WiP: The Intrinsic Dimensionality of IoT Networks

Matt Gorbett
Colorado State University
Fort Collins, Colorado, USA
matt.gorbett@colostate.edu

Hossein Shirazi
Colorado State University
Fort Collins, Colorado, USA
hossein.shirazi@colostate.edu

Indrakshi Ray
Colorado State University
Fort Collins, Colorado, USA
indrakshi.ray@colostate.edu

ABSTRACT

The Internet of Things (IoT) is revolutionizing society by connecting people and devices seamlessly and providing enhanced user experience and functionalities. However, the unique properties of IoT networks, such as heterogeneity and non-standardized protocol, have created critical security holes and network mismanagement. We propose a new measurement tool for IoT network data to aid in analyzing and classifying such network traffic. We use evidence from both security and machine learning research, which suggests that the *complexity* of a dataset can be used as a metric to determine the trustworthiness of data. We test the complexity of IoT networks using Intrinsic Dimensionality (ID), a theoretical complexity measurement based on the observation that a few variables can often describe high dimensional datasets. We use ID to evaluate four modern IoT network datasets empirically, showing that, for network and device-level data generated using IoT methodologies, the ID of the data fits into a low dimensional representation; this makes such data amenable to the use of machine learning algorithms for anomaly detection.

CCS CONCEPTS

• **Security and privacy → Intrusion detection systems**; • **Computing methodologies → Anomaly detection**; **Dimensionality reduction and manifold learning**.

KEYWORDS

Intrusion Detection, IoT, Internet of Things, Intrinsic Dimensionality, Data Complexity

ACM Reference Format:
Matt Gorbett, Hossein Shirazi, and Indrakshi Ray. 2022. WiP: The Intrinsic Dimensionality of IoT Networks. In *Proceedings of the 27th ACM Symposium on Access Control Models and Technologies (SACMAT) (SACMAT '22), June 8–10, 2022, New York, NY, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3532105.3535038

1 INTRODUCTION

With people, objects, sensors, and services all connected through devices ranging from household appliances to smartphones and PCs, the Internet of Things (IoT) network infrastructure faces the challenging task of managing heterogeneous devices and their communications in the absence of standardization. The proliferation of IoT systems has introduced new, and emerging security vulnerabilities [2, 5, 10, 38] which can be readily exploited to cause harm. Such vulnerabilities arise because of device manufacturers neglecting security for performance considerations [11], end-users not updating each device regularly [37], and a continually expanding marketplace of devices and manufacturers [10].

Further, the unique characteristics of IoT networks have introduced new complications. Notably, heterogeneity and non standardized protocol of IoT networks have been posited as critical challenges for enhancing the security of IoT systems [8, 18, 31] – networks with diverse devices ranging from single-purpose machines to robust servers, each with varied communication structures, are cumbersome to protect. Past work has proposed behavioral fingerprinting of devices [7], and further fine-tuning device-specific anomaly detection models depending on the complexity of devices [15]. Others propose supervised machine learning solutions [21, 33], utilizing modern network datasets such as Aposemat IoT-23 (IoT-23) [13].

In this work, we take a different approach by first examining the supposition that heterogeneous IoT networks have higher complexity than regular non-IoT datasets. We calculate Intrinsic Dimensionality (ID), a property that has been proposed to measure the complexity of a data set as a whole [39] and evaluate it on four IoT datasets and two non-IoT datasets. We analyze the datasets from two perspectives: network level and device level, and show that, despite the variability of IoT devices, the complexity of benign network activity is low.

1.1 Problem Statement

We focus on the question of heterogeneity and complexity in IoT networks by asking the following questions:

(1) Do the properties of multi-device heterogeneous IoT networks exhibit fundamentally more complex behavior?
(2) What devices are harder to protect in machine learning-based frameworks?

1.2 Proposed Approach

In this work, we measure the complexity of IoT network traffic using the novel perspective of ID, testing the hypothesis that IoT networks contain complex network packets as a result of their heterogeneous behavior. We first measure ID at the network dataset level, showing that, counter to intuition, several IoT datasets exhibit lower ID compared to non-IoT benchmarks. Additionally, we show that the ID measurement used in our experiments exhibit similar *rank order* complexity as Complex IoT [15], *i.e.*, the complexity measurements of devices are arranged in a similar order in both works.

