



**Budapest University of Technology and Economics**  
Faculty of Electrical Engineering and Informatics  
Department of Network Systems and Services

# Improving the robustness of similarity-based IoT malware detection methods against adversarial samples

MASTER'S THESIS

*Author*

József Sándor

*Advisors*

Dr. Levente Buttyán  
Roland Nagy

October 16, 2023

# Contents

<b>Abstract</b>	<b>i</b>
<b>Kivonat</b>	<b>ii</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 Similarity-based IoT malware detection</b>	<b>4</b>
2.1 Binary similarity hash function . . . . .	4
2.2 SIMBioTA . . . . .	5
2.2.1 Malware analysis at the backend . . . . .	7
2.2.2 Detection process on the IoT device . . . . .	7
2.3 SIMBioTA-ML . . . . .	8
2.3.1 Design choices for machine learning . . . . .	8
2.4 Performance . . . . .	9
2.4.1 True positive detection rate . . . . .	9
2.4.2 False positive detection rate . . . . .	10
<b>3 Evasion of IoT malware detection</b>	<b>12</b>
3.1 Overview of the adversarial examples problem . . . . .	12
3.2 Strategies for creating adversarial examples . . . . .	13
3.2.1 Overview of strategies . . . . .	13
3.2.2 Strategy 1: Chunker . . . . .	14
3.2.3 Strategy 2: Disguiser . . . . .	15
3.3 Measurement . . . . .	16
3.3.1 Dataset . . . . .	16
3.3.2 Setup . . . . .	17
3.3.3 Results . . . . .	19
3.3.4 Discussion . . . . .	19
<b>4 Adversarial training on SIMBioTA-ML</b>	<b>23</b>

4.1	Setup . . . . .	23
4.2	Results . . . . .	25
4.3	Discussion . . . . .	26
<b>5</b>	<b>PATRIoTA inspired by SIMBIoTA</b>	<b>31</b>
5.1	Design . . . . .	31
5.1.1	Overview . . . . .	32
5.1.2	Particle size and similarity threshold . . . . .	33
5.1.3	Detection threshold . . . . .	33
5.1.4	Optimal configurations . . . . .	33
5.2	Evaluation . . . . .	36
5.2.1	Experiment design . . . . .	36
5.2.2	Detection capability . . . . .	37
5.2.3	Storage requirement . . . . .	37
5.2.4	Run time performance . . . . .	39
5.3	Discussion . . . . .	39
<b>6</b>	<b>Related work</b>	<b>42</b>
6.1	ML-based (IoT) malware detection . . . . .	42
6.2	Adversarial examples and robustness analysis . . . . .	42
<b>7</b>	<b>Conclusion</b>	<b>44</b>
	<b>Acknowledgements</b>	<b>45</b>
	<b>Bibliography</b>	<b>46</b>

## HALLGATÓI NYILATKOZAT

Alulírott *Sándor József*, szigorló hallgató kijelentem, hogy ezt a diplomatervet meg nem engedett segítség nélkül, saját magam készítettem, csak a megadott forrásokat (szakirodalom, eszközök stb.) használtam fel. Minden olyan részt, melyet szó szerint, vagy azonos értelemben, de átfogalmazva más forrásból átvettem, egyértelműen, a forrás megadásával megjelöltem.

Hozzájárulok, hogy a jelen munkám alapadatait (szerző(k), cím, angol és magyar nyelvű tartalmi kivonat, készítés éve, konzulens(ek) neve) a BME VIK nyilvánosan hozzáférhető elektronikus formában, a munka teljes szövegét pedig az egyetem belső hálózatán keresztül (vagy autentikált felhasználók számára) közzétegye. Kijelentem, hogy a benyújtott munka és annak elektronikus verziója megegyezik. Dékáni engedéllyel titkosított diplomatervek esetén a dolgozat szövege csak 3 év eltelte után válik hozzáférhetővé.

Budapest, 2023. október 16.

---

*Sándor József*  
hallgató

# Abstract

Detecting malicious files before they are executed is a fundamental defense mechanism in computer-based systems, including in embedded systems and the Internet-of-Things (IoT). Indeed, with the dramatic increase of the number of deployed embedded IoT devices in the world, the number of known attacks against them has also increased in the recent past, and such attacks include infecting them with malware. Therefore, malware detection on embedded IoT devices became an active research area, resulting in several IoT malware detection methods. In this paper, we introduce two recently proposed solutions, SIMBIO TA and SIMBIO TA-ML. They both use binary similarity measures for detecting even previously unseen malware, and they have good detection performance, while being very resource efficient at the same time. In addition, SIMBIO TA-ML improves the malware detection capability of SIMBIO TA by taking advantage of machine learning.

The main problem addressed in this work is that current IoT malware detection methods, such as SIMBIO TA and SIMBIO TA-ML, are vulnerable to adversarial evasion techniques. This means that, knowing how the malware detection method works, attackers can create specifically crafted malware samples that mislead the detector and evade detection. Unfortunately, existing solutions are not necessarily robust against these type of attacks. We demonstrate that creating malware samples that evade detection is relatively easy by proposing two simple adversary example creation strategies and showing that the robustness of SIMBIO TA and SIMBIO TA-ML against them is rather weak. Both strategies append bytes to the end of an existing malware sample such that the malware remains functional, but the new sample becomes dissimilar to the original sample, hence, evading detection by SIMBIO TA and SIMBIO TA-ML. The first strategy appends chunks of the original sample, whereas the second strategy appends an entire benign file. We measure the robustness of SIMBIO TA and SIMBIO TA-ML against these strategies by measuring their detection accuracy on the crafted adversarial samples. It turns out that SIMBIO TA-ML is somewhat robust against the first strategy, but both SIMBIO TA and SIMBIO TA-ML completely fail against the second one.

To overcome this problem, we propose two different solutions for antivirus companies. The first approach is adversarial training, which means in our case that the training set of the malware detector algorithm is extended with samples that are crafted by using the adversarial evasion strategies that we proposed. The second approach, named PATRIO TA, involves searching for fixed-size parts of known malware within the binary of a suspicious file. Based on our measurements, both solutions are effective in achieving high detection accuracy for both the original malware samples and the adversarial samples. The price that we have to pay for this remarkable robustness is the increased training time and the increased size of the detection model, however, we argue that both are bearable in practice.

# Kivonat

A rosszindulatú fájlok futtatás előtti felismerése alapvető védelmi mechanizmus a számítógép-alapú rendszerekben, beleértve a beágyazott rendszereket és a tárgyak internetét (Internet-of-Things, IoT) is. A beágyazott IoT-eszközök számának drámai növekedésével az ellenük irányuló ismert támadások száma is megnőtt, melyek között hangsúlyosak a malware típusú támadások. Ezért a beágyazott IoT-eszközökön a malware programok detektálása aktív kutatási területté vált, és ennek eredményeképpen számos IoT malware detekciós módszer született. Ebben a tanulmányban két nemrégiben javasolt megoldást mutatunk be, a SIMBIO TA-t és a SIMBIO TA-ML-t. Mindkettő bináris hasonlósági mértékeket használ a korábban nem látott rosszindulatú programok felismerésére, továbbá azon kívül, hogy sikeresen felismerik a rosszindulatú fájlokat, hatékonyan bánnak az erőforrásokkal is. Ezen túlmenően a SIMBIO TA-ML a gépi tanulás előnyeit kihasználva javítja a SIMBIO TA detekciós képességét.

A fő probléma, amellyel ez a dolgozat foglalkozik, az, hogy a jelenlegi IoT malware detekciós módszerek, mint például a SIMBIO TA és a SIMBIO TA-ML, sebezhetőek az ún. adversarial technikákkal szemben. Ez azt jelenti, hogy a támadó a malware-t detektáló módszer működésének ismeretében olyan speciálisan kialakított malware mintákat hozhat létre, amelyek elkerülik a detekciót. Sajnos a meglévő megoldások nem feltétlenül robusztusak az ilyen típusú támadásokkal szemben. Ebben a tanulmányban két egyszerű adversarial minta létrehozási stratégiával megmutatjuk, hogy detektálást elkerülő rosszindulatú mintákat viszonylag könnyű létrehozni, és látni fogjuk, hogy a SIMBIO TA és a SIMBIO TA-ML robusztussága ezekkel szemben meglehetősen gyenge. Mindkét stratégia bájtokat csatol egy meglévő malware végéhez úgy, hogy a malware működőképes marad, de az új minta nem hasonlít az eredeti mintához, és így kikerüli a SIMBIO TA és a SIMBIO TA-ML általi észlelést. Az első stratégia az eredeti minta egy darabját, míg a második stratégia egy teljes jóindulatú fájlt csatol. A SIMBIO TA és a SIMBIO TA-ML robusztusságát ezekkel a stratégiákkal szemben úgy határozzuk meg, hogy megmérjük milyen pontossággal detektálják az adversarial mintákat. Kiderül, hogy a SIMBIO TA-ML viszonylag robusztus az első stratégiával szemben, de mind a SIMBIO TA, mind a SIMBIO TA-ML teljesen kudarcot vall a második stratégiával szemben.

