack detectability	Step 3: Have security experts complete the questionnaire	Obtain values for the taxonomy features (X) from M's spec sheet, FCC ID DB, MUD files, and captured traffic
Step 2: Devise a questionnaire for feature ranking	Step 4: Analyze the security experts' responses using the AHP method	Calculate the D-Score by multiplying M's taxonomy feature (normalized) values with the respective A's feature weights W
Output: Taxonomy (X) & questionnaire	Output: Feature weights (W)	Output: D-Score (A, M)

Fig. 7: Flow chart of the key steps in the two phases of our proposed method.

Phase A: Deriving a detectability model for an IoT attack scenario A. This model reflects the contribution (weight, denoted as W) of each feature from the taxonomy to the detectability of the corresponding attack scenario. Note that the weights W are calculated *only once for a given attack scenario* A, and can later be applied to multiple IoT models Ms in order to assess the detectablity of A on each of them.
Phase B: Computing the detectability score of a given IoT model M for an attack scenario A, using the detectability model derived in phase A. This is done by evaluating the relevant features (denoted as X) for the given IoT model, and using the feature weights to calculate the D-Score as a weighted sum ($W \cdot X$).

Once phase B is complete, if the calculated D-Score for an attack is low (i.e., it would be hard to detect the attack) and yet the organization decides to deploy the IoT device, proper countermeasures are encouraged to be taken in order to mitigate cyber-risks. These can include [18] *(a)* changing default passwords, *(b)* installing security patches, *(c)* disabling UPnP, *(d)* monitoring certain TCP/IP ports (e.g., 23 or 2323 to defend from attempts to gain unauthorized control via Telnet), *(e)* implementing additional network isolation, *(f)* updating firewall rules, etc. Further countermeasures can be found in, e.g., [88], [89], from which the organization should implement the ones most relevant to the anticipated attack.

Our proposed method is envisioned as a means of facilitating risk assessment and mitigation processes, and support IoT deployment decisions which are often made by a wide range of users (from home users to enterprise network administrators). When using our method, a simple comparison between the D-Score labels (Fig. 1b) of competing IoT models is all that is needed, similar to the manner in which a consumer would compare energy efficiency rating labels before purchasing an appliance. However, in order to associate an appropriate D-Score label with an IoT model, we advise implementing the

two phases of our proposed method, as outlined above and illustrated in Fig. 7. Phase A, the research-oriented phase, can be performed by a specialized cyber-security research institute. In comparison, phase B, which is more technical, should be performed for every newly released IoT model by a central standards institution, in order to ensure the uniformity and trustworthiness of the D-Score labels. Further details about the steps comprising the two phases of our method are provided in the subsections that follow; a step-by-step quantitative demonstration of our proposed method is provided in Section VII.

A. Deriving a Detectability Model for an IoT Attack Scenario

The first phase of our method consists of the following four steps, where steps 1 and 2 are generic (i.e., performed and fine-tuned only at the beginning), while steps 3 and 4 are performed for each A separately.

1) **Developing a taxonomy for IoT attack detectability**: The taxonomy (described in Section V) is a hierarchical decomposition of the attributes of IoT models into a list of categories, sub-categories and features. We designed our taxonomy to be generic so it could support a variety of IoT models and attack scenarios.

2) **Devising a respective questionnaire used for feature ranking**: The questionnaire is comprised of questions for cyber-security experts, where each question compares two elements (from the taxonomy) in terms of their relevance and relative importance to detecting a specific IoT attack scenario. For D-Score's taxonomy, the default AHP implementation would require the experts to compare three pairs of categories (i.e., $\frac{3 \cdot (3-1)}{2}$), five pairs of sub-categories ($1 + 1 + 3$), and also 57 pairs of features ($1 + 3 + 3 + 10 + 15 + 15 + 10$), resulting in a total of 65 comparisons. Although less than the number of all pairwise comparisons required in a naïve approach, it still requires a significant amount of expert time. Thus, to further reduce the number of comparisons, we propose the following process:

   - *Preliminary filtering of categories:* The categories are not assigned weights, but are only used by the experts to focus their attention on the sub-categories and features that really matter for attack detection. As discussed in Section VII-A, this novel procedure was used broadly by the experts in our study and saved effort on their part and much computation.
   - *Preliminary filtering of sub-categories in the selected categories:* Irrelevant sub-categories may remain within a potentially informative category (i.e., not filtered out), so we provide the ability to filter them out as well, prior to the comparison among sub-categories and features.
   - *Pairwise comparisons among sub-categories:* In order to assign weights to the remaining sub-categories, each expert is required to rank pairs of sub-categories.
   - *Pairwise comparisons among the features in each sub-category*: Similarly, pairwise comparison of features is required, however only in the sub-categories previously deemed as relevant by the expert.

3) **Having security experts answer the questionnaire**: For a given attack scenario A (e.g., DDoS flooding), cyber-security researchers and practitioners who are familiar with A are requested to use their experience and best judgment to respond to the pairwise comparisons in the questionnaire.

4) **Analysing the security experts' responses using the AHP method to assign feature weights**: The experts' responses are analyzed using AHP, resulting in a weight assigned to each feature of the taxonomy. This weight is the average of the respective weights provided by all of the experts who completed the questionnaire for the given attack scenario A. The (averaged) feature weights W determine: *How informative each feature is for detecting the selected attack scenario on any given IoT model? Which are the most informative features for each attack scenario?*

A challenge arises when the features X do not share the same range of values (see Table IV). To overcome this challenge, feature normalization was used in prior research by, e.g., expert-based ranking [86] (i.e., another level at the bottom in the taxonomy's hierarchy), manual division into equivalent groups [23], [90], or automatic binning [91]. In contrast, for normalization we propose using the hyperbolic tangent (tanh) function with some fine-tuning, as it is *(a)* systematic and consistent, *(b)* requires less human intervention, *(c)* is non-linear and *(d)* its slope can be easily adjusted. We use the following notation to describe our proposed method for feature normalization:

$i \in [1, |features|]$: Index of a feature in the taxonomy
$j \in [1, |attacks|]$: Index of an attack scenario of interest
x_i: Original value of the i^{th} feature
x_i^{min}, x_i^{max}: Minimum and maximum values, respectively, in the expected range of x_i
$x'_{ij} \in [0, 1]$: Normalized value of x_i for attack scenario j
$\alpha_{ij} \in \mathbb{R}$: Coefficient to control the shape of tanh function
$\beta_{ij} \in \mathbb{R}$: Coefficient to control the value of α_{ij}
$\delta_{ij} \in \{-1, 1\}$: Direction of influence on a given attack scenario j for increasing values of x_i
w_{ij}: Weight assigned to feature i and attack j using AHP

To normalize the original value of $x_i \in [x_i^{min}, x_i^{max}]$ into $x'_{ij} \in [0, 1]$, first we calculate α_{ij} as in Eq. (2), such that the range of x_i is taken into account. For example, x_i^{max} of percentage and binary features (e.g., ENCI and BATT, respectively) is $1 = \log_{10}(0)$, the number of sensors (NSNS) is typically lower than $10 = \log_{10}(1)$, the number of packets per inbound session (PCKI) is typically lower than $100 = \log_{10}(2)$, etc.

$$\alpha_{ij} = \frac{1}{\beta_{ij} \cdot 10^{\{\log_{10}(x_i^{max})-1\}}} \qquad (2)$$

Table IV outlines the feature ranges which are relevant for our research and the associated α values, calculated using Eq. (2). For simplicity, we used $\beta = 5$ for all of the feature ranges, so that each range of x_is (i.e., a row in this table) has its α.

After deriving the α values we determine δ_{ij}, as increasing values of a feature may increase the D-Score for a given attack scenario (and then $\delta_{ij} = 1$), but decrease it for others ($\delta_{ij} =$

TABLE IV: Derived α for varied feature value ranges ($\beta = 5$).

x_{min}	x_{max}	$Log_{10}(x_{max})$	α	*Features*
0	1	0	2	ENCI, BATT
0	10	1	0.2	NSNS, FINT
0	100	2	0.02	PCKI, IATO
0	1,000	3	0.002	CPUS, IATI

-1). Finally, for each original x_i we calculate the normalized value x'_{ij}, as in Eq. (3).

$$x'_{ij} = \begin{cases} tanh(\alpha_{ij} \cdot \delta_{ij} \cdot x_i) & \text{if } \delta_{ij} = 1 \\ 1 + tanh(\alpha_{ij} \cdot \delta_{ij} \cdot x_i) & \text{if } \delta_{ij} = -1 \end{cases} \qquad (3)$$

As shown in Fig. 8, all of the x's converge to one, each with a proper steepness based on its α. That is, the convergence is quicker for x_is whose x_i^{max} is lower.
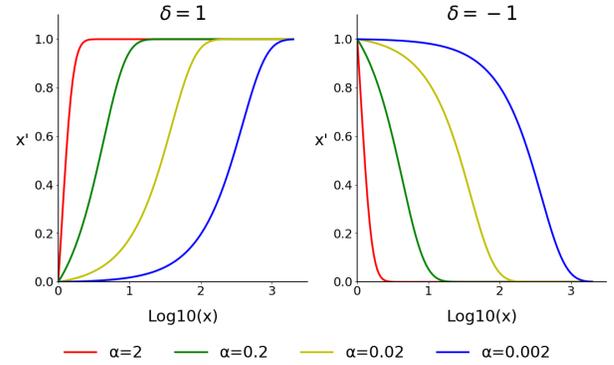
Fig. 8: Feature normalization for various ranges.