SACMAT '22, June 8–10, 2022, New York, NY, USA

ACM ISBN 978-1-4503-9357-7/22/06…$15.00
https://doi.org/10.1145/3532105.3535038

1.3 Contributions

Our contributions are as follows:

- We measure the complexity of several popular IoT, non-IoT, and IoT device datasets using ID, testing the hypothesis that IoT networks contain complex interactions. Despite being heterogeneous in nature, we show that IoT network activity has low ID measurements, with ID values similar to device-level traffic. *Low ID measurements provide strong evidence that we can build robust and secure Machine Learning (ML) models to protect IoT networks.* To the best of our knowledge, this is the first work that analyzes IoT datasets using ID.
- We show the usefulness of ID as a *device* complexity measurement, displaying empirical results across 17 devices and the intuition behind them.
- We measure the complexity of IoT traffic at the network level and show that even for networks with several devices ranging from single-purpose to multi-use servers, the data generally fits into low ID estimates.

The rest of the paper is organized as follows. We first summarize related and prior work in Section 2. We then detail ID concepts in Section 3 as well as discuss our methodology. We summarize the datasets we used in experiments in Section 4. Next is our analysis of device level ID estimates in Section 6. Finally, we conclude this paper with a discussion and pointers to future directions in Section 7.

2 RELATED WORK

In this section, we review Intrusion Detection Systems (IDS), IoT security research, current algorithms for IoT intrusion detection, and finally, open problems in deep learning anomaly detection.

IoT: IoT is a rapidly evolving field, with research being done at dozens of institutions across industry and academia [8, 18, 20, 31]. It is postulated that IoT increases the vulnerability of networks because the *attack surface* has increased, with many new entry and exit points with new devices available on networks [24, 32]. A heterogeneous IoT network is typically made of various sub-devices within a distributed network. It includes resource-constrained devices, such as a smart light bulb or garage door opener, and more powerful devices such as embedded and regular computers.

Existing research notes that IoT networks and devices have multiple intrusion sources: IoT backends, cloud services supporting an IoT device, and other hubs within the IoT system [1, 29], which makes it difficult to implement traditional intrusion detection approaches such as rule-based and signature-based methods.

Complexity in Deep Learning and Security: Several works in computer vision have shown that various classes of deep neural networks are susceptible to anomalous data [14, 16, 25, 28, 35]. Recently, Serra et al. showed that this was a result of data complexity [35]. They use image compression measurements to show that models developed with more complex data are susceptible to anomalous and out-of-distribution examples. Further, they show that simpler datasets are more robust to anomalies.

Pope et al. [31] showed that common computer vision datasets exhibit very low intrinsic dimension relative to their number of pixels. They also showed that the intrinsic dimension greatly impacts learning: the higher the intrinsic dimension of a dataset, the harder it is to learn from it. In addition, they showed that the extrinsic dimension of the dataset, i.e.the total number of pixels per image in a dataset, did not effect learning and generalization, indicating that sample complexity only depends on the intrinsic dimension rather than the total dimension of the dataset.

Interestingly, a similar complexity finding was found in a recent security paper Haefner and Ray [15]. Using data from various IoT devices, they find that each device has varying complexity. They formalize a complexity measure (IP Spread/IP Depth) per device in order to fine-tune an Isolation Forest anomaly detection algorithm. Their architecture, ComplexIoT, measures network traffic on a device level, which can be used in Host Intrusion Detection Systems.

This work is similar to ComplexIoT [15] in that we propose a complexity measurement; however, there are several key differences:

- We analyze IoT datasets both from the point of view of network-level and the device level, while ComplexIoT only looks at device level complexity.
- The ComplexIoT complexity score is based on IP spread and IP depth and does not consider other network features to compute its complexity estimate.

3 METHODOLOGY

In this section we first briefly explain concepts and mathematics behind ID. We will later use ID to measure both network and device level IoT and non-IoT datasets, showing how this complexity measurement is a strong tool of our ability to assess network data at multiple levels.

3.1 Intrinsic Dimensionality

The ID of a dataset is the minimum number of variables needed to retain a full approximation of the data [6]. It is based on the observation that high-dimensional data can often be described by a lower number of variables. The utility of lower dimensional representations is apparent throughout ML research, from data compression (such as autoencoders [17]) to dimensionality reduction (PCA). ID is akin to autoencoders and PCA, however quite distinct in that its an estimate of the lowest possible dimension of a dataset (e.g. the lowest possible bottleneck size in an autoencoder), and not a reduction technique in itself. ID can be thought of as a geometric property to measure complexity of a dataset as a whole [39].

Formally, the ID of dataset $\mathcal{X} \in \mathbb{R}^{mxn}$, with m samples and n features, lies on a lower dimensional manifold $\mathcal{M}$, where $ID = dim(\mathcal{M})$, i.e. ID is the dimension of the manifold $\mathcal{M}$ of the data. Usually, the ID measurement is significantly less than extrinsic dimension n, or number of features.