Ennek a problémának a leküzdésére két különböző megoldást javasolunk a vírusirtó cégek számára. Az első megközelítés az adversarial training, ami a mi esetünkben azt jelenti, hogy a rosszindulatú szoftvereket észlelő algoritmus tanító halmazát olyan mintákkal bővítjük, amelyeket az általunk javasolt adversarial stratégiák felhasználásával készítettünk. A második megközelítés, amelyet PATRIO TA-nak nevezünk, ismert rosszindulatú programok rögzített méretű részeit keresi egy gyanús fájl bináris állományán belül. Méréseink alapján mindkét megoldás hatékony, abban az értelemben, hogy képes magas felismerési pontosságot elérni mind az eredeti rosszindulatú szoftverminták, mind az adversarial minták esetében. Az ár, amelyet ezért a megnövekedett robusztusságért fizetnünk kell, a megnövekedett tanítási idő és a detektáló modell megnövekedett mérete, de kijelenthető, hogy mindkettő elviselhető mértékű a gyakorlatban.

# Chapter 1

## Introduction

An embedded device is a specialized device meant for specific purposes and it is usually embedded as part of a larger system. Nowadays, these devices are widespread and could be surprisingly diverse. The collection of these embedded devices that are connected to the Internet, together with their often cloud-based backend infrastructure, is called the Internet-of-Things (IoT).

Just like other types of computers, embedded IoT devices have security weaknesses. IoT devices can be found everywhere (e.g., healthcare, transportation, agriculture), therefore, their vulnerabilities represent a huge attack surface for attackers. IoT devices are desirable targets for attackers, because they handle sensitive information or they control critical processes. Indeed, with the dramatic increase of the number of deployed embedded IoT devices in the world, the number of known attacks against them has also increased in the recent past, and such attacks include infecting them with malware. Furthermore, they are highly connected over the Internet, so malware could spread easily from one device to the other. One of the most infamous examples is the Mirai malware [5], which infected hundreds of thousands of IoT devices and launched one of the largest distributed denial of service attacks against Internet-based services in 2016. But the IoT threat landscape includes other malware families as well, such as Gafgyt, Tsunami, and Dnsamp [11].

Malware detection is an essential part of modern defense mechanisms used in computer-based systems. Malware detection approaches can be categorized into signature-based, heuristic, and cloud-based approaches [7]. In the past, antivirus products only used signatures. A signature, in this context, is a short sequence of bytes that uniquely identifies a set of variants of a malware. Malware detection algorithms scan files and search for such signatures. If the given signature is found in the file, the file is considered malware. In practice, however, signature-based detection has significant disadvantages. It is expensive, because signatures are usually created manually by experts [1]. Furthermore, signature-based detection can be misled with various techniques (e.g., packing, encryption, obfuscation, code polymorphism). These techniques keep the functionality of the malware, while they make their characteristic signatures disappear.

Heuristic malware detection relies on rules, created by experts that capture more complex static patterns in malware than simple signatures do. Consequently, compared to signature-based approaches, heuristic techniques can detect a larger set of variants of the same malware. Yet, even this approach is unable to cope with obfuscation techniques. Furthermore, the threat landscape is constantly evolving with both new types of malware

and variations of existing malware<sup>1</sup>. Both cases require new signatures and heuristic rules to be generated constantly and this poses a major challenge for antivirus companies.

Therefore, there is significant effort to automate the detection process using machine learning [14, 40, 38]. In order to extract features for machine learning, static and dynamic program analysis techniques are used [32]. Features include instruction-level data, data related to control-flow, invoked API functions and system calls, and messages sent over the network. These features are used to train machine learning models for malware detection. Furthermore, machine learning requires lots of training data, i.e., benign and malicious samples in this case.

Regarding the system architecture, nowadays, antivirus products install a client-side component on the users' machines, which typically performs signature-based and heuristic detection. If this client component cannot determine whether a sample is malicious or not, then it sends the sample to a server, which performs a more in-depth analysis, including e.g., dynamic behavior analysis in a sandbox environment. We refer to this architectural setup as cloud-based malware detection. Thanks to using dynamic behavior analysis, it is very effective, it can even cope with advanced evasion techniques (e.g., obfuscation, code polymorphism). Cloud-based approach can be applied to resource constrained IoT devices as well [35].

While many attacks can be successfully prevented with these approaches, unfortunately, no matter how good a malware detection system is, attackers constantly work on methods to evade their detection (this is a cat-and-mouse game). In particular, they want to construct malware samples that have the same function as older samples but not recognized by detectors. In case of machine learning, such kind of inputs are called *adversarial examples*. The concept of adversarial examples can also be adopted in the domain of malware detection (being machine learning-based or otherwise): an adversarial example in this context would be a specially crafted malware sample that evades detection by a specific detection method. The degree of vulnerability against adversarial examples defines the robustness of the system. History shows that traditional signature and heuristic malware detection is not robust against adversarial samples, which explains the large number of polymorphic malware. And unfortunately, it has been shown in the literature [6, 34, 29] that machine learning based malware detectors can also be misled easily.

In this work, we compare two recent IoT malware detection solutions, SIMBIO TA and SIMBIO TA-ML, in terms of robustness. SIMBIO TA (SIMilarity Based IoT Antivirus) [37] is an effective and efficient IoT antivirus solution. SIMBIO TA is similar to traditional signature-based solutions, but it uses TLSH hash values of known malware instead of raw binary signatures. TLSH [27] is a similarity hash algorithm, which means it outputs similar hash values for similar inputs. SIMBIO TA-ML [28] improves the malware detection capability of SIMBIO TA with machine learning. In case of SIMBIO TA-ML, for training the machine learning model, feature vectors are extracted from the TLSH hash value of files.

For the comparison of SIMBIO TA and SIMBIO TA-ML in terms of robustness against adversarial examples, we design two adversarial example creation strategies. The purpose of each strategy is to create adversarial examples that evade similarity-based malware detection. Both strategies modify existing malware samples by appending extra bytes to them such that those bytes are never executed but they make the modified samples dissimilar to the original ones. The first strategy adds chunks of the original sample to the malware

---

<sup>1</sup><https://www.sophos.com/en-us/medialibrary/pdfs/technical-papers/sophoslabs-2019-threat-report.pdf>, (accessed: October 16, 2023)  
<https://www.ntsc.org/assets/pdfs/cyber-security-report-2020.pdf>, (accessed: October 16, 2023)



and ensures that a certain target difference is achieved by doing so. The second strategy embeds a malware into a known benign file and ensures that the resulting sample becomes similar to the benign file (and hence dissimilar to the original malware sample). We show by measurements that SIMBIO<sub>TA</sub>-ML is robust against the first strategy, but it can be misled by the second one, while SIMBIO<sub>TA</sub> has poor robustness against both strategies.

To overcome this problem and increase the robustness of both SIMBIO<sub>TA</sub> and SIMBIO<sub>TA</sub>-ML, we propose two solutions, one for each. In the case of SIMBIO<sub>TA</sub>-ML, we employ adversarial training, a concept originally used in the image recognition domain to enhance the robustness of machine learning-based models against adversarial examples. We adopt this approach in the domain of malware detection and demonstrate its effectiveness. Adversarial training in our case means that the training set of the malware detector algorithm is extended with samples that are crafted using the adversarial evasion strategies we proposed. We measure the detection accuracy of SIMBIO<sub>TA</sub>-ML trained on such an extended training set and show that it remains high both for the original malware samples and for the adversarial samples. Secondly, to increase the robustness of SIMBIO<sub>TA</sub> while preserving its advantages, we propose PATRIO<sub>TA</sub>, a modified (upgraded) version of SIMBIO<sub>TA</sub>. The main idea of PATRIO<sub>TA</sub> is to split malware samples known to the antivirus provider into multiple fixed-size parts, referred to as *particles*, and to perform the same operations on those particles as the operations performed by SIMBIO<sub>TA</sub> on entire samples. This involves the antivirus provider building a similarity graph of known particles, computing its dominating set, and distributing the TLSH hash values corresponding to the particles in the dominating set to the clients. The clients also split any file to be scanned into particles, computing the TLSH hash values for each particle. If a threshold number of these computed hashes are similar to the hashes in the dominating set, the file is detected as malware. The price that we have to pay for this remarkable robustness in both cases is the increased training time and the increased size of the detection model, however, we argue that both are bearable in practice. Additionally, while PATRIO<sub>TA</sub> enhances robustness, it also comes with an unfortunate trade-off of increased detection time (the time required for decision making).

This document is organized as follows. In Chapter 2 we take a closer look at the design and performance of SIMBIO<sub>TA</sub> and SIMBIO<sub>TA</sub>-ML. Chapter 3 presents the strategies for creating adversarial examples. Additionally, we describe our methodology for measuring the performance of the proposed adversarial example creation strategies and present the results of our measurements. In Chapter 4, we present our first proposed countermeasure, adversarial training on SIMBIO<sub>TA</sub>-ML, along with the measurement results demonstrating its effectiveness in enhancing the robustness of SIMBIO<sub>TA</sub>-ML against evasion by adversarial samples. In the first part of Chapter 5, we introduce PATRIO<sub>TA</sub> and present its design considerations in detail. In the second part, we compare PATRIO<sub>TA</sub> to SIMBIO<sub>TA</sub> in terms of malware detection capability and robustness against adversarial sample creation strategies. In Chapter 6 we show the current state of the art in this specific scientific field. Finally, Chapter 7 concludes our work.

## Chapter 2

# Similarity-based IoT malware detection

Although, IoT malware detection is a challenging problem, there have been solutions proposed in the literature [35, 20]. In this work, we are interested in the similarity-based IoT malware detection solutions called SIMBIO TA and SIMBIO TA-ML, which have been proposed recently, and which have remarkable malware detection capabilities, while being resource efficient at the same time. Before we delve into the architecture of these solutions, we get to know the meaning of similarity hashes that form the basis of the mentioned systems.