B. Computing the Detectability Score of a Given IoT Model for an Attack Scenario

Given combinations of *(a)* the characteristics of an IoT model (as specified in the taxonomy), and *(b)* an attack scenario of interest, this step quantitatively assesses the ease of attack detection. This step's input (i.e., the required technical and behavioral information), can be obtained from the device's spec sheet, the FCC ID database [92], MUD files [93], and from locally captured network traffic. Then, for an attack scenario A_j and IoT model M with features X'_j (normalized for A_j), the D-Score is calculated using Eq. (4).

$$D - Score(A_j, M) = W_j \cdot X'_j = \sum_{i=1}^{|X|}(w_{ij} \cdot x'_{ij}) \qquad (4)$$

VII. QUANTITATIVE EVALUATION

A. Implementation of the Questionnaire

In order to quantitatively evaluate our method for D-Score calculation, we designed and implemented an online question-naire [94] using the open-source survey tool LimeSurvey [95]. The first part of the questionnaire presented some background information and guidelines to the respondent, and in the second part, the respondent was asked to anonymously provide

TABLE V: Evaluated IoT attack scenarios

Abbreviation	Description of the IoT attack scenario	Ref.	Responses
C&C communication	The IoT device is communicating with a C&C server	[2]	10
DDoS flooding	The IoT device is executing a flood attack as part of a DDoS campaign	[2]	13
Data exfiltration	The IoT device is exfiltrating data	[30]	12
Bot scanning	The IoT device is being scanned / brute-forced by a bot (as used by Mirai for propagation)	[18]	5

demographic details regarding their professional experience and educational level. In the third part, in order to strengthen the validity of the results, we asked the respondent to select an IoT attack scenario from the list of attacks presented in the threat model (see Table V), based on his/her knowledge.

In the fourth part of the questionnaire, as part of the novel process of preliminary filtering (described in Section VI-A), the respondent was asked to *(a)* select only the categories relevant for the detection of the attack scenario he/she selected and then to *(b)* select only the relevant sub-categories within the categories selected in *(a)*. As summarized in Table VI, this ability to select specific relevant categories and sub-categories was widely exercised by the respondents. For example, of those who addressed the *Data exfiltration* scenario, only 50% kept the *SB* category, (i.e. the remaining 50% filtered it out), leading to a reduction of 11 AHP-style pairwise comparisons among the sub-categories, three comparisons among the features in the *FNC* sub-category and 10 more among the features in the *INT* sub-category, i.e., a total reduction of 24 (!) pairwise comparisons. Moreover, even if a category was not filtered out, often the respondent did not select all of the sub-categories in the category. For example, for *Data exfiltration* 100% of the respondents kept the *NT* category, however 0% kept the *INB* sub-category, thus reduced their total comparisons by 12.

In the fifth part of the questionnaire each respondent was asked to perform pairwise comparisons among the remaining sub-categories as well as among the features within these sub-categories. The scale of each comparison ranges from -5 to +5, where "5" signifies extreme importance, "3" reflects essential or strong importance, and "0" means equal importance. The "-" and "+" signs determine which is more important, i.e., the left-hand side or the right-hand side of the comparison, respectively.

B. Collection and Preprocessing of the Responses

To collect responses, we reached out via email to cyber-security experts from our academic and professional circles in Europe, the Middle East, and the Far East. Each expert received a link to our questionnaire [94], accompanied by concise instructions. After the experts completed the online

Fig. 9: Background of the cyber-security experts.

questionnaires, we exported their responses to a central .csv file. Due to the preliminary filtering process used in the fourth part, there were missing values in this .csv file for the categories and sub-categories filtered out by the experts. Therefore, as part of our data preprocessing, we filled in the missing values as follows: If a category or sub-category was kept but not the other one, a value of ± 5 was added, depending on the side of the pairwise comparison. If neither category or sub-category was kept, then a value of 0 was added, meaning they were equally unimportant.

The median amount of time it took for a respondent to complete the questionnaire was only 13 minutes. Based on the collected demographic details, in Fig. 9 we can see that our respondents have more academic experience (7.64 years, on average) than industry experience (4.33), and more experience in defensive cyber-security (2.75 years, on average) than offensive cyber-security (1.23). In addition, of the 40 respondents, 18 have a doctoral degree, and 22 have a master's degree (many of whom are currently enrolled in Ph.D. programs) or less (mostly M.Sc. or B.Sc. students at the time of response collection).

C. Calculation of the Weights Using AHP

To analyze the data using the AHP methodology, we implemented a Python script which is capable of receiving the pairwise comparisons of a respondent as input and producing the weights of the features and sub-categories as output. We applied this script to each response separately and then calculated the average weight for each of the four groups of responses corresponding to the four attack scenarios.

As noted above, we collected multiple responses from various cyber-security experts for each of the four IoT attack

TABLE VI: Percentage of respondents who kept each category and sub-category during the preliminary filtering process.

Attack scenario	Categories			Sub-categories						
	HW	SB	NT	SNA	RSR	FNC	INT	INB	OUT	SRD
C&C communi.	40%	60%	**90%**	10%	20%	30%	20%	70%	**90%**	**80%**
DDoS flooding	54%	62%	**100%**	15%	31%	46%	38%	46%	**85%**	**85%**
Data exfiltration	50%	50%	**100%**	33%	33%	25%	25%	0%	75%	**83%**
Bot scanning	40%	80%	80%	0%	40%	**80%**	20%	60%	80%	80%
Average	46%	63%	**93%**	15%	31%	45%	26%	44%	**82%**	**82%**

Fig. 10: Levels of agreement among the cyber-security experts for the evaluated attack scenarios.

scenarios. As an indication of the responses' validity and robustness, for each attack scenario we measured the *agreement level* between each pair of weight vectors W in terms of the average cosine similarity, defined in Eq. (5), similarly to [96].

$$Cosine\, Similarity(A, B) = \frac{A \cdot B}{\|A\|\, \|B\|} \quad (5)$$

As evident in Fig. 10, the agreement level as to the relative importance of the sub-categories is very high. It is also evident that the agreement level for the *Data exfiltration* scenario is moderately lower than the other three, maybe due to a lower level of attack scenario clarity. That is, *C&C communication*, *DDoS flooding* and *Bot scanning* may all be very unambiguous and well-known to cyber-security experts. In contrast, there is a chance that the respondents interpreted the *Data exfiltration* scenario in different ways, thus ranked the related sub-categories differently from one another. Regarding the weighing of features, it is apparent how the agreement levels for all of the attack scenarios are shifted downwards, compared to the sub-categories. The reason may be that pairwise comparisons among a large number of distinct elements (30 features) may be more complicated than comparing seven sub-categories to one another.

The feature weights W_j for each evaluated IoT attack scenario A_j are presented in Fig. 11. Each weight is an average of the corresponding weights given by all of the experts. In this heat map, the color is normalized for every row separately so that for every attack scenario we can easily observe the

features that are found most informative by the experts. On the one hand, features like *CCOM* seem very informative for all of the attack scenarios evaluated, while on the other hand, features like *NUSR* or *ENCI* may be of negligible importance for detecting any of the attack scenarios evaluated. Though uninformative for these scenarios, we advise against removing them completely from the taxonomy, since they may be relevant for other attacks to be addressed in the future. In between, there are features that are deemed important for some attacks and not for others. For instance, *PCVI* has the largest weights for *bot scanning* detection and relatively low weights for the other scenarios, likely because this is where anomalies would be found once this attack scenario is realized. Interestingly, for all of the evaluated attacks, the indirect features under the *Hardware* category (i.e., NSNS, NACT, CPUS, MEMS, and BATT) were assigned relatively high Ws, meaning that this novel part of our taxonomy was found to be informative by the experts for IoT attack detection, even though they are not direct/detectable attack features [78].

Fig. 12 presents the weights calculated for the sub-categories. As expected, the sub-categories which contain the direct/detectable features (i.e., the *Inbound* and *Outbound network traffic* as well as the *Communication sources and destinations*) were ranked as more informative for attack detection than the indirect features. Among them, the *Inbound network traffic* was given the highest weight for *Botnet scanning*, likely due to the nature of this scenario, in which bots try to communicate with (i.e., scan) the IoT device. Nevertheless,

	NSNS	NACT	CPUS	MEMS	BATT	STOS	ADAP	CCOM	NUSR	FINT	DINT	NUIS	SRNG	IATI	PCKI	PCSI	PCVI	ENCI	UDPI	IATO	PCKO	PCSO	PCVO	ENCO	UDPO	DSIP	DSPR	WLPR	SRIP	OPPR
Bot scanning	0.027	0.027	0.034	0.036	0.046	0.054	0.049	0.082	0.017	0.02	0.02	0.016	0.015	0.028	0.033	0.039	0.06	0.019	0.025	0.024	0.025	0.041	0.022	0.02	0.024	0.042	0.031	0.038	0.027	0.058
DDoS flooding	0.039	0.03	0.046	0.056	0.046	0.033	0.036	0.059	0.022	0.037	0.028	0.02	0.017	0.031	0.022	0.021	0.03	0.02	0.026	0.035	0.049	0.024	0.033	0.016	0.033	0.053	0.032	0.028	0.044	0.035
Data exfiltration	0.077	0.048	0.048	0.05	0.053	0.03	0.04	0.07	0.02	0.027	0.023	0.029	0.032	0.016	0.016	0.016	0.016	0.016	0.016	0.026	0.043	0.024	0.037	0.045	0.024	0.047	0.026	0.032	0.023	0.029
C&C communication	0.03	0.036	0.045	0.042	0.043	0.05	0.043	0.061	0.021	0.025	0.021	0.02	0.02	0.02	0.024	0.026	0.024	0.027	0.032	0.021	0.032	0.024	0.033	0.035	0.049	0.059	0.029	0.03	0.039	0.038