As an intuitive example, points $x_1 ... x_m$ exist on a piece of paper in three dimensional space. We can describe the points relative to the three dimensional space, (d_1, d_2, d_3), or we can describe them relative to their position on the piece of paper, where only two variables are needed. Here, the representation of points $x_1 ... x_m$ in 3D space is the *extrinsic* dimension, whereas their points relative to the piece of paper are their ID.

The main approach to estimate ID involves examining the neighborhood around a reference point x_i for each x in $\mathcal{X}$. A common

equation used in existing research was proposed by Levina and Bickel [22]:

$$ID(X) = \left(\frac{1}{m(k-1)} \sum_{i=1}^{m} \sum_{j=1}^{k-1} log \frac{T_k(x_i)}{T_j(x_i)} \right)^{-1} \quad (1)$$

where m is the number of samples, x_i is a sample in the dataset, k and j are the k^{th} and j^{th} nearest neighbors. $T_k(x_i)$ is the distance between x_i and x_k, similarly, $T_j(x_i)$ is the distance between x_i and x_j. Intuitively, Equation 1 measures the rate that new neighbors are encountered as we move out from the reference point x_i. We use this equation for all ID estimates in Figure 1.

Outside of deep learning (as described in Section 2), intrinsic dimensionality has been used in applications such as anomaly detection[36], clustering, similarity search, and deformation in complex materials.

Distance Metric. For the distance metric required in Equation 1, we use Hamming Distance to compute similarity between both categorical and continuous feature points. While Euclidean Distance is typically used in Equation 1 to measure LID, Ma et al. [23] suggested not using Euclidean Distance as the underlying distance metric. Choosing the Hamming Distance metric over Euclidean Distance for continuous variables showed better experimental results. Effectively, this turns each pairwise feature distance into a binary metric: 0 for same, 1 for different. We compute Hamming Distance as:

$$\mathcal{H}(x_i, x_j) = \frac{\text{Number of mismatching features}}{\text{Total Features}} \quad (2)$$

Entropy. We calculate entropy of each feature and set it as the weight. In a dataset with n features, we set weight w_i for feature i to $n/Entropy(i)$, where the entropy of a feature i is:

$$-\sum_{j=1}^{m} p_j log_2(p_j) \quad (3)$$

and n is the total number of features and j is the number of classes in feature i. p_j is number of occurrences of class j in feature i. For example, the protocol feature may have TCP and UDP classes, we compute the counts for each to calculate entropy. We find that features with low entropy should be weighted more since they are stable properties of benign samples. For example, if benign samples come from TCP protocol 99% of the time, we can theorize that new samples matching the TCP protocol may be similar to a benign one.

4 DATASETS

This section summarizes the datasets we use in our experiments. We use two common non-IoT network intrusion datasets, UNSW-NB15 and KDD Cup, and four IoT related datasets (*TON_IoT*, NetFlow Bot-IoT (NF Bot-IoT), IoT-23, IoT Sense).

4.1 Non-IoT: UNSW-NB15

UNSW-NB15 [27] is a standard and commonly used network intrusion dataset from the USNW at Canberra Cyber Range lab. The dataset provides modern network traffic scenarios compared to the KDD datasets, which are more than a decade old. There are 47 features (of which we use 42), ranging from basic features to content and time-related features. Nine types of attacks are included in the dataset.

4.2 Non-IoT: KDD Cup 1999

We use a variation of the KDD Cup 1999 dataset [12] located on Kaggle. The dataset consists of 13,449 benign instances and 41 features, which we use to measure ID.

4.3 TON IoT

TON_IoT (*TON_IoT*) [3, 26], published in 2020, comprises heterogeneous IoT data across several devices. The work uses several data source types, including sensor, raw, and log data. Additionally, it includes several infrastructure layers in the testbed architecture, such as the edge, fog, and cloud layers with nine types of generated attacks: Distributed Denial-of-Service (DDoS), Scanning, Ransomware, Backdoor, and Injection attacks. The dataset has 41 total features; however, the authors recommend not to use source IP/port and destination IP/port. The dataset simulates traffic from seven IoT sensors: weather, smart garage door, smart fridge, smart TCP/IP Modbus, GPS tracker, motion-enabled light, and a smart thermostat. For measuring ID, we deduplicated data instances, and as a result, 61.8% of instances have been removed.

4.4 NF Bot-IoT

NF Bot-IoT [34], published in 2020, is a dataset based on the Bot-Net IoT dataset [19, 20]. Botnets are an important attack vector to protect against as they have been the source of several breaches over the past few years [19]. NF Bot-IoT converts four common network Network Intrusion Detection Systems (NIDS) datasets into network flow datasets using the commonly deployed NetFlow [9] protocol for network traffic collection. Authors argue NetFlow's features are easier to extract compared to the complex features used in the original NIDS datasets since NetFlow's features are usually extracted from packet headers. The dataset includes several attacks, including Denial-of-Service (DoS), Distributed Denial-of-Service (DDoS), Service Scan, Keylogging, and Data exfiltration attacks.