### 2.1 Binary similarity hash function

Cryptographic hashes such as MD5 and SHA-1 are used for many data mining and security applications. The collision resistance property of cryptographic hash functions make them suitable for unique identification of files in practice. However, if a single byte of a file is changed, then its cryptographic hash will be a completely different hash value. The situation is different for similarity hashes: for similar inputs, binary similarity hash functions output similar hash values. SIMBIO TA and SIMBIO TA-ML use TLSH hash function as the basis of their system [27], which is also a similarity based hash function.

TLSH has a lightweight calculation time in the range of milliseconds on contemporary personal computers, which makes it suitable in the context of malware detection even on resource constrained IoT devices. A TLSH value is relatively short, it can be represented in 36 bytes (35 byte hash value +1 byte for versioning), which is also an advantage. More specifically, computing a TLSH hash value involves the following steps [27]:

1. Process the raw byte string using a sliding window of size 5 to populate an array of bucket counts.
2. Calculate quartile points  $q_1$ ,  $q_2$ , and  $q_3$  based on the buckets' values.
3. Construct the hash value's header based on the quartile points, the length of the byte string, and a checksum on its content.
4. Construct the hash value's body.

The first three bytes of the resulting TLSH hash value is a header with following parts:

- the first byte is a checksum value;
- the second byte stores the so-called L value, which is calculated from (the logarithm of) the length of the original byte sequence;
- the two nibbles of the third byte are called the Q1 and Q2 ratios, and they are computed from the quartile points q1 and q3, and the quartile points q2 and q3, respectively.

The rest of the bytes are the binary representations of the 128 buckets that TLSH uses during the construction of the hash value quantized to two bits.

As an illustration, let us consider the following prefix of a TLSH hash value, represented in hexadecimal format:

T1 B3 C3 09 A5 BC 43 9B 4F CA C3 DB F6 . . .

The first two character of the TLSH hash (T1) is the version number. The version number is followed by the 3-byte header. The first byte of the header is a checksum, which has the value of hexadecimal B3 in our example. This is followed by the L value, which is hexadecimal C3 in this case. Next come the Q1 and Q2 ratios, which are hexadecimal 0 and 9, respectively, in the example. The remaining bytes are the binary representations of the buckets turned into hexadecimal numbers. As each bucket value is represented by two bits, the next hexadecimal number A, in the example, encodes the 2-bit values 10 and 10 of the first two buckets. The same way, the next hexadecimal number 5 encodes the 2-bit values 01 and 01 of the next two buckets, etc.

We can measure the similarity of two inputs by comparing their TLSH hash values with a comparison algorithm [27]. The output of the comparison algorithm is an integer number. The minimum value of this number is 0, which means that the two files of the compared hashes are almost identical. A higher score should represent that there are more differences between the inputs. This algorithm calculates the similarity value of the two given TLSH hash with various weighting. E.g., differences in hash value’s header are taken into account with a larger weight than the differences in the hash value’s body. For more details on the calculation of TLSH values and TLSH differences, we refer the reader to [27].

It is important to note that the similarity hash functions and comparison algorithms operate only on raw byte sequences, therefore they are suitable for capturing static byte level similarity of binary files, but nothing more.

## 2.2 SIMBIO TA

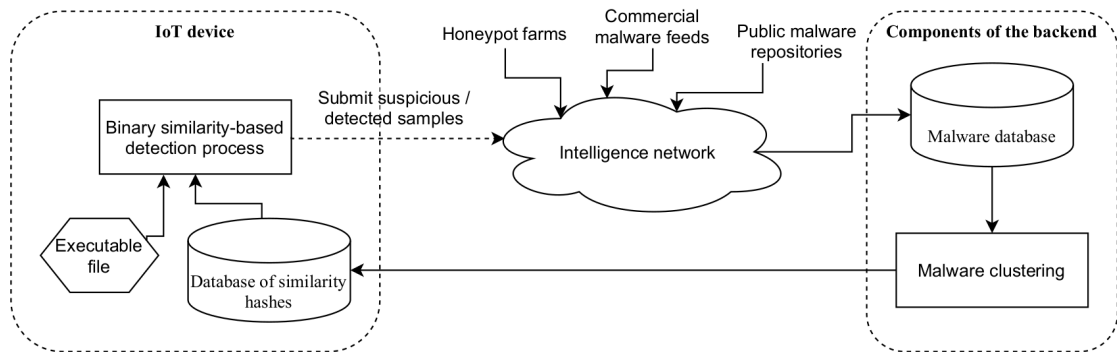
SIMBIO TA is an IoT malware detection solution that was proposed in [37]. It saves storage, memory, computation, and bandwidth, which resources are constrained in IoT field. In a certain sense, SIMBIO TA is a hybrid solution, it combines preferably the properties of signature-based and cloud-based solutions.

SIMBIO TA efficiently detects malware by leveraging the similarity between scanned files and known malware samples. The approach takes advantage of the fact that malware belonging to the same family tends to exhibit similarities, while samples from different families, as well as benign programs, are dissimilar. This leads to a distinct separation between the similarity graph of known malware samples and that of benign programs, as depicted in Figure 2.1. In this graph, vertices represent binary files, and an edge connects two vertices if the corresponding files share a certain level of similarity. Typically, each



**Figure 2.1:** Illustration of a similarity graph of a set of binaries. Each vertex represents a binary, and an edge connects two vertices if the corresponding files exhibit similarity based on a specified similarity measure. In this specific illustration, the files are considered similar if their TLSH difference score is below the threshold of 40. Malware binaries are denoted by red nodes, while benign binaries are depicted as green nodes. Notably, there is a clear distinction between malware and benign files, and the similarity graph of malware exhibits significant clustering.

cluster in the similarity graph can be represented by a few representative samples, where all cluster members are similar to at least one of these representatives. Therefore, for malware detection purposes, knowledge of the representative samples is sufficient. Scanned files that resemble any of the representatives are likely to be malware, while files that do not exhibit similarity to any of the representatives are likely to be benign.



**Figure 2.2:** Architecture of SIMBIoTA.

We can see the high-level architecture of SIMBIO TA in Figure 2.2. It consists of 3 major components: IoT device (i.e., client), backend (i.e., server), and intelligence network. The intelligence network (e.g., honeypot farms, commercial malware feeds, and public malware repositories) provides malware samples for the malware database of the backend. SIMBIO TA is similar to the signature-based approach, but it uses TLSH hash values instead of signatures. SIMBIO TA is also similar to the cloud-based approach, because the backend does the resource-intensive tasks. The backend processes the malware samples and it creates a representative subset of TLSH hashes that is pushed to the IoT device. The IoT device uses this small database of TLSH hashes, and it runs a lightweight algorithm to detect malware, based on binary similarity.

Similarity hashes are a good alternative to signatures in IoT malware detection. Firstly, the similarity hash used by SIMBIO TA (i.e., the TLSH hash) is represented in a very short sequence of bytes (36 bytes). Secondly, based on binary similarity property, one hash can fully represent a group of malware. Hence, all malware samples on the backend can be covered with a relatively small database on the IoT device. In addition, computing similarity hashes does not require manual work of experts, but it can be completely automated. Furthermore, another advantage of similarity hashes is their short calculation time. Processing a single file (i.e., hash generation time + hash comparison time) takes only a few milliseconds even on CPU constrained devices.

### 2.2.1 Malware analysis at the backend

The backend database of SIMBIO TA is built from the malware samples collected from the intelligence network. Typically, thousands of samples are collected each day. The IoT device could not handle this number of TLSH hashes, so the backend can transmit only a few TLSH hash values. These TLSH hash values form a representative subset that represents the whole backend malware database.

We can imagine the malware database as a graph, where nodes are the malware samples, and two nodes are connected if the TLSH similarity score of their samples is below a selected threshold<sup>1</sup>. More precisely, the mentioned representative subset of samples must form a dominating set<sup>2</sup> for the imagined graph. SIMBIO TA uses a simple greedy algorithm for constructing the dominating set: if a new sample received by the backend is not similar to any of the samples in the current dominating set, it adds the new sample to the dominating set, otherwise it moves on to the next new sample.

Based on the above SIMBIO TA can create the representative subset of the malware database and it can extend if it is necessary. The backend sends the TLSH hashes of this dominating set to the client. If the dominating set on the server side is expanding, then the server informs the client with the updates.

### 2.2.2 Detection process on the IoT device

On IoT device the detection process is called before executing a file. If the TLSH distance of the given file's TLSH hash and one of the hashes in the dominating set is below the

---

<sup>1</sup>In case of SIMBIO TA, 40 is used as threshold value, which was selected by extensive empirical analysis described in [10].

<sup>2</sup>A dominating set for a graph  $G = (V, E)$  is a subset  $D$  of  $V$  such that every vertex not in  $D$  is adjacent to at least one member of  $D$ . In graph theory minimal dominating set problem is a classical NP-complete decision problem. Therefore it is believed that there may be no efficient algorithm that finds a smallest dominating set for all graphs.

selected threshold, then it is considered as a malicious file. According to the client’s needs and possibilities, it can further customize his policy of use. SIMBIO TA’s client has a great advantage with respect to other cloud-based malware detection system’s clients, because with its local database it can operate even if the backend is unavailable.