Feature

Fig. 11: Weights of the features across the evaluated attack scenarios
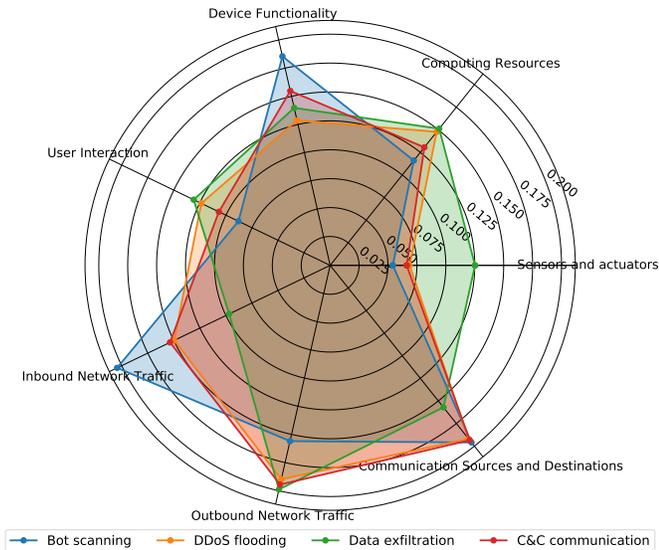
Fig. 12: Weights of the sub-categories across the evaluated attack scenarios.

the sub-categories which involve the indirect features were assigned non-negligible weights. This is most apparent for the *Sensors and Actuators* sub-category, which was found to be informative for detecting *Data exfiltration* attacks. This is probably because higher values of the *Number of sensors* feature lead to more frequent and versatile outbound traffic (thus less predictable), which could hide this attack.

D. Utilizing the Weights to Calculate D-Scores

Having calculated the weights W using AHP, in this subsection we demonstrate how an organization can utilize them to calculate D-Scores for alternative, competing, IoT models and decide accordingly which to deploy. For that, we purchased four security cameras whose models were studied in [12]. We deployed them in our lab and collected their static and dynamic values which correspond to the 30 features in the taxonomy. Assuming that *DDoS flooding* is the primary attack scenario to defend from, we calculated the respective D-Scores. Table VII presents the resultant D-Scores and also compares them to a key attack detection metric obtained from [12]. In that paper, several anomaly detection algorithms were used to detect flooding attacks carried out by Mirai and BASHLITE botnets from commercial IoT devices, including Isolation Forest [103] (abbreviated as IF in Table VII), Local

Outlier Factor [104] (LOF) and One-Class Support Vector Machine (OCSVM) [105]. The anomaly decisions were taken in that paper based on a *sequence* of traffic metadata feature vectors, using a majority vote on a moving window whose length was optimized for each IoT model separately. The authors of [12] demonstrated excellent detection results in terms of true positive rate (TPR) and false positive rate (FPR), however with a cost of relatively lengthy moving windows (equivalent to long detection times), treated here as a metric for attack detectability.

As evident in Table VII, for the DDoS flooding attack we found a strong negative correlation between the (optimized) window sizes and the D-Scores for each evaluated algorithm. This means that, as expected, D-Scores which we attained using an expert-based method are aligned with an attack detectability metric obtained using a data-driven method, such that longer window sizes (i.e., extended detection times) are expressed via lower D-Scores. In addition, among the four security camera models in Table VII, the D-Score of [26] was found to be the lowest, meaning that worse detection performance is expected from this IoT model for DDoS flooding attacks, compared with the other models. This finding corresponds well with [12], where the optimized window size is the highest for this model for all the evaluated algorithms.

As described in Section I, a (continuous) D-Score can be translated into a (categorical) detectability label. If we employ the same number of categories as in Fig. 1b (inspired by [8]), and use equal-width binning for discretization, then seven bins are attained, each is $1/7 = 0.143$ wide. In that case, an IoT model whose $D - Score \in [0.857, 1.00]$ would be labeled as "A" (most highly detectable, thus advised for deployment). Regarding the IoT models in our evaluation (i.e., the four security cameras), Table VII indicates that two are labaled as "E" since their $D - Score \in [0.286, 0.429]$, and the other two as "D" (superior detectability of DDoS flooding attacks, $D - Score \in [0.429, 0.571]$).

Table VII also provides key features whose values differentiate among the four alternative security cameras. For instance, when comparing the security camera with the lowest D-Score [26] to the one with the highest [99], it is evident how in most cases lower D-Scores are associated with network traffic that is more frequent and variable, e.g., a higher tendency to communicate continuously (*CCOM*), shorter inter-arrival times of outbound sessions (*IATO*), higher variability of outbound packet size (*PCVO*), and higher number of destination ports (*DSPR*). Another interesting finding is that among the 30

TABLE VII: D-Scores of alternative security cameras studied in [12] in the context of a data-driven anomaly-based method to detect DDoS flooding attacks executed by IoT devices.

Evaluated IoT models			Detectability		Algorithms' window sizes			Key differentiating features			
Ref.	Manufacturer	Model num.	Score	Label	IF	LOF	OCSVM	CCOM	IATO	PCVO	DSPR
[26]	Provision	PT-838	0.429	E	170	1,300	350	0.900	8.655	142.033	12
[97]	Simple Home	XCS7-1003	0.433	E	170	600	250	0.820	25.035	131.895	10
[98]	Provision	PT-737E	0.462	**D**	96	**200**	91	0.960	**15.673**	**122.780**	**9**
[99]	Simple Home	XCS7-1002	**0.477**	**D**	**61**	300	**86**	**0.370**	25.060	125.956	40
Correlation with the D-Score					-0.997	-0.809	-0.943				

TABLE VIII: DDoS flooding attack detection: Correlation between expert-based D-Scores of various IoT device types and models, and data-driven anomaly-based performance metrics obtained in [12].

Evaluated IoT models				Detectability		Algorithms' window sizes		
Ref.	Type	Manufacturer	Model num.	Score	Label	IF	LOF	OCSVM
[100]	Baby monitor	Philips	B120N/10	0.424	E	170	1,800	300
[26]	Camera	Provision	PT-838	0.429	E	170	1,300	350
[97]	Camera	Simple Home	XCS7-1003	0.433	D	170	600	250
[101]	Doorbell	Danmini	WF 720P	0.445	D	16	300	141
[98]	Camera	Provision	PT-737E	0.462	D	96	200	91
[102]	Doorbell	Ennio	Bell	0.470	D	11	450	36
[99]	Camera	Simple Home	XCS7-1002	0.477	D	61	300	86
Correlation with the D-Score						-0.762	-0.771	-0.930

taxonomy features, the ones which differentiate the most between the normal behavior of the above two camera models also conform with past research. For instance, the top three features in terms of absolute difference in x' (namely IATI, PCKI and PCVI) have been found informative by [79] for IoT attack detection. Similarly, IATI and PCVI have been selected by [17] as two of the three most important features, using Gini Score.

E. Extension of the DDoS Flooding Attack Detectability Evaluation to Non-Competing IoT Models

In subsection VII-D we focused on a use-case in which D-Scores are calculated for alternative (competing) IoT models of the same type, to prioritize and determine: which smart camera to deploy, assuming that the organization aims at quick and accurate detection of DDoS flooding attacks. In the current subsection, summarized in Table VIII, we extend our evaluation to another use-case, in which the D-Score is used for black/white listing of non-competing IoT devices. That is, an organization's security policy can set a threshold on the D-Score, and only the IoT devices whose D-Scores surpass this threshold are allowed to be connected to the organization network. As evident in Table VIII, which adds a smart baby monitor and two doorbells to the four previously-analyzed cameras, strong negative correlations are maintained between our expert-based D-Scores and the data-driven performance metric from N-BaIoT [12]. Despite the larger variability in IoT device types and models, lower D-Scores are still associated with longer window sizes (detection times), most strongly for the OCSVM-based AIDS.

F. Extension of the Detectability Evaluation to Additional IoT Attack Scenarios

Until now, in Subsections VII-D and VII-E, our quantitative evaluation concentrated only on one of the four attack scenarios we address in this paper (outlined in Table V), namely DDoS flooding attack. The reason is that this is the only attack scenario for which we have data-driven performance metrics [12] to compare with our expert-based approach. In this subsection, for completeness, we calculate the D-Score for the same IoT models evaluated in Subsection VII-E, this time for the remaining attack scenarios, namely bot scanning, data exfiltration and C&C communication. The summary of this evaluation is portrayed in Table IX, where for each IoT type (camera, doorbell and baby monitor) we highlighted *(a)* the highest D-Score for each IoT attack scenario, *(b)* the highest average D-Score, and *(c)* the highest minimum D-Score. The latter is highlighted in order to recognize the IoT model whose minimum D-Score for any attack scenario of interest is higher than the other (competing) IoT models. In real-world settings, this *maximin criterion* [106] could support procurement and deployment decisions, based on maximizing the minimum chance of detecting an attack carried out.

The results (see Table IX) show that when contemplating which smart camera to deploy, Simple Home XCS7-1002 consistently has the maximum chances for attack detection, and so is Ennio Bell (for smart doorbells) and the Philips B120N/10 baby monitor. This may be due to the fact that Simple Home XCS7-1002 is the least continuously-communicating camera (CCOM), and also has the largest inter-arrival time of outbound sessions (IATI). Similarly, Ennio Bell has the lowest CCOM and IATI, as well as the smallest number of unique destination ports (DSPR). Overall, the maximin D-Score of Philips B120N/10 was the lowest of all IoT models evaluated,

TABLE IX: Calculation of D-Scores for IoT models studied in [12] and the IoT attack scenarios outlined in Table V.