4.5 IoT-Sense

IoT Sense [7], published in 2018, is a dataset of benign examples generated based on 14 real IoT devices. Authors activated different functionalities of each device using controller apps and captured packets. There are 21 features captured in the dataset, with labeled devices for each sample. We use this dataset for both ID measurements as well as device-specific ID measurements in Sections 5 and 6. We categorize devices in this dataset into three categories, namely, Light, Appliance, and Hub/Outlet.

Devices include TCP Light, Avox Light, Musiac Music Player, DLink Camera, iDeviceSocket, iView Light, LutronHub, Netamo Climate, Omna camera, Phillips HUE, Tplink Light, Wemo Outlet, Wink Hub, and Smart Things Hub. We use this dataset for both ID measurements as well as device-specific ID measurements in Sections 5 and 6. Lights (TCP light, iView Light, AWOX Light, Phillips Hue Light, TP-Link Light), Appliances (Musiac Music player, D-Link Camera, Omna Camera, Netamo Climate), Hub/Outlet (iDevice Socket, WEMO Outlet, Lutron Hub, Wink Hub, and Smart Things Hub).

Table 1: Dataset summaries including total number of samples, percentage of benign samples, percentage of malicious samples, percent duplicates, number of features, and number of attacks.

Name	#To.	%Ben.	%Mal.	%Dup.	#Fea.	#Att
UNSW-NB15	82K	45%	55%	12.5%	42	-
KDD Cup	24K	45.8%	54.2%	0%	41	-
TON_IoT	461K	65%	35%	62%	38	9
IoT-23	1M	50%	50%	2%	19	7
IoT Sense	54K	100%	0	63%	21	-
NF Bot-IoT	599K	21.7%	78.3	0%	12	4

4.6 IoT-23

IoT-23 [13] was released in 2020. The dataset has 23 different scenarios, of which three are benign traffic scenarios captured on real IoT devices. The dataset contains almost 11 million total records; however, with the difficulty of modeling this much data, we sample a million records with the following logic: From the entire dataset, we sample 500K malicious records and 500K benign examples from simulated files that contain a source IP or destination IP with an internal IP address and have at least 50 samples belonging to that specific IP address. We find that 99.99% of internal IPs have at least 50 samples. We also included all samples from three real devices (Philips HUE smart LED lamp, Amazon Echo, and a Somfy smart door lock) with 1,634 total packets.

We categorize these devices similar to IoT-Sense as a light (Philips HUE smart LED lamp), Smart Controller (Amazon Echo), and Appliance (Somfy smart door lock). Philips HUE light is in both datasets so that it can be used for comparison device-specific ID measurements.

This dataset also captures 20 simulated scenarios of both benign and malicious traffic. It offers several attack examples: DDoS, FileDownload (to infected device), HeartBeat (indicates packets sent on the connection are used to keep track of infected host by CC server), Mirai, Torii, and Okiru BotNets (new common attacks), and HorizontalPortScan (used to gather information for future attacks).

4.7 IoT Dataset Features

We use the available features for each dataset, except we exclude source and destination IP and port as well as any ID or timestamp columns for *TON_IoT* and IoT-23. We include IPs and ports in NF Bot-IoT because of its small number of available features to provide more discriminability. Features common among the datasets include protocol, source, and destination bytes, connection state, service, duration, missed bytes, number of packets, window size, payload, entropy, DNS, SSL, and HTTP properties.

5 INTRINSIC DIMENSIONALITY OF NETWORK DATASETS

Our first ID analysis is measuring benign networks at the full dataset level, using the approach explained in Section 3.1. We measure the ID over several K estimates (K=3, K=5, K=10, and K=20), where K is the number of neighbors to use to measure ID. Results of this experiment are depicted in Figure 1.

IoT datasets vs. Non-IoT Datasets. First, we look at the benign subsets of four IoT datasets (*TON_IoT*, IoT-23, NF Bot-IoT, IoT-Sense). Each dataset contains an ID estimate under 2. Comparatively, the non-IoT network datasets of UNSW-NB15 and KDD Cup '99 have ID estimates between 3.61 and 7.1, substanially higher than the IoT network data. The relative simplicity of IoT network data indicates it will easier to estimate its behavior, leading to better Network Intrusion Detection (NID) models and more robust detection of attacks.