## 2.3 SIMBIO TA-ML

SIMBIO TA-ML was proposed in [28]. The purpose of SIMBIO TA-ML is to improve the malware detection capability of SIMBIO TA with machine learning. In order to do so, SIMBIO TA-ML replaces the dominating set construction by machine learning. The modified architecture is shown in Figure 2.3. On embedded IoT device, the database of TLSH hash values is replaced with a machine learning model. Hence, in place of the lightweight detecting algorithm, the client only have to give the examined file to the stored machine learning model for detection. The machine learning model is trained on the backend using both malicious and benign samples. Therefore, SIMBIO TA’s intelligence network is extended with sources that also provide benign samples for the backend. Benign samples could be received from IoT software vendors or from public software repositories.

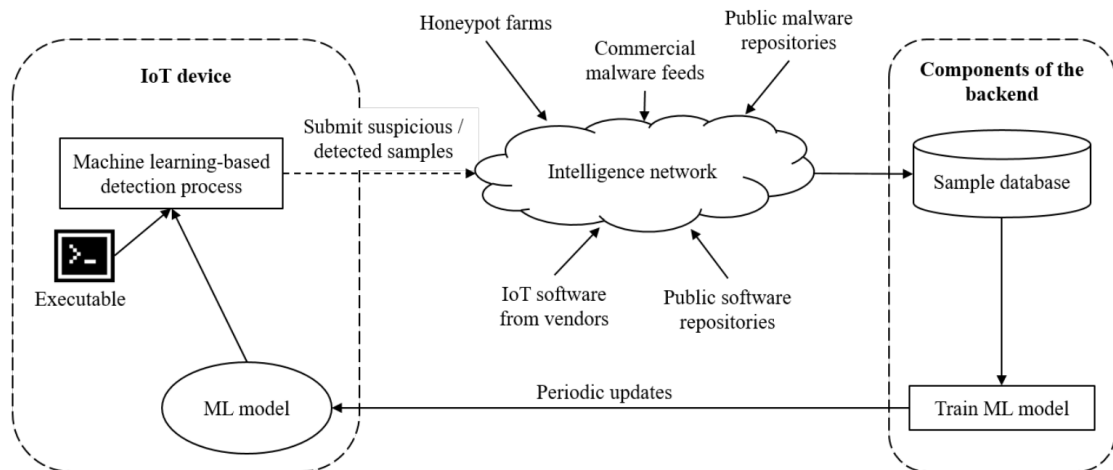


Figure 2.3: Architecture of SIMBIO TA-ML.

### 2.3.1 Design choices for machine learning

Machine learning models for malware detection must be trained using features that represent important qualities of executable files. In general, features can be derived using static or dynamic program analysis. However, dynamic program analysis (i.e., monitoring a program’s execution) is not very economical, and that would be important in case of IoT devices. Therefore, the extraction of feature vectors has to be lightweight and has to be done statically.

TLSH hash values can be considered static features because their calculation involves only the processing of the raw bytes in the program file. The TLSH hash value is transformed into 131 features by splitting the hash value into smaller parts. Specifically, from the header the L value, the Q1 ratio, and the Q2 ratio is taken. The bytes representing buckets are split into bit pairs, which gives 128 2-bit features for the 128 buckets. A random forest classifier is trained on these extracted features. Choosing a random forest classifier is advantageous because it automatically filters non-predictive features [9].

## 2.4 Performance

Before we show the performance of SIMBIO TA and SIMBIO TA-ML on adversarial examples, we take a look at their results on the unmodified samples. As for the results, we mention only the false and true positive rates of the two systems, because in the context of this paper only these are relevant. In addition to the measurements presented below, one can read about the experiment in more detail in [28].

For the experiments of SIMBIO TA and SIMBIO TA-ML the same data set is used. We discuss the origin and composition of the samples in Section 3.3.2. For now, it is enough to know that there is a relatively large sample data set. This data set contains malicious and benign files and these files are executables written for either the ARM or the MIPS architecture.

The same evaluation is taken separately on ARM and MIPS samples, however in the following we refer to them together. SIMBIO TA and SIMBIO TA-ML use the same experiment design. The experiment uses samples created between January 1st, 2018 and September 15th, 2019. This time interval is divided into weeks. Both SIMBIO TA and SIMBIO TA-ML receive updates for their detection methods at the beginning of each week.

Malicious samples are organized into weekly batches based on the date they were first seen. 10% of each weekly batch is available to the backend, this training set represents the knowledge provided by the intelligence network. The other 90% is never shown to the backend, this is the test set. The malware detection rate on the test set gives the system’s true positive detection rate.

SIMBIO TA-ML requires a balanced data set for training and testing the machine learning model, so the same number of malicious and benign samples per week is needed. However, there is no information about when the benign samples were first seen. Therefore, benign samples are randomly assigned to be part of either the training or test set. Each week, the same number of benign samples are selected from the benign training set as the number of malicious samples in the malware training set. Selected benign samples are available to the backend for training purposes. The test set of benign samples is selected in the same way. Samples of the benign test set are never shown to backend and their evaluation gives the system’s false positive detection rate.

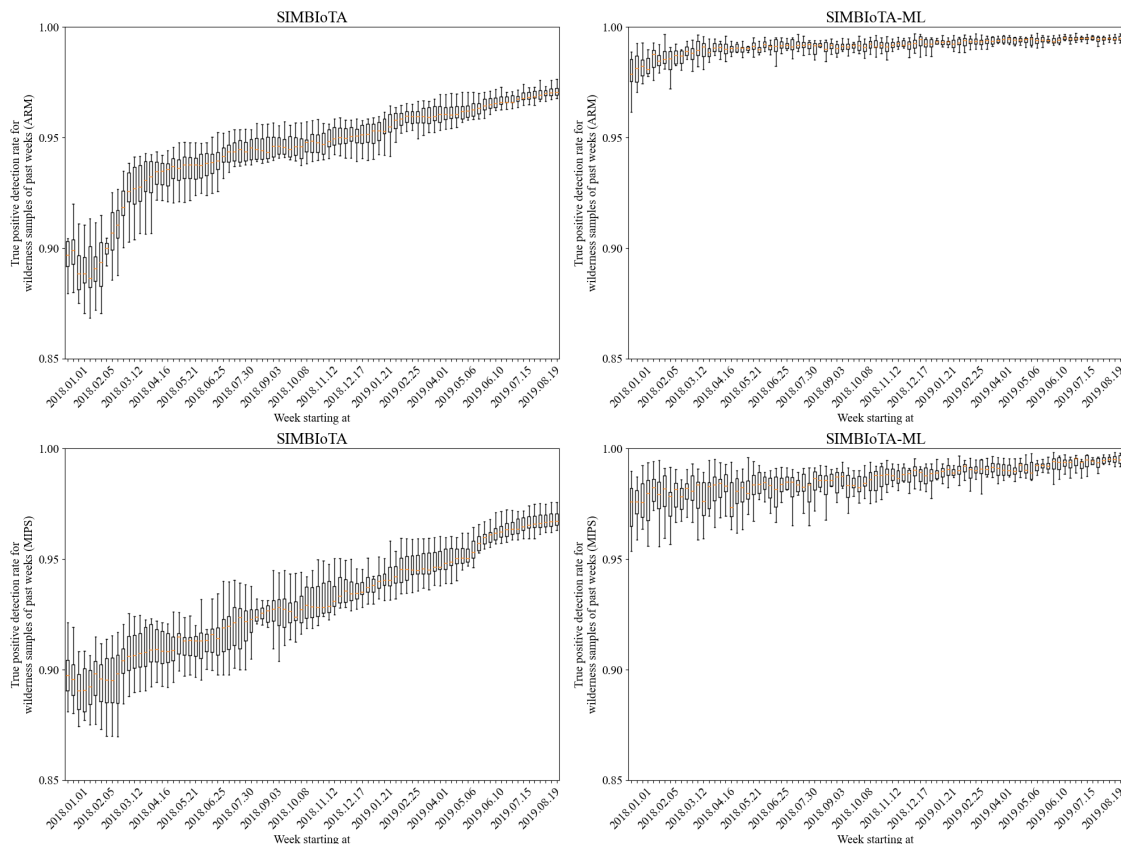
The method of assigning benign samples to the training and test sets introduces randomness into the experiment. To eliminate the effects of this randomness on the measurement results, the experiment is repeated 12 times and traditional box plots are used to present the results. The data points of box plots show the results of the 12 runs of the experiment for each week.

### 2.4.1 True positive detection rate

There are two approaches for measuring the true positive detection rate of SIMBIO TA and SIMBIO TA-ML. The first approach evaluates the test sets of all previous weekly batches, while the second approach takes only the current weekly batch.

Results of the first approach is shown on Figure 2.4. The left-hand side of the figure shows the performance of SIMBIO TA and the right-hand side shows the performance of SIMBIO TA-ML. Both solutions show a learning curve for both the MIPS and the ARM architectures, i.e., their true positive detection rate improves as time passes and more samples are made available to the backend. However, SIMBIO TA-ML consistently out-

performs SIMBIO<sub>TA</sub> by having a true positive detection rate above 95% throughout the measurement.