IoT model				D-Score					
Type	Ref.	Manufacturer	Model num.	DDoS flooding	Bot scanning	Data exfiltration	C&C communication	Avg.	(Maxi-)min.
Camera	[98]	Provision	PT-737E	0.462	0.426	0.493	0.470	0.463	0.426
	[26]	Provision	PT-838	0.429	0.408	0.478	0.446	0.440	0.408
	[99]	Simple Home	XCS7-1002	**0.477**	**0.460**	**0.502**	**0.489**	**0.482**	**0.460**
	[97]	Simple Home	XCS7-1003	0.433	0.418	0.466	0.461	0.445	0.418
Doorbell	[101]	Danmini	WF 720P	0.445	0.418	0.460	0.449	0.443	0.418
	[102]	Ennio	Bell	**0.470**	**0.429**	**0.498**	**0.480**	**0.469**	**0.429**
Baby monitor	[100]	Philips	B120N/10	**0.424**	**0.403**	**0.437**	**0.398**	**0.416**	**0.398**

meaning that any of the attack scenarios addressed in this paper would be the hardest to detect. We attribute this low attack detectability score mostly to the largest percentage of outbound encrypted traffic (ENCO) and to the highest number of sensors (NSNS), which include video, motion detection, noise detection, temperature sensor and humidity sensor.

Regarding the findings in this subsection, we note that except for the DDoS flooding attack, the D-Score results for the remaining attack scenarios are expert-based estimations and need to be validated using a data driven approach, once proper data and labels are available. We leave this for future research.

G. Analysis of the Sensitivity to the Quality of Questionnaire Responses

The experimental results discussed thus far were attained based on *all* of the questionnaire responses we received, as reported in Table V. However, not all of the respondents have the same background, experience and level of expertise in cyber-security, IoT or attack detection. Therefore, a question arises regarding the proper selection of responses, such that safe feature weights are eventually drawn, to be used for D-Score calculation. To address this question, in this subsection we analyze the sensitivity of our method to using only *subsets* of the responses, based on varying thresholds of response quality criteria. As the range of the explicit criteria such as the years of experience (see Fig. 9) is relatively limited, we chose to use the more implicit response quality criterion known as *consistency ratio* (CR). The CR [23] is calculated for a group of pairwise comparisons (e.g., comparisons among multiple features within a sub-category), and it reflects the extent of inconsistent judgments made by evaluators in tasks of multiple pairwise comparisons. Typically, CR levels of up to 0.1 are acceptable [87], [48].

In Table VII, strong negative correlations are demonstrated between the (expert-based) D-Scores of the four security camera models and their (data-driven) optimized window sizes in N-BaIoT [12], based on all thirteen questionnaire responses to the *DDoS flooding* attack scenario. To examine the sensitivity of our method to CR, we *(a)* calculated the CR for each of the seven sub-categories in each response, *(b)* calculated the Mean CR for each response, *(c)* calculated D-Scores for varying thresholds of Mean CR, i.e., based on a subset of responses whose Mean CR is below the threshold, and *(d)* calculated the correlation of the D-Scores with the window sizes optimized for each algorithm. The results of this sensitivity analysis are summarized in Table X, and they show some variability of the correlation strength. On the one hand, for the *IF* algorithm the strongest correlation was attained when the D-Scores were calculated using all thirteen responses, including two responses whose Mean CR are 0.172 and 0.198. On the other hand, for the two other algorithms, no more than four (!) responses (the most consistent ones) were required to reach the best results. Looking forward, for feature weights calculation in future applications of the D-Score methodology, remaining with the commonly used threshold of 0.1 seems like a reasonable option. In our case it yields sub-optimal yet strong correlation with the data-driven approach.

TABLE X: Correlation of the D-Scores calculated for the *DDoS flooding* attack scenario with the optimized window sizes attained in N-BaIoT [12] as a function of the Mean CR threshold used to filter questionnaire responses

Range of Mean CR	Number of questionnaire responses	Correlation of the D-Scores with the window sizes		
		IF	LOF	OCSVM
0 - 0.033	4	-0.981	**-0.874**	**-0.978**
0 - 0.067	9	-0.989	-0.841	-0.965
0 - 0.100	11	-0.994	-0.829	-0.956
All	13	**-0.997**	-0.809	-0.943

VIII. DISCUSSION

A. IoT Model Complexity, Traffic Predictability and Attack Detectability

The motivation to conduct this research originated from the insight that in many cases, variability exists in the capability of an AIDS to detect the same cyber-attack on differing IoT models. In order to guide the investigation of the source of this variability in IoT attack detection performance, we raised two groups of research questions in Section I, namely *(a)* network traffic predictability and *(b)* IoT attack detectability, which complement one another. Our suspicion was that the (indirect) technical characteristics of an IoT Model (e.g., the number of sensors and actuators), as well as its designated functionality and intended modes of user interaction, have an effect on the (direct) network traffic features that are typically analyzed by AIDSs. We argued that IoT models which *(a)* are more technically complex, *(b)* have a more versatile functionality and *(c)* are interacted with more frequently and more diversely, are likely to express less deterministic traffic patterns. Consequently, a traffic-based profile of normal behavior would be less stable, such that the task of anomaly-based attack detection (which is the basic principle in AIDSs) becomes more difficult. In Section IV (and Appendix A) we explored the stability/predictability of IoT traffic patterns and, indeed, found correlations between them and the above-mentioned device complexity factors. For example, the number of sensors is positively correlated with the number of hourly outbound flows, and the CPU speed is correlated with the number of hourly unique destination ports. In Section V we introduced a novel taxonomy which incorporates the various static and dynamic features of an IoT model. Then, in Sections VI and VII, we presented a method to assess IoT attack detectability, and quantitatively evaluated it, respectively.

Altogether, we found statistical evidence to the link between the complexity of IoT models and their traffic predictability. In addition, our expert-based method for assessing attack detectability, which utilized the above-mentioned IoT model complexity and traffic predictability, demonstrated strong correlation with three anomaly detection algorithms from past (independent, data-driven) research.

B. Trade-Offs in Designing the Taxonomy

Due to the taxonomy design considerations that we outlined in Section V (e.g., being informative, flexible, self-explanatory,

and feasible to deploy), the selection and definition of some features might be sub-optimal for assessing IoT attack detectability. For instance, the *CPU* and *memory utilization per time frame* [19] could be more informative than the actual *CPU speed* or the *memory size* we eventually incorporated into the taxonomy. The reason for this is that in the case of a DDoS flooding attack, the utilization of these key computing resources might (suddenly) increase regardless of the maximum CPU speed or the total memory size. However, measuring the CPU and memory utilization for an IoT device (unlike a PC) is exceedingly challenging and might therefore make the calculation of D-Score infeasible. Another example of a potentially more informative feature is the type of sensors (motion detectors, microphones, etc.) as opposed to their quantity, since different sensor types typically have different traffic ratios and contents. The problem is that *sensor type* is a categorical feature, and encoding it into multiple binary features (e.g., *has a motion detector*, *has a microphone*, etc.) would result in a much more complex taxonomy, and would accordingly require many more pairwise comparisons and thus overburden the questionnaire respondents. Yet another feature that could assist in attack detectability quantification is the deployment location of the IoT device, e.g., a normally-crowded room or a rarely-attended warehouse. However, on the one hand this is another categorical feature that would increase the taxonomy's complexity, while on the other hand we already incorporated user interaction features in our taxonomy, so the deployment location itself becomes redundant.

All things considered, the taxonomy we present and use in this paper is the result of several trial-and-error design iterations, conducted in order to fine-tune the taxonomy while balancing the above-mentioned considerations and trade-offs. In future research, additional fine-tuning can be accomplished by re-defining existing features and/or adding possibly-missing ones. In any case, making such feature changes in the taxonomy (and accordingly also in the questionnaire) is an achievable task using the web-based questionnaire we developed.

C. Comparison of Approaches for IoT Attack Detectability Assessment: Expert-Based vs. Data-Driven

Hypothetically, an organization which already deploys an AIDS and wishes to use the AIDS to quantify the detectability of a certain attack A in advance could opt to employ a data-driven approach. Naturally, this option would allow a straightforward preliminary assessment which is based on the same conditions as in the projected deployment. Nonetheless, despite the anticipated high accuracy, employing a data-driven (anomaly-based) approach comes with a cost of a non-trivial overhead, as elaborated in Table XI. Among the substantial components of this overhead is the need to acquire a high-quality set of network traffic data which is representative of M as comprehensively as possible. Key steps in acquiring such a dataset are the implementation and execution of A, followed by feature extraction and rigorous ground truth labeling of every instance as "benign" or "malicious". We note that in theory, anomaly-based attack detection methods (as in AIDSs) necessitate only benign instances for training [9], as they are intended to capture only the normal traffic patterns of M, and then label any severe-enough abnormality as an attack. However, in practice, AIDSs tend to suffer from non-negligible FPR [9], [107], [108], [109], [110], [111], meaning that in too many cases, *rare-yet-benign* abnormal events are falsely identified as attacks. Consequently, redundant and potentially harsh countermeasures might be activated in response. To decrease the FPR which is associated with such (unsupervised) anomaly-based methods, one could use labeled data for model calibration or anomaly threshold tuning, as in semi-supervised [109], [110] or hybrid [111] approaches.