Effects of K Value. One other observation is related to K values. Several works note that the ID estimate is sensitive to K [4, 30], so estimating ID values over several K's gives us a robust picture. As Figure 1 shows, rank order of each dataset does not change substantially given the choice of K. Hamming Distance is used for all distance computations.

Effects of Feature Size Another notable finding across all datasets is that the extrinsic dimensionality, or number of features in the datasets, does not appear to be correlated with its intrinsic dimensionality. For example, *TON_IoT* contains almost as many features as UNSW-NB15 and KDD Cup 1999, however, its ID estimates are substantially lower than either. This indicates that the features of IoT network data are more simplistic in nature than non-IoT datasets.

IoT vs. Alternative Datasets Finally, we compare these values to more difficult modeling on common computer vision datasets. Pope et al. [30] show that MNIST has the lowest ID, estimated between 7 and 13, with the state-of-the-art accuracy of 99.84%. In contrast, ImageNet has an ID between 26 and 43 with a state-of-the-art accuracy of 88.5%, indicating that datasets with a higher ID may be difficult to model. This is further examined in [30]. Relating these values to ID estimates on network data, we see that UNSW-NB15 has a similar ID to MNIST. While this indicates that UNSW-NB15 can still be modeled with very high accuracy, its ID is substantially higher than IoT network datasets, indicating IoT networks may be easier to model.

6 INTRINSIC DIMENSIONALITY OF IOT DEVICES

In this section, we analyze benign IoT traffic for specific devices. The purpose of this analysis is to compare device-specific complexity via ID in order to (i) verify the behavior of various devices as described in [15] and (ii) reason about the ID values across different IoT devices.

Haefner and Ray [15] defined a spectrum for the complexity for IoT devices, starting with simple devices such as single-purpose machines with low variability in their network interactions to complex devices (like Amazon Echo) with high variability in their network interactions. Similarly, and to simplify understanding of our results, we split ID measures into three categories: Low (0 to 0.5), Medium (0.5 to 0.7), and High ID (more than 0.7+).

In our datasets, 6 of these devices (Omna Camera, Smart Things huB, Netmao Climate, Lutron Hub, Wemo Outlet, and iDevice Socket) are labeled as low complexity devices. Three devices of iView Light, TP-Link Light, and D-Link Camera are labeled as medium complexity. Seven devices are classified with high complexity measures (Wink Hub, Philips HUE light, Door Lock, Musiac

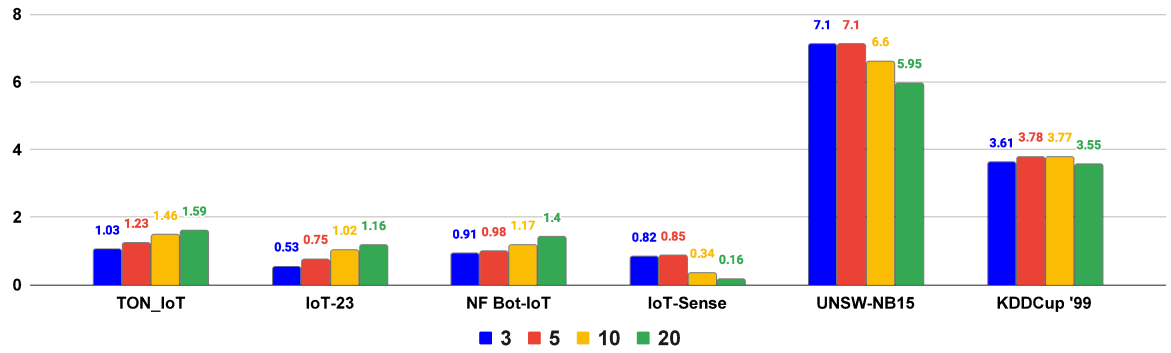

Figure 1: ID measurement for six different datasets. The four left datasets are IoT datasets and the two right datasets are non-IoT datasets. This figure also shows different K values for ID using different colors.

Table 2: Results of ID experiments for 17 different devices on IoT-23 and IoT-Sense.

Category	Device	Dataset	ID
Light	TCP light	IoTSense	0.787
Light	iView Light	IoTSense	0.543
Light	AWOX Light	IoTSense	0.845
Light	Phillips Hue	IoTSense	0.796
Light	TP-Link Light	IoTSense	0.55
Light	Phillips Hue	IoT-23	0.73
Appliance	Musiac Music player	IoTSense	0.866
Appliance	D-Link Camera	IoTSense	0.595
Appliance	Omna Camera	IoTSense	0.323
Appliance	Netamo Climate	IoTSense	0.439
Appliance	Somfy Door Lock	IoT-23	0.81
Hub/Outlet	iDevice Socket	IoTSense	0.474
Hub/Outlet	WEMO Outlet	IoTSense	0.483
Hub/Outlet	Lutron Hub	IoTSense	0.447
Hub/Outlet	Wink Hub	IoTSense	0.881
Hub/Outlet	Smart Things Hub	IoTSense	0.353
Smart Controller	Amazon Echo	IoT-23	1.14

Music, AWOX light, and Amazon Echo). Multiple observations can be made here.