**Figure 2.4:** Box plot of the true positive detection rate for malware test sets of all previous weeks.

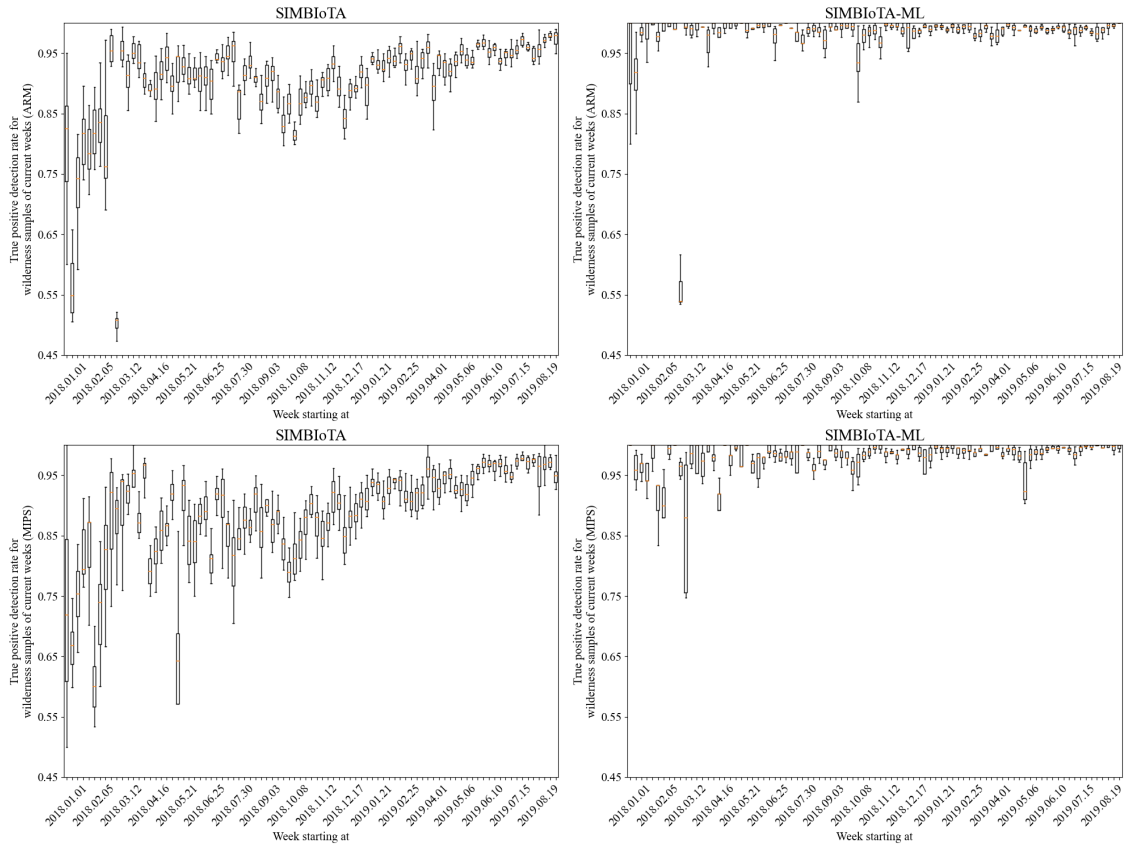
Results of the second approach is shown in Figure 2.5. The left-hand side of the figure shows the performance of SIMBIO<sub>TA</sub> and the right-hand side shows the performance of SIMBIO<sub>TA</sub>-ML. SIMBIO<sub>TA</sub>'s performance varies in time and its performance reaches 90-95% only for the second half of the measurement. SIMBIO<sub>TA</sub>-ML also shows variations in its true positive detection rate but the variation is smaller than that of SIMBIO<sub>TA</sub>, and performance stays above and around 95% for the majority of the experiment. Therefore, it can be stated that SIMBIO<sub>TA</sub>-ML outperforms SIMBIO<sub>TA</sub> in this case as well.

## 2.4.2 False positive detection rate

In order to measure the false positive detection rate of the systems, the following experiment is executed. SIMBIO<sub>TA</sub> does not use benign samples for learning, thus, all benign samples are submitted to SIMBIO<sub>TA</sub> for measuring false positive detection rate. In the case of SIMBIO<sub>TA</sub>-ML, however, the benign test set of actual weakly batch is given to detection process.


As reported in [37] SIMBIO<sub>TA</sub> did not detect any benign samples as malicious, hence achieved a false positive rate of 0. However, SIMBIO<sub>TA</sub>-ML has a false positive detection rate 1% on average, as Figure 2.6 shows. This phenomenon is common in machine learning field. Interestingly, MIPS samples show higher false positive detection rate than ARM samples, but it decreases in both cases as time goes on. Overall, it can be stated that





**Figure 2.5:** Box plot of the true positive detection rate for malware test set of the current week.

while SIMBIoTA-ML’s false positive detection rate is higher than SIMBIoTA’s, it is still acceptable for malware detection.



**Figure 2.6:** Box plot of the false positive detection rate for benign samples in the test set for SIMBIoTA-ML.


## Chapter 3

# Evasion of IoT malware detection

We can say based on what we have seen so far that SIMBIO TA and SIMBIO TA-ML appropriately recognize malicious files. However, both of our systems detection mechanism is based on binary similarity. What if the attacker knows this? By using this information, could the attacker increase the chances of evading detection of his malware? Can the attacker achieve that his malware is classified as a benign file by the detecting system? If so, what strategies might he have to do this? In this chapter we are looking for answers to these questions.

### 3.1 Overview of the adversarial examples problem

Before we delve into answering the previous questions, we take a look at the adversarial examples problem in general. Besides the malware detection context, the adversarial example problem appear in several other machine learning field. A prominent example is image recognition, which creates adversarial example with so-called image perturbation. Here, we add some adversarial perturbation noise to an image, such that no difference is visible for the human eye, but the given machine learning model confidently does a wrong prediction. This popular example on Figure 3.1 shows how a Deep Neural Network (in this case, GoogLeNet) can be misled this way [15].



**Figure 3.1:** Adversarial example against image recognition machine learning model that achieves misclassification of panda as gibbon.

Even though the concept of adversarial examples were first defined in the machine learning field, they can be interpreted in the context of non-ML-based classifiers too. Hence, we can measure the robustness of SIMBIO TA and SIMBIO TA-ML against the same adversarial examples. In general, adversarial examples are those inputs that were specifically designed to cause the given model to make a mistake [8]. In our case this mistake would be

the misclassification of malicious samples as benign. The attacker’s interest is to create successful adversarial examples that evade malware detection.

In the machine learning field there are two major approaches for creating adversarial examples. We can craft adversarial examples in the feature space, which means we construct feature vectors instead of real inputs that mislead the model. Furthermore, we can create real inputs as adversarial examples by creating completely new samples or by modifying existing ones. In case of computer programs we would like to create adversarial examples that not only mislead the classifier but that remain executable too. Since there is no guarantee that an input reconstructed from a feature vector would be a meaningful computer program, constructing in feature space is excluded, so we have to create real inputs. Malware files are also computer programs, however creating a brand-new malware can be expensive, therefore, it looks more reasonable to modify existing ones.

## 3.2 Strategies for creating adversarial examples

To answer the questions asked at the beginning of the chapter, we have to think with the head of the attacker. Firstly, we assume that we have already a fancy malware and we do not want to spoil its functionality. The only problem is that SIMBIO TA-ML, but even SIMBIO TA, indicates correctly that it is malicious. However, we know that the similarity hash generation and comparison algorithms (which are also used by our detection systems) do not take into account the format of the inputs, they only consider raw sequences of bytes. Furthermore, we can even find out that this similarity hash function is the TLSH.

Our idea is that we can mislead the detection system if we could manipulate the TLSH hash value of our malware. Knowing the nature of TLSH, to do so, we have to modify the raw binary. Targeted modification of the binary by manipulating the source code without spoiling the original functionality is not so trivial task<sup>1</sup>. It would be much easier to add a few extra bytes to the end of the binary that actually will never be executed, but will change the TLSH hash value. However, we have to do this expansion of the binary carefully, because too much growth can be noticeable for the defenses system.

We can increase the success rate of our attack in the IoT field, especially against malware detection, with the quantity of the adversarial examples. So the attacker needs economical solutions for creating adversarial examples. Simply adding bytes requires no particular sophistication from the attacker, and it is also economical in terms of resources.

### 3.2.1 Overview of strategies

For creating adversarial examples, we developed two strategies. The first one is called Chunker, the second is called Disguiser. These represent two different approaches. In case of Chunker we add to the malware chunks of itself and the goal is to increase the TLSH difference between the malware and the crafted adversarial example. In case of Disguiser we want to hide our malware in a benign file, so the goal is to decrease the TLSH difference between the benign file and the adversarial example. Moreover, these strategies are relatively simple, an attacker can easily implement them in a real-world situation.

---

<sup>1</sup>Techniques like obfuscation preserve the functionality of programs, but changes completely their binary. We do not deal with such techniques, because for SIMBIO TA and SIMBIO TA-ML obfuscated malware would seem totally new malware and their detection performance in this case has already been measured in [28].

### 3.2.2 Strategy 1: Chunker

As mentioned in Section 2.2, SIMBIO TA uses 40 as TLSH similarity threshold. So, our intuition is that if the TLSH distance between the original malware and our crafted adversarial example is at least 40, SIMBIO TA misclassifies it. Chunker’s main idea is to simply add bytes to the end of the malware binary. The question is, how many and what kind of bytes need to be added to reach our goal. If all attached bytes are constant or random, the byte entropy<sup>2</sup> would change and a static analyzer would easily detect it. It seems a reasonable solution to add some chunks from the original malware to itself. With this solution, if we choose properly the chunks, the byte entropy of the modified file will be almost the same as the original.