The substantial cost of acquiring labeled data with malicious instances for FPR reduction, can be saved if an organization chooses to use our expert-based method. Additional considerations should also be taken into account while contemplating between data-driven and expert-based approaches:

- *Ability to cope with zero-day attacks.* The above-mentioned cost of acquiring labeled data relies on the assumption that the malware of interest is available for local execution in a

TABLE XI: Comparison of the overhead required by the data-driven and the expert-based approaches

Description of overhead component	Approach		Remarks
	Data-driven	Expert-based	
Procurement of M device(s)	✓	✓	
Deployment of M device(s) in a lab	✓	✓	
Extraction of static features from, e.g., the spec sheet		✓	
Capturing of benign traffic data	✓	✓	
Implementation and execution of A attack	✓		For model calibration
Capturing of malicious traffic data	✓		For model calibration
Precise labeling of traffic data (malicious / benign)	✓		For model calibration
Extraction of dynamic features from the captured traffic	✓	✓	
Procurement and/or installation of an AIDS OR design, training and testing of an anomaly detection algorithm	✓		
Design and implementation of the questionnaire		✓	Can use our questionnaire
Collection of questionnaire responses from experts		✓	
Analysis of the questionnaire responses (i.e., weights calculation)		✓	

controlled environment, to be followed by capturing network traffic from M and labeling the extracted feature vectors. However, this is not always the case, as for new variants of malware or attacks, it might take time to capture A's code or replicate it.

- *Generalizability.* Even if the malware (whether original or replicated) *is* available for execution, the data-driven approach might over-fit the specific variant of the attack scenario, the (simulated) manner of attack execution, the network structure, the normal user interaction, etc., thus it might fail in generalizing to other environments or conditions. Also, with the data-driven approach, the detectability assessment is based on the specific algorithm(s) used for anomaly detection. In contrast, the feature weights that are trained for the expert-based approach are agnostic of any algorithm.

- *Scalability.* Calculating the D-Score of the same A for multiple Ms using a data-driven method would require to repeat the entire process of device deployment / attack execution / data capturing and labeling / model training and calibration, etc., multiple times, i.e., once for each M. In comparison, our method is more efficient in the sense that *(a)* the existing taxonomy and online questionnaire are generic and ready to be used, *(b)* phase A in our method (deriving a detectability model, see Subsection VI-A) is performed just once for an IoT attack scenario A, and *(c)* only phase B (computing the D-Score of an IoT model M for an attack scenario A, see Subsection VI-B) recurs for each M separately. Note that phase B merely requires collecting normal traffic data and some technical characteristics of M, and then running an automated script to calculate the final D-Score using AHP. Conversely, calculating D-Score of multiple As for a given M could be instantaneous in cases where phase A (attack-oriented) has already been performed for other Ms, and phase B (device-oriented) has been performed for the same M in the past, regardless of A's identity. In any case, we assume that any organization which chooses to incorporate D-Scores / labels into its security policies, whether calculated using a data-driven or expert-based method, would focus only on a limited set of As and Ms, i.e., the top most prevalent or impactful attack scenarios, and a relatively small number of alternative Ms.

- *Objectivity.* Product reviews posted in e-commerce websites could sometimes be manipulated or even fake [112]. On the one hand, our expert-based method also (partially) relies on human inputs. On the other hand, the participants in the envisioned deployment of our method are independent cyber-security experts, e.g., knowledgeable scholars, experienced practitioners and/or regulatory standardization organizations such as the FCC [113]. Thus, a high level of reliability and objectivity is anticipated.

When considering the use of a classic data-driven approach versus our expert-based approach, we make a distinction between two scenarios: In the first scenario, the organization interested in assessing the detectability of IoT attacks has *everything* it needs to perform this task. This includes, e.g., a list of IoT models to consider, a list of IoT-related attacks that are anticipated by and/or dangerous to the organization, enough data collected from the IoT models of interest (which is also accurately labeled with the specific malware variants), an existing attack detection system to be used continuously in research and production, dedicated research staff, the time and funding needed to conduct research while fine-tuning the detection models, and so forth. In this scenario, the organization is more likely to favor the data-driven approach. However, the second scenario may better reflect the reality of many organizations today; in this scenario the organization lacks sufficient IoT security expertise, because, for example, it is a small office or a non-tech company, and is willing to rely on researchers and expert judgements, accumulated and analyzed by an official standards organization. As described above, in this scenario there are various advantages to employing the expert-based approach, including the ability to cope with zero-day attacks, generalizability, scalability and objectivity.

IX. CONCLUSION AND FUTURE WORK

Based on device complexity and the associated traffic predictability, this study presents the first step toward explaining and quantifying the detectability of IoT-related attacks. In accordance with the ideas presented in [114] and in order to reduce the overhead associated with the data-driven approach, we opted for an expert-based approach. Our proposed method leverages the collective wisdom and experience of cyber-security researchers and practitioners, and can serve as an interim solution to *cold start* problems, in which preliminary attack detectability assessments are required, yet a large-scale high-quality correctly labeled dataset is not available.

In this study, we *(a)* quantitatively explored the predictability of IoT network traffic, which was taken for granted thus far, *(b)* designed a detailed taxonomy for IoT attack detection, *(c)* proposed a means of quantifying the values of the taxonomy's features for a given IoT model, *(d)* derived four detectability models which correspond to common attack scenarios, *(e)* evaluated our method for *DDoS flooding* attacks using commercial IoT devices, and *(f)* compared our results with past research.

Regarding network traffic predictability, we found that *(a)* the traffic of IoT devices is significantly more predictable than of non-IoT devices, *(b)* variability in traffic predictability does exist among disparate IoT models, and *(c)* the IoT traffic predictability is correlated with static device characteristics (e.g., the number of sensors). Concerning our non-trivial selection of an expert-based approach (rather than a data-driven one) to quantify in advance the detectability of IoT-related attacks, we found that *(d)* despite differing backgrounds and experience levels, the agreement among the cyber-security experts was relatively high. In addition, *(e)* the weights which were assigned to the sub-categories and features using our expert-based method align with the importance of corresponding features from past data-driven studies on IoT attack detection. Additional findings are that *f)* the D-Scores which were attained using our expert-based method are aligned with data-driven attack detection performance from past research, and *(g)* higher detectability is associated with

IoT models whose baseline traffic is easier to profile, i.e., they communicate less frequently, with less variability in packet sizes and with fewer destination ports.

In future research we will take additional steps toward further explaining and quantifying the detectability of IoT-related attacks, while focusing on the following challenges:

- Broaden the scope of quantitatively evaluating our method by means of, e.g., additional attack scenarios, more questionnaire responses to be collected from further security experts, and additional data-driven methods and experimental detection results to compare with.
- Analyze the ability of malicious actors to execute adversarial attacks. That is, an attacker who is familiar with the D-Score calculation method and parameters (α, β, δ, etc.) might use this knowledge to optimize an attack by, e.g., maximizing the frequency of outbound packets as part of a DDoS attack, while remaining "under the radar".
- If the calculated D-Score of a certain IoT model to an attack scenario of interest is relatively low, yet the organization decides to deploy it, an option we consider to explore in future research is to leverage an IoT model's D-Score and take its value into account while setting the anomaly detection threshold of *this* IoT model. That is, a deployed IoT model with a low D-Score is advised to use a low anomaly threshold, such that attacks would not be missed (i.e., to increase the TPR), while a different IoT model with a high D-Score can be associated with a relatively high threshold (to decrease the FPR).

ACKNOWLEDGMENT

REFERENCES

[1] L. Da Xu, W. He, and S. Li, "Internet of things in industries: A survey," *IEEE Transactions on industrial informatics*, vol. 10, no. 4, pp. 2233–2243, 2014.

[2] R. Hallman, J. Bryan, G. Palavicini, J. Divita, and J. Romero-Mariona, "IoDDoS-the internet of distributed denial of service attacks," in *2nd international conference on internet of things, big data and security*. Porto, Portugal: SCITEPRESS, 2017, pp. 47–58.

[3] Y. Yang, L. Wu, G. Yin, L. Li, and H. Zhao, "A survey on security and privacy issues in Internet-of-Things," *IEEE Internet of Things Journal*, vol. 4, no. 5, pp. 1250–1258, 2017.

[4] D. Demeter, M. Preuss, and Y. Shmelev, "IoT under fire: Kaspersky detects more than 100 million attacks on smart devices in H1 2019," Oct 2019 (accessed October 24, 2019). [Online]. Available: https://www.kaspersky.com/about/press-releases/2019_iot-under-fire-kaspersky-detects-more-than-100-million-attacks\-on-smart-devices-in-h1-2019

[5] H. Thompson and S. Trilling, "Cyber security predictions: 2019 and beyond," November 2018 (accessed October 24, 2019). [Online]. Available: https://www.symantec.com/blogs/feature-stories/cyber-security-predictions-2019-and-beyond

[6] L. Goasduff, "Gartner Says 5.8 Billion Enterprise and Automotive IoT Endpoints Will Be in Use in 2020," Oct 2019 (accessed October 24, 2019). [Online]. Available: https://www.gartner.com/en/newsroom/press-releases/2019-08-29-gartner-says-5-8-billion-enterprise-and-automotive-io

[7] M. Shirer, "The Growth in Connected IoT Devices Is Expected to Generate 79.4ZB of Data in 2025, According to a New IDC Forecast," 2019 (accessed October 24, 2019). [Online]. Available: https://www.idc.com/getdoc.jsp?containerId=prUS45213219

[8] Union, European, "Tyre label — Wikipedia, the free encyclopedia," 2019. [Online]. Available: https://en.wikipedia.org/wiki/Tyre_label#/media/File:EC_tyre_label_CA.svg

[9] P. Garcia-Teodoro, J. Diaz-Verdejo, G. Maciá-Fernández, and E. Vázquez, "Anomaly-based network intrusion detection: Techniques, systems and challenges," *computers & security*, vol. 28, no. 1-2, pp. 18–28, 2009.