First, Amazon Echo has the highest ID value among all devices, which intuitively makes sense. In addition, while Haefner and Ray [15] do not measure Amazon Echo directly, they measured a similar device in Alexa. Their results confirmed that Amazon Alexa had a high complexity measure, which matches our findings. Further, the rank order of ComplexIoT is similar to ours: both sets had TP-Link Light, Philips HUE, Smart Hub, and Alexa/Echo devices. ComplexIoT measured TP-Link as the lowest complexity, followed by the HUE, Smart Things Hub, and Alexa. Our measurements indicated that the Smart Things Hub had the lowest complexity, followed by the TP-Link Light, HUE, and Echo. The only inconsistency in the ranks was the Smart Things Hub, where Complex IoT measured a higher complexity value. Otherwise, the order of complexities of each device was the same. The Smart Things Hub could have yielded different measurements between the two results because of varying network captures between the two datasets.

All devices in the low-complexity range are aligned with our understanding of the simple-functionalities of these devices, except the Omna Camera. Like a camera with a high volume of sending and receiving data, we expect it will fall in the medium complexity range, similar to the D-Link Camera. One reason could be that the camera uses a specific protocol for sending images/videos (like UDP) that are not captured in the dataset we had, and only command packets have been captured. Consequently, the ID model does not give it a high value.

In our experiment, we had six lights which fell into medium and high complexity, a consistent result among them. We had two instances of the same device (Philips Hue light) in two datasets with different sets of features. The results of these two devices are very close to each other, 0.796 and 0.73 for IoT-Sense and IoT-23 datasets, respectively. This shows that regardless of the different features in the two datasets, our approach consistently evaluated the same device, which is promising.

7 CONCLUSION AND FUTURE WORK

In this work, we view several network datasets through the lens of complexity and show that IoT datasets exhibit a lower ID complexity estimate than standard network collections. Our complexity analysis provides a novel mathematical look into the details surrounding IoT network datasets, showing the relative simplicity of network collections through ID estimates. Additionally, we make connections between complexity in IoT security and open problems in deep learning, such as the difficulty in modeling increasingly complex data such as large images. Connecting the dots between security and anomaly detection in machine learning remains an essential facet of developing secure systems, and we hope this paper can provide researchers with a unique perspective towards building more robust and secure frameworks.

ACKNOWLEDGEMENT

This work was supported in part by funding from NSF under Award Number CNS 1822118, NIST, ARL, Statnett, AMI, Cyber Risk Research, NewPush, and State of Colorado Cybersecurity Center.

REFERENCES

[1] Rasheed Ahmad and Izzat Alsmadi. 2021. Machine learning approaches to IoT security: A systematic literature review. *Internet of Things* 14 (2021), 100365. https://doi.org/10.1016/j.iot.2021.100365

[2] F. A. Alaba, M. Othman, I. A. T. Hashem, and F. Alotaibi. 2017. Internet of Things security: A survey. *Journal of Network and Computer Applications* 88 (2017), 10–28. https://doi.org/10.1016/j.jnca.2017.04.002

[3] A. Alsaedi, N. Moustafa, Z. Tari, A. Mahmood, and A. Anwar. 2020. TON_IoT Telemetry Dataset: A New Generation Dataset of IoT and IIoT for Data-Driven Intrusion Detection Systems. *IEEE Access* 8 (2020), 165130–165150. https://doi.org/10.1109/ACCESS.2020.3022862

[4] L. Amsaleg, O. Chelly, T. Furon, S. Girard, M. E. Houle, K. Kawarabayashi, and M. Nett. 2015. Estimating Local Intrinsic Dimensionality. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '15)*. Association for Computing Machinery, 29–38. https://doi.org/10.1145/2783258.2783405

[5] I. Andrea, C. Chrysostomou, and G. Hadjichristofi. 2015. Internet of Things: Security vulnerabilities and challenges. In *IEEE Symposium on Computers and Communication (ISCC)*. 180–187. https://doi.org/10.1109/ISCC.2015.7405513

[6] D. Bernal. 2014. 3 - Analytical techniques for damage detection and localization for assessing and monitoring civil infrastructures. In *Sensor Technologies for Civil Infrastructures*, M. L. Wang, J. P. Lynch, and H. Sohn (Eds.). Woodhead Publishing Series in Electronic and Optical Materials, Vol. 56. Woodhead Publishing, 67–92. https://doi.org/10.1533/9781782422433.1.67