---

**Algorithm 1** *Chunker*

---

**Input:** malware binary  $MW$

**Output:** adversarial example  $ADVEX$

- 1:  $CHUNKS \leftarrow$  split  $MW$  to 20 equal parts
  - 2:  $CH \leftarrow$  element from  $CHUNKS$  w/ the closest entropy to  $MW$
  - 3:  $ADVEX \leftarrow MW$
  - 4: **repeat**
  - 5:      $ADVEX \leftarrow ADVEX + CH$
  - 6: **until**  $TLSH\_difference(MW, ADVEX) > 40$
  - 7: **return**  $ADVEX$
- 

The pseudocode of Chunker is presented as Algorithm 1, whose steps are as follows.

- Select an arbitrary malware binary file.
- Split the raw byte sequence of the given file to 20 equal parts. With this, we get 20 chunks, each is 5% of the original file size.
- Select the chunk with the entropy closest to the entropy of the original file.
- Add the selected chunk to the end of the binary file as many time as it is necessary to reach our goal (i.e., to get TLSH difference large enough between the crafted sample and the original one).

To find out how many chunks are needed to be added, we performed an empirical study. We take randomly 2000 samples from the whole malware data set (for description of whole data set see Subsection 3.3.1). To each malware we add different number of chunks and we calculate the TLSH difference between the original and the modified malware. The result is shown in Figure 3.2. One can see that in most cases, 4 chunks (i.e., 20% of the original malware) are enough to be added for reaching TLSH difference 40.

We can observe many boxplot outliers in Figure 3.2. The reason is that, especially in case of the small malware files (<1kB), there are some special samples, where only a few added bytes can cause a large TLSH difference (>80).


---

<sup>2</sup>Measure of disorder or uncertainty in the byte distribution of the given binary file.

$H = -\sum_{i=0}^{255} \frac{n_i}{N} \cdot \log_2\left(\frac{n_i}{N}\right)$ , where  $N$  is the length of file in bytes and  $n_i$  is the number of occurrences of byte value  $i$ .

0 is the minimum value of byte entropy, this occurs when all bytes in the binary have the same value.

8 is the maximum value of byte entropy, this occurs when the byte values are distributed uniformly at random.



**Figure 3.2:** TLSH difference between the original sample and the adversarial example created from it by strategy Chunker as a function of the amount of bytes added.

### 3.2.3 Strategy 2: Disguiser

The false positive rate of a good malware detection system is as low as possible. This means that only in very few cases a benign file will be classified as a malware. Our idea is that we could mislead the detection system, if we hide a malware inside of a benign file. So, the Disguiser strategy concatenates a benign file to the end of a malware binary. Hence when this file is executed, the malware will run, although the TLSH hash of the file may be determined by the added benign content.

Here, our intuition is that a small malware in a large benign file can be hidden easier, because the TLSH difference between the benign content and the adversarial example will be small. So our goal is to maintain this TLSH difference under the explained threshold of 40.

To find out what is the sufficient size ratio between the benign and the adversarial example to remain under the threshold of 40, we take a little empirical study. We select randomly 100 malware and 300 benign files. We concatenate each benign to each malware, and we calculate the TLSH difference between the hosting benign file and the crafted adversarial example. The result of this measurement is shown on Figure 3.3. On the left side of the Figure 3.3 are represented all adversarial examples. On the right side of the Figure 3.3 (which is a zoomed-in version of the left figure) are represented only those adversarial examples that have a TLSH distance less than 50 from the hosting benign file. We can observe (on the left side of the figure) that in general the higher the size ratio is, the higher the TLSH difference is. We can also observe (on the right side of the figure) that there are quite a few points in the area where the TLSH difference is below 40 and the size ratio is

above 1.2. Practically, this means that only the pairs with size ratio below 1.2 have the chance to have TLSH difference below 40.

---

**Algorithm 2** *Disguiser*

---

**Input:** malware binary  $MW$ , pool of benign files  $BN\_POOL$

**Output:** adversarial examples  $ADVEX\_POOL$

```

1: for all  $BN \in BN\_POOL$  do
2:   if  $sizeof(MW)/sizeof(BN) < 0.2$  then
3:      $ADVEX \leftarrow MW + BN$ 
4:     if  $TLSH\_difference(BN, ADVEX) < 40$  then
5:       add  $ADVEX$  to  $ADVEX\_POOL$ 
6:     end if
7:   end if
8: end for
9: return  $ADVEX\_POOL$ 

```

---

The pseudocode of Disguiser is presented as Algorithm 2, whose steps are as follows.

1. Select an arbitrary malware binary  $M$ .
2. Search a benign file  $B$ , so that the size ratio of  $(M+B)$  and  $B$  is below 1.2.
3. Concatenate  $B$  to the end of  $M$ .
4. Calculate the TLSH difference between  $(M+B)$  and  $B$ .
5. If the TLSH difference is below the required threshold we got an adversarial example.
6. Continue with Step 2, as long as there are still unscanned benign files.

### 3.3 Measurement


In this section we present in detail the setup and results of the measurement of the robustness against the crafted adversarial examples in case of SIMBIoTA and SIMBIoTA-ML. Before we look at the results of SIMBIoTA and SIMBIoTA-ML against the set of crafted adversarial examples, we present the background of our experiment. We say a few words about the used malware and benign dataset, we give a more detailed description about the specific parameters, and we show the selection of base material for adversary examples. Finally, we dive into the measurement methodology.

#### 3.3.1 Dataset

We perform all experiments using the same data set as used for the evaluation of SIMBIoTA and SIMBIoTA-ML in [28]. This dataset is called CrySyS-Ukatemi benchmark dataset of IoT malware 2021 (or CUBE-MALIoT-2021 for short). The dataset consists of 29,209 malicious ARM samples and 18,715 malicious MIPS samples, extended with 4,727 benign ARM samples and 9,392 benign MIPS samples. For malicious samples, metadata is also available, which details, among others, the date the sample was first seen in the wild (i.e., submitted to VirusTotal). CUBE-MALIoT-2021 is publicly available<sup>3</sup> for use by the IoT malware research community.

---

<sup>3</sup><https://github.com/CrySyS/cube-maliot-2021> (accessed: October 16, 2023)



**Figure 3.3:** Illustration of the effect of the size ratio between the adversarial example and the hosting benign file on the TLSH difference between them when strategy Disguiser is used. The x-axis shows the ratio of the sizes and the y-axis shows the TLSH difference.

### 3.3.2 Setup

In order to get a more accurate picture of the SIMBIoTAs’ robustness against adversarial examples we worked out a quite sophisticated parameter set for testing. We divided the malware sample dataset into three equal parts by size (Small-Medium-Large). The exact intervals are given in Table 3.1.

Disguiser and Chunker have further special parameters. Chunker creates adversarial examples with TLSH difference 40 and 60 from the original malware. Based on the information of Figure 3.2, practically to reach the threshold difference 40, 4 chunks are needed to be added, and to reach the threshold difference 60, 6 chunks are needed to be added. As explained in Section 3.2.3, Disguiser always uses 0.2 as the maximum of the  $M/B$  size

Architecture	S	M	L
ARM	1 - 59,900	59,901 - 120,875	120,876 - 1,942,729
MIPS	1 - 71,508	71,509 - 104,712	104,713 - 2,423,149

**Table 3.1:** Maximum and minimum limits of small (S), medium (M), and large (L) size intervals of malware dataset (in bytes), in the ARM and MIPS cases.

ARM			
Target TLSH diff.	S	M	L
40	3823	3619	3626
60	3647	3380	3392
MIPS			
Target TLSH diff.	S	M	L
40	3708	3265	2935
60	3288	3065	2593

**Table 3.2:** Number of adversarial examples created with strategy Chunker from the set of small (S), medium (M), and large (L) samples, with target TLSH difference values 40 and 60, in the ARM and MIPS cases.

ARM			
Target TLSH diff.	S	M	L
40	3429	1565	512
30	3868	1298	390
MIPS			
Target TLSH diff.	S	M	L
40	5046	3720	736
30	5055	2370	829

**Table 3.3:** Number of adversarial examples created with strategy Disguiser from the set of small (S), medium (M), and large (L) samples, with target TLSH difference values 30 and 40, in the ARM and MIPS cases.

ratio (this 20% size increase is still acceptable). Moreover, Disguiser creates adversarial examples with TLSH difference of maximum 40 and 30 from the hosting benign file.

The size of the sample data set is relatively large. So, we randomly select a subset as base material for adversarial examples. For Chunker we select 4000 samples from each size category. When creating adversarial examples from these samples by using the Chunker strategy with 4 or 6 added chunks, the criteria described previously in this subsection are not met by all of these selected samples: some of the resulting samples were not far enough in TLSH difference from the original sample. We ignored these samples, and kept only those that satisfy our constraints. The exact numbers of samples obtained in this way are shown in Table 3.2. The data in Table 3.2 indicates that it is more difficult to create adversarial examples from larger malware samples than it is from smaller ones. Furthermore, it is more difficult to reach TLSH threshold 60 than it is to reach TLSH threshold 40. These observations are consistent with intuition.

In case of Disguiser we randomly select 100 malware from each size category and we pair them with 300 randomly selected benign files. Table 3.3 shows how many pairs meet the criteria defined previously in this subsection. The data of Table 3.3 shows that it is more difficult to hide a larger malware into a benign file than a smaller one<sup>4</sup>.