[10] K. Zhao and L. Ge, "A survey on the Internet of things security," in *2013 Ninth international conference on computational intelligence and security*. Leshan, China: IEEE, 2013, pp. 663–667.

[11] F. A. Alaba, M. Othman, I. A. T. Hashem, and F. Alotaibi, "Internet of Things security: A survey," *Journal of Network and Computer Applications*, vol. 88, pp. 10–28, 2017.

[12] Y. Meidan, M. Bohadana, Y. Mathov, Y. Mirsky, A. Shabtai, D. Breitenbacher, and Y. Elovici, "N-BaIoT—Network-Based Detection of IoT Botnet Attacks Using Deep Autoencoders," *IEEE Pervasive Computing*, vol. 17, no. 3, pp. 12–22, Jul 2018.

[13] H. Bahşi, S. Nõmm, and F. B. La Torre, "Dimensionality reduction for machine learning based IoT botnet detection," in *2018 15th International Conference on Control, Automation, Robotics and Vision (ICARCV)*. Singapore: IEEE, 2018, pp. 1857–1862.

[14] S. Nõmm and H. Bahşi, "Unsupervised Anomaly Based Botnet Detection in IoT Networks," in *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*. Orlando, FL: IEEE, 2018, pp. 1048–1053.

[15] J. Blythe and S. Johnson, "The consumer security index for IoT: A protocol for developing an index to improve consumer decision making and to incentivize greater security provision in IoT devices," in *IET Conference Publications*, vol. 2018. London, UK: IEEE, 2018, pp. 1–7.

[16] P. G. Kelley, J. Bresee, L. F. Cranor, and R. W. Reeder, "A nutrition label for privacy," in *Proceedings of the 5th Symposium on Usable Privacy and Security*. Mountain View, CA: ACM, 2009, p. 4.

[17] R. Doshi, N. Apthorpe, and N. Feamster, "Machine learning DDoS detection for consumer internet of things devices," in *2018 IEEE Security and Privacy Workshops (SPW)*. San Francisco, CA: IEEE, 2018, pp. 29–35.

[18] E. Bertino and N. Islam, "Botnets and internet of things security," *Computer*, vol. 50, no. 2, pp. 76–79, 2017.

[19] V. H. Bezerra, V. G. T. da Costa, S. Barbon Junior, R. S. Miani, and B. B. Zarpelão, "IoTDS: A One-Class Classification Approach to Detect Botnets in Internet of Things Devices," *Sensors*, vol. 19, no. 14, p. 3188, 2019.

[20] I. Cvitic, D. Perakovic, M. Perisa, and M. Botica, "Definition of the IoT Device Classes Based on Network Traffic Flow Features," in *4th EAI International Conference on Management of Manufacturing Systems*. Krynica Zdroj, Poland: Springer, 2020, pp. 1–17.

[21] A. Singla, E. Bertino, and D. Verma, "Overcoming the lack of labeled data: Training intrusion detection models using transfer learning," in *2019 IEEE International Conference on Smart Computing (SMART-COMP)*. Washington DC, USA: IEEE, 2019, pp. 69–74.

[22] P. Alaei and F. Noorbehbahani, "Incremental anomaly-based intrusion detection system using limited labeled data," in *2017 3rd International Conference on Web Research (ICWR)*. IEEE, 2017, pp. 178–184.

[23] T. L. Saaty, "What is the analytic hierarchy process?" in *Mathematical models for decision support*. Berlin, Heidelberg: Springer, 1988, pp. 109–121.

[24] P. Anand, Y. Singh, A. Selwal, M. Alazab, S. Tanwar, and N. Kumar, "IoT Vulnerability Assessment for Sustainable Computing: Threats, Current Solutions, and Open Challenges," *IEEE Access*, vol. 8, pp. 168 825–168 853, 2020.

[25] Y. Meidan, V. Sachidananda, H. Peng, R. Sagron, Y. Elovici, and A. Shabtai, "A Novel Approach for Detecting Vulnerable IoT Devices Connected Behind a Home NAT," *Computers & Security*, p. 101968, 2020.

[26] Provision, "Provision PT-838E," 2019. [Online]. Available: https://www.provision-isr.com/index.php?option=com_virtuemart&view=productdetails&virtuemart_product_id=516&virtuemart_category_id=41&Itemid=123

[27] 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 *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2017, pp. 2177–2184.

[28] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis

et al., "Understanding the Mirai botnet," in *26th {USENIX} security symposium ({USENIX} Security 17)*, 2017, pp. 1093–1110.

[29] Y. Perwej, K. Haq, F. Parwej, M. Mumdouh, and M. Hassan, "The internet of things (IoT) and its application domains," *International Journal of Computer Applications*, vol. 975, p. 8887, 2019.

[30] C. J. D'Orazio, K.-K. R. Choo, and L. T. Yang, "Data exfiltration from Internet of Things devices: iOS devices as case studies," *IEEE Internet of Things Journal*, vol. 4, no. 2, pp. 524–535, 2016.

[31] K. Irion and N. Helberger, "The Weeping Angels are back, and they attack our privacy via smart TVs," 2017.

[32] I. Ullah, M. A. Shah, A. Wahid, and A. Waheed, "Protection of enterprise resources: A novel security framework," in *2017 International Conference on Communication Technologies (ComTech)*. MCS, Pakistan: IEEE, April 2017, pp. 98–103.

[33] Allot, "IoT Smart Home Security," October 2019 (accessed October 24, 2019). [Online]. Available: https://www.allot.com/products-service-providers/connected-home-security-services/

[34] Avast, "IoT security," October 2019 (accessed October 24, 2019). [Online]. Available: https://www.avast.com/technology/iot-security

[35] T. Varshney, N. Sharma, I. Kaushik, and B. Bhushan, "Architectural Model of Security Threats & their Countermeasures in IoT," in *2019 international conference on computing, communication, and intelligent systems (ICCCIS)*. IEEE, 2019, pp. 424–429.

[36] Q. Jing, A. V. Vasilakos, J. Wan, J. Lu, and D. Qiu, "Security of the Internet of Things: perspectives and challenges," *Wireless Networks*, vol. 20, no. 8, pp. 2481–2501, 2014.

[37] M. Roopak, G. Y. Tian, and J. Chambers, "Multi-objective-based feature selection for DDoS attack detection in IoT networks," *IET Networks*, vol. 9, no. 3, pp. 120–127, 2020.

[38] K. Shaukat, T. M. Alam, I. A. Hameed, W. A. Khan, N. Abbas, and S. Luo, "A review on security challenges in Internet of Things (IoT)," in *2021 26th International Conference on Automation and Computing (ICAC)*. IEEE, 2021, pp. 1–6.

[39] S. Sutar and P. Mekala, "An extensive review on IoT security challenges and LWC implementation on tiny hardware for node level security evaluation," *International Journal of Next-Generation Computing*, vol. 13, no. 1, 2022.

[40] R. R. Papalkar and A. S. Alvi, "Analysis of Defense Techniques for DDoS Attacks in IoT - A Review," *ECS Transactions*, vol. 107, no. 1, p. 3061, 2022.

[41] J. L. Leevy, J. Hancock, T. M. Khoshgoftaar, and J. M. Peterson, "IoT information theft prediction using ensemble feature selection," *Journal of Big Data*, vol. 9, no. 1, pp. 1–48, 2022.

[42] D. Uroz and R. J. Rodríguez, "Characterization and Evaluation of IoT Protocols for Data Exfiltration," *IEEE Internet of Things Journal*, 2022.

[43] P. Jayalaxmi, G. Kumar, R. Saha, M. Conti, T.-h. Kim, and R. Thomas, "DeBot: A deep learning-based model for bot detection in industrial internet-of-things," *Computers and Electrical Engineering*, vol. 102, p. 108214, 2022.

[44] V. Shivraj, M. Rajan, and P. Balamuralidhar, "A graph theory based generic risk assessment framework for internet of things (IoT)," in *2017 IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS)*. IEEE, 2017, pp. 1–6.

[45] Y. Mathov, N. Agmon, A. Shabtai, R. Puzis, N. O. Tippenhauer, and Y. Elovici, "Challenges for Security Assessment of Enterprises in the IoT Era," *arXiv preprint arXiv:1906.10922*, 2019.

[46] N. Agmon, A. Shabtai, and R. Puzis, "Deployment optimization of IoT devices through attack graph analysis," in *Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks*. New York, NY: ACM, 2019, pp. 192–202.

[47] S. Siboni, C. Glezer, A. Shabtai, and Y. Elovici, "A Weighted Risk Score Model for IoT Devices," in *International Conference on Security, Privacy and Anonymity in Computation, Communication and Storage*. Atlanta, USA: Springer, 2019, pp. 20–34.

[48] S. Siboni, C. Glezer, R. Puzis, A. Shabtai, and Y. Elovici, "Security Ranking of IoT Devices Using an AHP Model," in *Cyber Security Cryptography and Machine Learning*, S. Dolev, V. Kolesnikov, S. Lodha, and G. Weiss, Eds. Cham: Springer International Publishing, 2020, pp. 29–44.

[49] A. Shaghaghi, D. D. Nguyen, S. S. Kanhere *et al.*, "Is this IoT Device Likely to be Secure? Risk Score Prediction for IoT Devices Using Gradient Boosting Machines," *arXiv preprint arXiv:2111.11874*, 2021.