[7] B. Bezawada, M. Bachani, J. Peterson, H. Shirazi, I. Ray, and I. Ray. 2018. IoTSense: Behavioral Fingerprinting of IoT Devices. *arXiv:1804.03852 [cs]* (2018). http://arxiv.org/abs/1804.03852

[8] D. Choudhary. 2019. Security Challenges and Countermeasures for the Heterogeneity of IoT Applications. *Journal of Autonomous Intelligence* 1 (2019), 16. https://doi.org/10.32629/jai.v1i2.25

[9] Benoît Claise. 2004. *Cisco Systems NetFlow Services Export Version 9.* Request for Comments RFC 3954. Internet Engineering Task Force. https://doi.org/10.17487/RFC3954

[10] M. Conti, A. Dehghantanha, K. Franke, and S. Watson. 2018. Internet of Things Security and Forensics: Challenges and Opportunities. *Future Generation Computer Systems* 78 (2018), 544–546. https://doi.org/10.1016/j.future.2017.07.060 arXiv: 1807.10438.

[11] E. M. Dovom, A. Azmoodeh, A. Dehghantanha, D. E. Newton, R. M. Parizi, and H. Karimipour. 2019. Fuzzy pattern tree for edge malware detection and categorization in IoT. *Journal of Systems Architecture* 97 (2019), 1–7. https://doi.org/10.1016/j.sysarc.2019.01.017

[12] D. Dua and C. Graff. [n. d.]. UCI Machine Learning Repository. http://archive.ics.uci.edu/ml

[13] Sebastian Garcia, Agustin Parmisano, and Maria Jose Erquiaga. [n. d.]. IoT-23 Dataset: A labeled dataset of Malware and Benign IoT Traffic.(Version 1.0.0) [Data set]. Zenodo. https://www.stratosphereips.org/datasets-iot23

[14] M. Gorbett and N. Blanchard. 2020. Utilizing Network Properties to Detect Erroneous Inputs. *arXiv:2002.12520 [cs]* (2020). http://arxiv.org/abs/2002.12520

[15] K. Haefner and I. Ray. 2019. ComplexIoT: Behavior-Based Trust For IoT Networks. In *2019 First IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA)*. 56–65. https://doi.org/10.1109/TPS-ISA48467.2019.00016

[16] D. Hendrycks and K. Gimpel. 2018. A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks. *arXiv:1610.02136 [cs]* (2018). http://arxiv.org/abs/1610.02136

[17] G. E. Hinton and R. R. Salakhutdinov. 2006. Reducing the Dimensionality of Data with Neural Networks. *Science* 313, 5786 (July 2006), 504–507. https://doi.org/10.1126/science.1127647 Publisher: American Association for the Advancement of Science.

[18] R. Kollolu. 2020. *A Review on Wide Variety and Heterogeneity of IoT Platforms.* SSRN Scholarly Paper ID 3912454. Social Science Research Network, Rochester, NY. https://doi.org/10.2139/ssrn.3912454

[19] N. Koroniotis, N. Moustafa, E. Sitnikova, and J. Slay. 2018. Towards Developing Network Forensic Mechanism for Botnet Activities in the IoT Based on Machine Learning Techniques. In *Mobile Networks and Management (Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering)*. Springer International Publishing, Cham, 30–44. https://doi.org/10.1007/978-3-319-90775-8_3

[20] N. Koroniotis, N. Moustafa, E. Sitnikova, and B. Turnbull. 2019. Towards the development of realistic botnet dataset in the Internet of Things for network forensic analytics: Bot-IoT dataset. *Future Generation Computer Systems* 100 (2019), 779–796. https://doi.org/10.1016/j.future.2019.05.041

[21] R. Kozik, M. Pawlicki, and M. Choraś. 2021. A new method of hybrid time window embedding with transformer-based traffic data classification in IoT-networked environment. *Pattern Analysis and Applications* 24, 4 (2021), 1441–1449. https://doi.org/10.1007/s10044-021-00980-2 Number: 4.

[22] E. Levina and P. J. Bickel. 2004. Maximum Likelihood estimation of intrinsic dimension. In *Proceedings of the 17th International Conference on Neural Information Processing Systems (NIPS'04)*. MIT Press, Cambridge, MA, USA, 777–784.