In order to measure the robustness of SIMBIO TA and SIMBIO TA-ML against the adversarial examples, we need to simulate their behavior. Therefore, we train both SIMBIO TA

<sup>4</sup>We could use a more efficient algorithm for pairing malware files with benign files. However, in this study we rather concentrate on the robustness of the detection system against the adversarial examples. For now we showed that is also possible to create sufficient amount of adversarial examples with this simple method.



and SIMBIO-TA-ML on 10% of the dataset introduced in Subsection 3.3.1. When we have the trained SIMBIO-TA and SIMBIO-TA-ML, we can give them an adversarial example, and we can observe whether it is detected as a malware or not. Systematically, we give our adversarial examples (grouped by their parameters) to the detection systems, and we measure their detection accuracy. Similar to the experiment in [28], we repeated all measurements 12 times to eliminate the effects of randomly splitting the dataset into a 10% size training and 90% size testing part. In the next section we will see the results.

### 3.3.3 Results

We arrived to the presentation of measurement results. As described in previous sections, we prepared our adversarial examples. Now we see how SIMBIO-TA and SIMBIO-TA-ML react to these malicious files. Similar to the performance evaluation, we repeat the whole experiment 12 times, and we show the box plot of the accuracy results obtained in the 12 runs in Figures 3.4 and 3.5 for strategy Chunker and strategy Disguiser, respectively. Some of the results are according to expectations, some of them are somewhat surprising.


Firstly, we measure the robustness of SIMBIO-TA and SIMBIO-TA-ML against adversarial examples created by strategy Chunker with the parameter set described in Subsection 3.3.2. In case of Chunker our intuition could be that SIMBIO-TA-ML is more robust than SIMBIO-TA. This is true, as Figure 3.4 shows, SIMBIO-TA-ML has higher detection rate in all cases than that of SIMBIO-TA. Moreover, this higher accuracy is actually close to 1 for medium and large size samples, while we can observe a much lower accuracy (but still higher than SIMBIO-TA's) for small samples. It is clear that SIMBIO-TA can be misled with this strategy, because the classifier system of SIMBIO-TA operates directly with TLSH differences, and Chunker takes advantage of it. It seems that such a big TLSH difference (40-60) does not really matter for SIMBIO-TA-ML. Furthermore, in case of SIMBIO-TA it is also true that larger TLSH difference causes lower accuracy, as expected.

Secondly, we measure the robustness of SIMBIO-TA and SIMBIO-TA-ML against adversarial examples created by strategy Disguiser with the parameter set described in Subsection 3.3.2. The robustness of the two systems against adversarial examples of Disguiser are surprisingly poor. Basically, these adversarial examples are constructed by concatenating a malware and benign file in such a way that the size of the benign part is much larger than the size of malware part. Thus, the TLSH values of these adversarial examples are more similar to the TLSH values of benign files than to the TLSH values of malware samples. Therefore, SIMBIO-TA and SIMBIO-TA-ML misclassify these adversarial examples as a benign file. As Figure 3.5 shows, the accuracy of SIMBIO-TA-ML is not exactly 0. This may be because in case of SIMBIO-TA-ML there is a small false positive detection rate (see Subsection 2.4.2), where benign files are detected as malware.

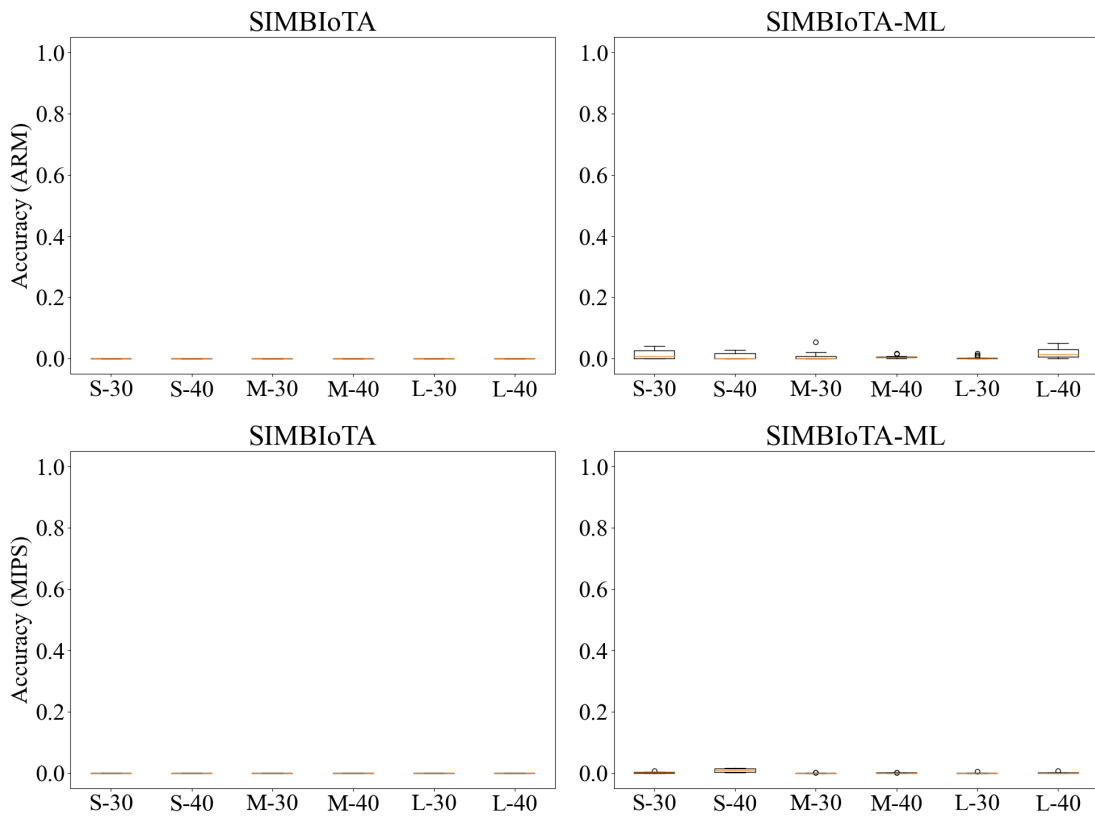
### 3.3.4 Discussion

In this section we have seen the results of the two strategies. We can state that SIMBIO-TA-ML is more robust against adversarial examples of Chunker than SIMBIO-TA is. Unfortunately, the robustness of SIMBIO-TA-ML is not preserved at all against strategy Disguiser. Moreover, in the perspective of these measurements, there does not seem to be any significant difference in the results in the ARM and the MIPS cases.

As in engineering work in general, these presented solutions also apply some trade-offs, e.g., Chunker can create relatively more adversarial examples than Disguiser, but the adversarial examples of Disguiser are more successful at evading the detection of SIMBIO-TA



**Figure 3.4:** Comparison of true positive detection rate of SIMBioTA and SIMBioTA-ML against strategy **Chunker**, in the ARM and MIPS cases.



**Figure 3.5:** Comparison of true positive detection rate of SIMBIO TA and SIMBIO TA-ML against strategy **Disguiser**, in the ARM and MIPS cases.

and SIMBioTA-ML. Furthermore, the Chunker strategy is a trade-off in itself. Chunker appends some bytes to the end of an existing malware, and we want to produce the added content relatively fast, on the other hand, we want the added content to look like some meaningful program code. We consider two approaches: Firstly, adding some constant or random bytes would be simple and fast, but easy to spot with simple static analysis; Secondly, if we add some bytes that have exactly the same byte distribution as the original malware, then it would be hard to spot with simple static analysis, but the generation of these bytes is too slow. The Chunker strategy adds some chunks from the original malware to itself, so it is a trade-off between the two previous alternatives, because in case of Chunker, the appended content looks like the binary of some meaningful program code, furthermore, we can add chunks faster than generating bytes with a specific byte distribution.

## Chapter 4

# Adversarial training on SIMBIO-TA-ML

Unfortunately, no matter how good a malware detection system is, attackers constantly work on methods to evade their detection, this is a cat-and-mouse game. Attackers have many advantages over antivirus companies, one of these is that attackers need only one successful adversarial example construction strategy to reach their goal, but antivirus companies should prepare for all possible adversarial strategies. In this chapter we present a possible solution that antivirus companies could use to increase the robustness of existing malware detection systems against adversarial examples.

In the previous chapter, we saw two possible methods that attackers could use to create adversarial examples that evade detection of SIMBIO-TA and SIMBIO-TA-ML. We showed by measurements that SIMBIO-TA-ML is robust against the Chunker strategy, but it can be misled by the Disguiser strategy, while SIMBIO-TA has poor robustness against both strategies. To overcome this problem, we propose to antivirus companies that they use SIMBIO-TA-ML with adversarial training.

Adversarial training has been used in the image recognition domain to increase the robustness of machine learning-based models against adversarial examples. We adopt this approach in the domain of malware detection and demonstrate its effectiveness. Adversarial training in our case means that the training set of the malware detector algorithm is extended with samples that are crafted by using the adversarial evasion strategies that we proposed.

We apply adversarial training only on SIMBIO-TA-ML, because based on the measurements so far, SIMBIO-TA-ML was more robust against the created adversarial examples than SIMBIO-TA. So it seems like a reasonable decision to improve the system, which is inherently better in terms of robustness.