[50] N. Apthorpe, D. Y. Huang, D. Reisman, A. Narayanan, and N. Feamster, "Keeping the smart home private with smart(er) iot traffic shaping," *Proceedings on Privacy Enhancing Technologies*, vol. 2019, no. 3, pp. 128–148, 2019.

[51] A. Subahi and G. Theodorakopoulos, "Detecting IoT User Behavior and Sensitive Information in Encrypted IoT-App Traffic," *Sensors*, vol. 19, no. 21, p. 4777, 2019.

[52] Y. Lu and L. Da Xu, "Internet of Things (IoT) cybersecurity research: a review of current research topics," *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 2103–2115, 2018.

[53] M. Hossain, R. Hasan, and A. Skjellum, "Securing the Internet of Things: A Meta-Study of Challenges, Approaches, and Open Problems," in *2017 IEEE 37th International Conference on Distributed Computing Systems Workshops (ICDCSW)*. GA, USA: IEEE, 6 2017, pp. 220–225. [Online]. Available: http://ieeexplore.ieee.org/document/7979820/

[54] M. Nawir, A. Amir, N. Yaakob, and O. B. Lynn, "Internet of Things (IoT): Taxonomy of security attacks," in *2016 3rd International Conference on Electronic Design (ICED)*. Phuket, Thailand: IEEE CAS, Aug 2016, pp. 321–326.

[55] S. A. Al-Qaseemi, H. A. Almulhim, M. F. Almulhim, and S. R. Chaudhry, "IoT architecture challenges and issues: Lack of standardization," in *2016 Future Technologies Conference (FTC)*. SF, USA: SAI, Dec 2016, pp. 731–738.

[56] ISO, "Open Systems Interconnection - Basic Reference Model," Nov 1989. [Online]. Available: https://standards.iso.org/ittf/PubliclyAvailableStandards/s014258_ISO_IEC_7498-4_1989(E).zip

[57] S. Rizvi, A. Kurtz, J. Pfeffer, and M. Rizvi, "Securing the Internet of Things (IoT): A Security Taxonomy for IoT," in *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/ 12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*. NY, USA: IEEE, Aug 2018, pp. 163–168.

[58] Y. Meidan, V. Sachidananda, H. Peng, R. Sagron, Y. Elovici, and A. Shabtai, "IoT-deNAT: Outbound flow-based network traffic data of IoT and non-IoT devices behind a home NAT," Jul. 2020. [Online]. Available: https://doi.org/10.5281/zenodo.3924770

[59] M. Skowron, A. Janicki, and W. Mazurczyk, "Traffic Fingerprinting Attacks on Internet of Things Using Machine Learning," *IEEE Access*, vol. 8, pp. 20 386–20 400, 2020.

[60] A. Sivanathan, D. Sherratt, H. H. Gharakheili, A. Radford, C. Wijenayake, A. Vishwanath, and V. Sivaraman, "Characterizing and classifying IoT traffic in smart cities and campuses," in *2017 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2017, pp. 559–564.

[61] M. R. Shahid, G. Blanc, Z. Zhang, and H. Debar, "IoT Devices Recognition Through Network Traffic Analysis," in *2018 IEEE International Conference on Big Data (Big Data)*, 2018, pp. 5187–5192.

[62] R. A. Bailey, *Design of comparative experiments*. Cambridge University Press, 2008, vol. 25.

[63] D. Semenick, "Tests and measurements: The T-test," *Strength & Conditioning Journal*, vol. 12, no. 1, pp. 36–37, 1990.

[64] 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 the symposium on applied computing*. Marrakesh, Morocco: ACM, 2017, pp. 506–509.

[65] C. Shi, J. Liu, H. Liu, and Y. Chen, "Smart user authentication through actuation of daily activities leveraging wifi-enabled iot," in *Proceedings of the 18th ACM International Symposium on Mobile Ad Hoc Networking and Computing*. Chennai, India: ACM, 2017, p. 5.

[66] Y. Ashibani and Q. H. Mahmoud, "A user authentication model for IoT networks based on app traffic patterns," in *2018 IEEE 9th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON)*. IEEE, 2018, pp. 632–638.

[67] N. Apthorpe, D. Reisman, and N. Feamster, "A smart home is no castle: Privacy vulnerabilities of encrypted IoT traffic," 2017.

[68] H. F. Nweke, Y. W. Teh, M. A. Al-Garadi, and U. R. Alo, "Deep learning algorithms for human activity recognition using mobile and wearable sensor networks: State of the art and research challenges," *Expert Systems with Applications*, vol. 105, pp. 233–261, 2018.

[69] Amcrest, "Amcrest IPM-721." [Online]. Available: https://support.amcrest.com/hc/en-us/categories/202176008-IPM-721

[70] Amazon, "Ring video doorbell." [Online]. Available: https://www.amazon.com/dp/B08N5NQ869?ref=tirvd_rc_nd_ucc

[71] H. E. Hurst, "Long-term storage capacity of reservoirs," *Transactions of the American society of civil engineers*, vol. 116, no. 1, pp. 770–799, 1951.

[72] B. Qian and K. Rasheed, "Hurst exponent and financial market predictability," in *IASTED conference on Financial Engineering and*

Applications. Proceedings of the IASTED International Conference Cambridge, MA, 2004, pp. 203–209.

[73] D. P. Subha, P. K. Joseph, R. Acharya U, C. M. Lim *et al.*, "EEG signal analysis: a survey," *Journal of medical systems*, vol. 34, no. 2, pp. 195–212, 2010.

[74] P. Dymora and M. Mazurek, "Anomaly detection in IoT communication network based on spectral analysis and Hurst exponent," *Applied Sciences*, vol. 9, no. 24, p. 5319, 2019.

[75] D. Mottl, "hurst 0.0.5: Hurst exponent evaluation and R/S-analysis," *PyPI*, 2019.

[76] 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, 2019.

[77] C. Majumdar, M. López-Benítez, and S. N. Merchant, "Real Smart Home Data-Assisted Statistical Traffic Modeling for the Internet of Things," *IEEE Internet of Things Journal*, vol. 7, no. 6, pp. 4761–4776, 2020.

[78] T. Kovanen, G. David, and T. Hämäläinen, "Survey: Intrusion detection systems in encrypted traffic," in *Internet of Things, Smart Spaces, and Next Generation Networks and Systems*. Cham: Springer, 2016, pp. 281–293.

[79] F. Shaikh, N. Ghani, and E. Bou-Harb, "IoT Threat Detection Leveraging Network Statistics and GAN," University of South Florida, Tech. Rep., 09 2019.

[80] O. Brun, Y. Yin, J. Augusto-Gonzalez, M. Ramos, and E. Gelenbe, "IoT attack detection with deep learning," in *ISCIS Security Workshop*. London, UK: Springer, 2018, pp. 1–11.

[81] K.-C. Lin, S.-Y. Chen, and J. C. Hung, "Botnet detection using support vector machines with artificial fish swarm algorithm," *Journal of Applied Mathematics*, vol. 2014, pp. 1–10, 2014.

[82] Cisco, *Capacity and Performance Management: Best Practices*. CA, USA: Cisco, 2005. [Online]. Available: https://www.cisco.com/c/en/us/support/docs/availability/high-availability/20769-performwp.html

[83] O. Hahm, E. Baccelli, H. Petersen, and N. Tsiftes, "Operating Systems for Low-End Devices in the Internet of Things: A Survey," *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 720–734, 2016.

[84] A. Khraisat, I. Gondal, P. Vamplew, J. Kamruzzaman, and A. Alazab, "A novel Ensemble of Hybrid Intrusion Detection System for Detecting Internet of Things Attacks," *Electronics*, vol. 8, no. 11, p. 1210, 2019.

[85] J. Habibi, D. Midi, A. Mudgerikar, and E. Bertino, "Heimdall: Mitigating the Internet of Insecure Things," *IEEE Internet of Things Journal*, vol. 4, no. 4, pp. 968–978, 2017.

[86] T. L. Saaty, "Decision making with the analytic hierarchy process," *International journal of services sciences*, vol. 1, no. 1, pp. 83–98, 2008.

[87] Y.-Y. Shih and C.-Y. Liu, "A method for customer lifetime value ranking - combining the analytic hierarchy process and clustering analysis," *Journal of Database Marketing & Customer Strategy Management*, vol. 11, no. 2, pp. 159–172, 2003.

[88] A. Roohi, M. Adeel, and M. A. Shah, "Ddos in iot: A roadmap towards security countermeasures," in *2019 25th International Conference on Automation and Computing (ICAC)*.

[89] A. Wahab, O. Ahmad, M. Muhammad, and M. Shah, "A Comprehensive Analysis on the Security Threats and their Countermeasures of IoT," *International Journal of Advanced Computer Science and Applications*, vol. 8, 01 2017.

[90] N. Saardchom, "Credit Scoring Model by Analytic Hierarchy Process (AHP)," *Global Review of Accounting and Finance*, vol. 3, no. 2, pp. 58–73, 2012.

[91] C. J. Yean, T. C. Yee, and I. K. Tan, "Relative trust management model for Twitter: An Analytic Hierarchy Process approach," in *International Conference on Frontiers of Communications, Networks and Applications (ICFCNA 2014)*. Malaysia: IET, 2014, pp. 1–6.

[92] U. S. F. C. Commission *et al.*, "Searchable fcc id database," 2020. [Online]. Available: https://fccid.io/

[93] E. Lear, D. Romascanu, and R. Droms, "Manufacturer usage description specification," March 2019.