[23] X. Ma, B. Li, Y. Wang, S. M. Erfani, S. Wijewickrema, G. Schoenebeck, D. Song, M. E. Houle, and J. Bailey. 2018. Characterizing Adversarial Subspaces Using Local Intrinsic Dimensionality. *arXiv:1801.02613 [cs]* (2018). http://arxiv.org/abs/1801.02613

[24] Mujahid Mohsin, Zahid Anwar, Ghaith Husari, Ehab Al-Shaer, and Mohammad Ashiqur Rahman. 2016. IoTSAT: A formal framework for security analysis of the internet of things (IoT). In *2016 IEEE Conference on Communications and Network Security (CNS)*. 180–188. https://doi.org/10.1109/CNS.2016.7860484

[25] W. Morningstar, C. Ham, A. Gallagher, B. Lakshminarayanan, A. Alemi, and J. Dillon. 2021. Density of States Estimation for Out of Distribution Detection. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*. PMLR, 3232–3240. https://proceedings.mlr.press/v130/morningstar21a.html ISSN: 2640-3498.

[26] N. Moustafa. 2021. A new distributed architecture for evaluating AI-based security systems at the edge: Network TON_IoT datasets. *Sustainable Cities and Society* 72 (2021), 102994. https://doi.org/10.1016/j.scs.2021.102994

[27] N. Moustafa and J. Slay. 2015. UNSW-NB15: a comprehensive data set for network intrusion detection systems (UNSW-NB15 network data set). In *2015 Military Communications and Information Systems Conference (MilCIS)*. 1–6. https://doi.org/10.1109/MilCIS.2015.7348942

[28] E. Nalisnick, A. Matsukawa, Y. W. Teh, D. Gorur, and B. Lakshminarayanan. 2018. Do Deep Generative Models Know What They Don't Know? https://openreview.net/forum?id=H1xwNhCcYm

[29] 18th October 2021. [n. d.]. Tapping AI for Intrusion Detection Systems. https://www.iotworldtoday.com/2021/10/18/tapping-ai-for-intrusion-detection-systems/

[30] P. Pope, C. Zhu, A. Abdelkader, M. Goldblum, and T. Goldstein. 2020. The Intrinsic Dimension of Images and Its Impact on Learning. https://openreview.net/forum?id=XJk19XzGq2J

[31] B. M. Rashma, S. Macherla, A. Jaiswal, and G. Poornima. 2021. Handling Heterogeneity in an IoT Infrastructure. In *Advances in Machine Learning and Computational Intelligence (Algorithms for Intelligent Systems)*. Springer, Singapore, 635–643. https://doi.org/10.1007/978-981-15-5243-4_60

[32] S. Rizvi, R. J. Orr, A. Cox, P. Ashokkumar, and M. R. Rizvi. 2020. Identifying the attack surface for IoT network. *Internet of Things* 9 (2020), 100162. https://doi.org/10.1016/j.iot.2020.100162

[33] A. K. Sahu, S. Sharma, M. Tanveer, and R. Raja. 2021. Internet of Things attack detection using hybrid Deep Learning Model. *Computer Communications* 176 (2021), 146–154. https://doi.org/10.1016/j.comcom.2021.05.024

[34] M. Sarhan, S. Layeghy, N. Moustafa, and M. Portmann. 2021. NetFlow Datasets for Machine Learning-Based Network Intrusion Detection Systems. In *Big Data Technologies and Applications (Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering)*. Springer International Publishing, Cham, 117–135. https://doi.org/10.1007/978-3-030-72802-1_9

[35] J. Serrà, D. Álvarez, V. Gómez, O. Slizovskaia, J. F. Núñez, and J. Luque. 2019. Input Complexity and Out-of-distribution Detection with Likelihood-based Generative Models. https://openreview.net/forum?id=SyxIWpVYvr

[36] Bernadette J. Stolz, Jared Tanner, Heather A. Harrington, and Vidit Nanda. 2020. Geometric anomaly detection in data. *Proceedings of the National Academy of Sciences* 117, 33 (2020), 19664–19669. https://doi.org/10.1073/pnas.2001741117

[37] S. Tweneboah-Koduah, K. E. Skouby, and R. Tadayoni. 2017. Cyber Security Threats to IoT Applications and Service Domains. *Wireless Personal Communications* 95, 1 (2017), 169–185. https://doi.org/10.1007/s11277-017-4434-6

[38] K. Zhao and L. Ge. 2013. A Survey on the Internet of Things Security. In *9th International Conference on Computational Intelligence and Security (CIS)*. 663–667. https://doi.org/10.1109/CIS.2013.145

[39] Shuo Zhou, Antoinette Tordesillas, Mehdi Pouragha, James Bailey, and Howard Bondell. 2021. On local intrinsic dimensionality of deformation in complex materials. *Scientific Reports* 11, 1 (2021), 10216. https://doi.org/10.1038/s41598-021-89328-8