[94] D.-S. Authors, "Online Questionnaire: Detecting IoT attacks," 2019. [Online]. Available: https://dscore.limequery.com/915153

[95] C. Schmitz, "Limesurvey," Nov 2017. [Online]. Available: https://www.limesurvey.org/

[96] R. Bitton, A. Finkelshtein, L. Sidi, R. Puzis, L. Rokach, and A. Shabtai, "Taxonomy of mobile users' security awareness," *Computers & Security*, vol. 73, pp. 266–293, 2018.

[97] SimpleHome, "Simple Home XCS7-1003," 2016. [Online]. Available: http://gosimplehome.com/images/specs/XCS7-1003-WHT-OutdoorSecurityCam_Manual.pdf

[98] Provision, "Provision PT-737E," 2019. [Online]. Available: https://www.provision-isr.com/index.php?option=com_virtuemart&view=productdetails&virtuemart_product_id=462&virtuemart_category_id=41&Itemid=123

[99] SimpleHome, "Simple Home XCS7-1002," 2016. [Online]. Available: http://gosimplehome.com/images/specs/XCS7-1002-PTCam_Manual.pdf

[100] Philips, "In.Sight wireless HD baby monitor B120N/10." [Online]. Available: https://www.philips.co.uk/c-p/B120N_10/in.sight-wireless-hd-baby-monitor

[101] Danmini, "DANMINI WF - Doorbell 720P WiFi Video Doorbell," 2019. [Online]. Available: https://www.dansarosa.com/danmini-720p-wifi-video-doorbell.html

[102] Ennio, "ENNIO Bell Smart WiFi Video Camera Door Phone Rainproof Intercom Doorbell IR LED," 2019. [Online]. Available: https://www.proyu-tech.com/download/16071516530693.html

[103] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in *2008 eighth ieee international conference on data mining*. IEEE, 2008, pp. 413–422.

[104] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, "LOF: identifying density-based local outliers," in *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, 2000, pp. 93–104.

[105] Z. Noumir, P. Honeine, and C. Richard, "On simple one-class classification methods," in *2012 IEEE International Symposium on Information Theory Proceedings*. IEEE, 2012, pp. 2022–2026.

[106] S. Barbarà and M. Jackson, "Maximin, leximin, and the protective criterion: characterizations and comparisons," *Journal of Economic Theory*, vol. 46, no. 1, pp. 34–44, 1988.

[107] K. Goeschel, "Reducing false positives in intrusion detection systems using data-mining techniques utilizing support vector machines, decision trees, and naive Bayes for off-line analysis," in *SoutheastCon 2016*, 2016, pp. 1–6.

[108] Nguyen Thanh Van, Tran Ngoc Thinh, and Le Thanh Sach, "An anomaly-based network intrusion detection system using deep learning," in *2017 International Conference on System Science and Engineering (ICSSE)*, 2017, pp. 210–214.

[109] S. Fitriani, S. Mandala, and M. A. Murti, "Review of semi-supervised method for intrusion detection system," in *2016 Asia Pacific Conference on Multimedia and Broadcasting (APMediaCast)*, 2016, pp. 36–41.

[110] W. Li, W. Meng, and M. H. Au, "Enhancing collaborative intrusion detection via disagreement-based semi-supervised learning in IoT environments," *Journal of Network and Computer Applications*, p. 102631, 2020.

[111] M. H. Bhuyan, D. K. Bhattacharyya, and J. K. Kalita, *Network Traffic Anomaly Detection Techniques and Systems*. Cham: Springer International Publishing, 2017, pp. 115–169. [Online]. Available: https://doi.org/10.1007/978-3-319-65188-0_4

[112] M. Zhuang, G. Cui, and L. Peng, "Manufactured opinions: The effect of manipulating online product reviews," *Journal of Business Research*, vol. 87, pp. 24–35, 2018.

[113] FCC, "The federal communications commission," 2020, https://www.https://www.fcc.gov/, Last accessed on 2020-08-14.

[114] J. Surowiecki, *The wisdom of crowds*. New York, NY: Anchor, 2005.

Yair Meidan is a PhD candidate in the Department of Software and Information Systems Engineering (SISE) at Ben-Gurion University of the Negev (BGU). His research interests include machine learning and IoT security. Contact him at yairme@post.bgu.ac.il.

Daniel Benatar is an MSc student in the SISE Department at BGU. His research interests include machine learning and cyber-security. Contact him at benatar@post.bgu.ac.il.

Ron Bitton is a PhD candidate in the SISE Department at BGU. His research interests include machine learning, cyber-security, and cyber-risk management. Contact him at ronbit@post.bgu.ac.il.

Dan Avraham is an MSc student in the SISE Department at BGU. His research interests include machine learning and cyber-security. Contact him at danavra@post.bgu.ac.il.

Asaf Shabtai is a professor in the SISE Department at BGU. His research interests include computer and network security, and machine learning. Shabtai received a PhD in information systems from BGU. Contact him at shabtaia@bgu.ac.il.

APPENDIX A
ANALYSIS OF TRAFFIC PREDICTABILITY USING AN ADDITIONAL DATASET (UNSW) [76]
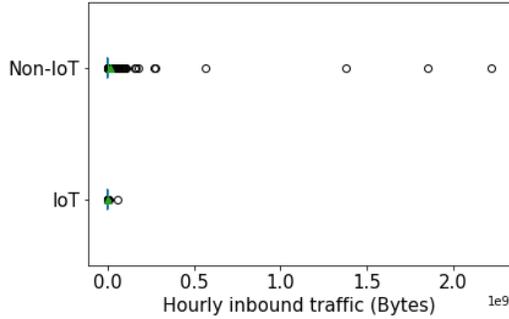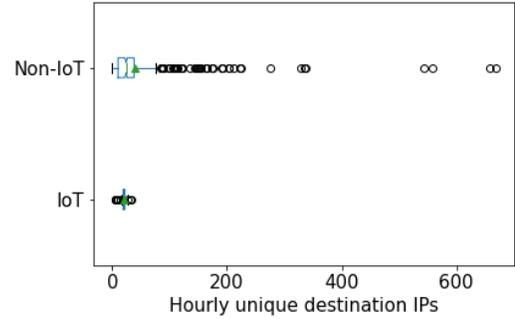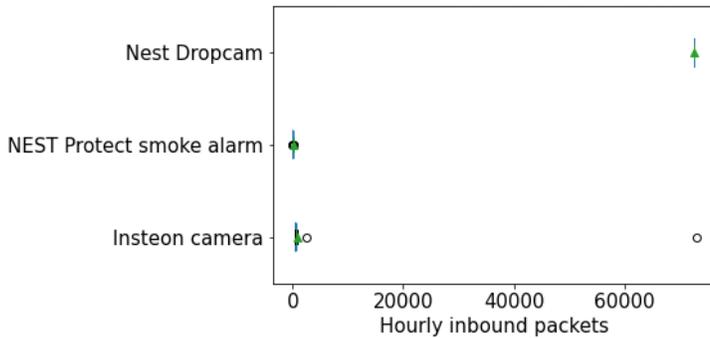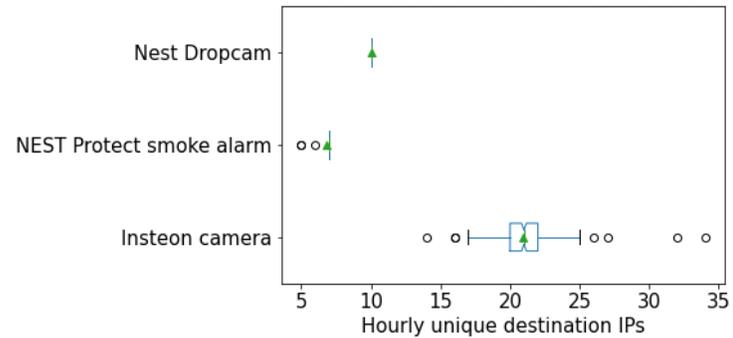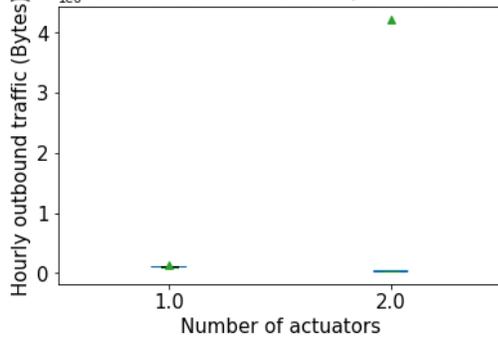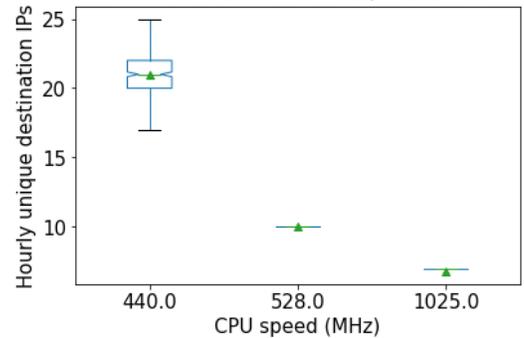
Fig. 13: Traffic predictability analysis using UNSW data [76]: Comparison between IoT and non-IoT devices (Fig. 13a and 13b), comparison among IoT models (Fig. 13c and 13d), and correlation with IoT model complexity (Fig. 13e and 13f). In this dataset, the IoT devices comprised of Nest Dropcam, Nest Protect smoke alarm, and Insteon camera, while the non-IoT devices comprised of Samsung Galaxy Tab, iPhone, a laptop, two Macbooks, and two Android phones.