### 4.1 Setup

As a first step for adversarial training on SIMBIO-TA-ML, we split the malware samples into a 10% train set and a 90% test set. To do this, we use K-folds cross-validation [30], which is a reliable and frequently used model checking technique. We use K-folds with 10 folds and repeat each measurement 10 times, where the samples of each fold belong to the train set once, and the test set consists of the samples of the other 9 folds. This ensures that each malware appears exactly once in the train set and 9 times in the test set.

To use adversarial training on SIMBIO-TA-ML we have to extend the original training sample set with adversarial examples. Originally, SIMBIO-TA-ML is trained on 10% of the malware dataset introduced in Subsection 3.3.1. The samples in the training set represent malware samples known to the antivirus company. Therefore, we construct adversarial examples for adversarial training from malware samples only from the training set, because the antivirus company has knowledge only about these files. After training SIMBIO-TA-ML on this extended set of training samples, we test its performance on the original test set and adversarial examples generated from the test set. In the following, SIMBIO-TA-ML is referred to as the updated SIMBIO-TA-ML after adversarial training and the original SIMBIO-TA-ML before adversarial training. Furthermore, we apply adversarial training separately in case of the Chunker and Disguiser strategies.

For adversarial training we have to determine how many adversarial examples should be included in the training set. In case of Chunker this is somewhat simpler than in case of Disguiser, because the Chunker strategy creates one adversarial example from one malware<sup>1</sup>. Therefore, in case of Chunker we use for training all adversarial examples created from malware files from the training set. While in case of Disguiser, the standard deviation of the number of adversarial examples created from a single malware is much larger, because the Disguiser strategy pairs one malware with all possible benign files and selects the pairs that meet the constraints described in Subsection 3.2.3.

To overcome this problem we created the LooseDisguiser strategy that is similar to Disguiser. The LooseDisguiser strategy, like the Disguiser strategy, pairs benign files with malicious files and creates an adversarial example from a pair if the ratio of the size of the malicious file to the size of the benign file is below 0.2. Unlike the Disguiser strategy, the LooseDisguiser strategy does not consider the TLSH distance between the hosting benign file and the constructed adversarial example. The LooseDisguiser strategy has a so-called multiply factor parameter (instead of TLSH threshold) that define the maximum number of adversarial examples created from a malware. With LooseDisguiser we can create constant<sup>2</sup> number of adversarial examples per malware.

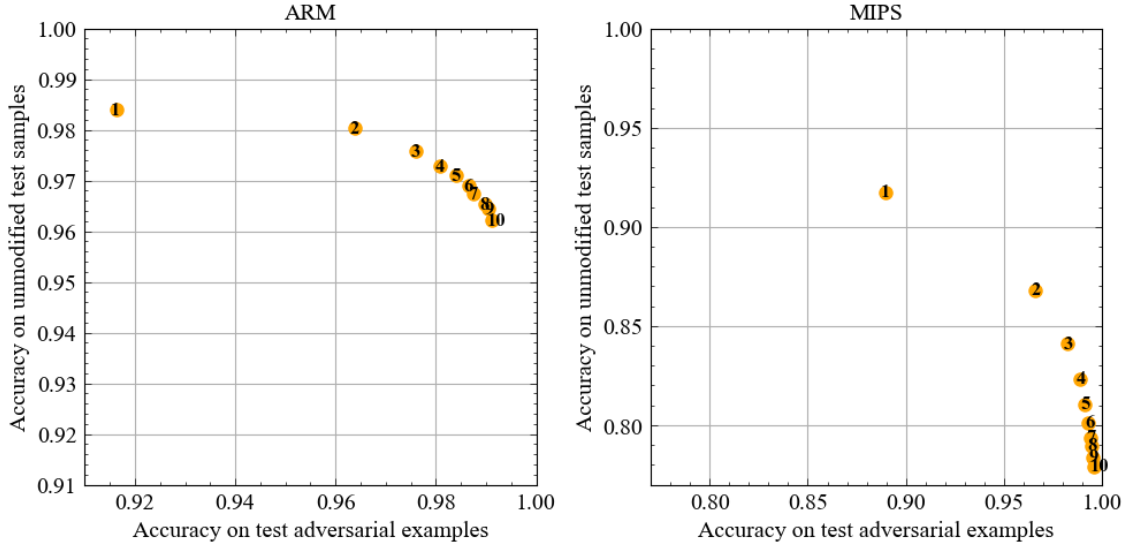
Depending on how many adversarial examples are added to the training set, the accuracy of SIMBIO-TA-ML changes on the test adversarial example set and the original test set. In case of Chunker we use for training all adversarial examples created from malware files from the training set. In case of LooseDisguiser, we measure the accuracy of the updated SIMBIO-TA-ML on the test adversarial example set and on the original test set with different multiply factors. In Figure 4.1 we see the results of this measurement. In this figure the ideal point is (1,1) which means 100% accuracy on adversarial example test set and 100% on the original test set. In case of ARM samples, the point corresponding to multiply factor 4 is the closest to this ideal point, however, in case of MIPS samples this multiply factor is 2. Seemingly, SIMBIO-TA-ML is more sensitive to noise (i.e., adversarial examples in the training set) in case of MIPS samples. This phenomenon requires further investigation and may lead to further research directions. To keep it simple, we choose multiply factor 4 for LooseDisguiser in case of both architectures<sup>3</sup>.

---

<sup>1</sup>In a small number of cases it occurs that the sample created with Chunker strategy does not reach the required TLSH distance from the original malware. In such a case, from this original malware the Chunker strategy cannot create an adversarial example.

<sup>2</sup>Or close to constant, because e.g., at multiply factor 10, we may not find 10 benign files that are five times the size of a very large malware.

<sup>3</sup>A different multiply factor can be chosen depending on which is more important: higher accuracy on the test adversarial example set or higher accuracy on the original test set. In addition, different multiply factors can be selected even for the ARM and MIPS cases.



**Figure 4.1:** The effect of the multiply factor of the LooseDisguiser strategy on the accuracy of SIMBIO TA-ML with adversarial training in the ARM (on the left) and MIPS (on the right) cases. The multiply factors are shown in the small rectangles, where each rectangle corresponds to a given combination of accuracy values.

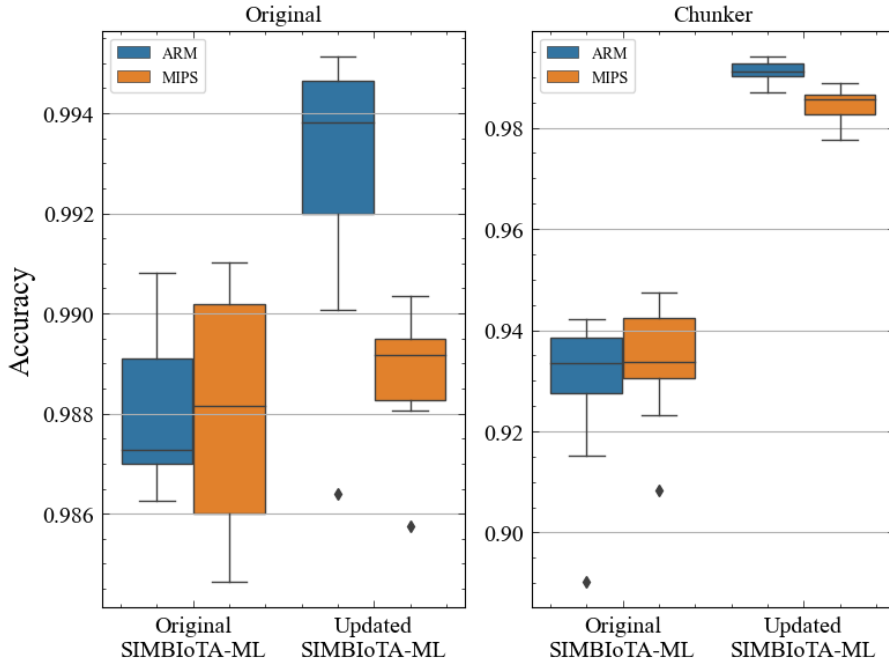
The exact numbers of samples obtained in this way are shown in Table 4.1. In case of Disguiser the train adversarial sample set is empty, because we use the adversarial examples of this strategy only for testing the original and the updated SIMBIO TA-ML. Furthermore, some cells of Table 4.1 contain intervals, rather than specific values, because the number of elements in the train and test sets may differ slightly in the 10 measurements.

## 4.2 Results

In this section we present the results of adversarial training on SIMBIO TA-ML. We measure the detection accuracy of SIMBIO TA-ML trained on the extended training set and show that it remains high both for the original malware samples and for the adversarial samples.

ARM			
Adversarial sample set	Chunker	LooseDisguiser	Disguiser
Training	2,684 - 2,727	11,663 - 11,674	–
Test	24,285 - 24,328	26,288 - 26,289	26,288 - 26,289
MIPS			
Adversarial sample set	Chunker	LooseDisguiser	Disguiser
Training	1,515 - 1,569	7,464 - 7,488	–
Test	13,862 - 13,916	16,813 - 16,819	16,813 - 16,819

**Table 4.1:** Number of elements in the train and test adversarial sample set constructed with the Chunker, LooseDisguiser, and Disguiser strategies, in the ARM and MIPS cases.



**Figure 4.2:** Comparison of the accuracy of the original and the updated SIMBIO-TA-ML on the original test sample set, and on the adversarial test sample set constructed with Chunker strategy, in the ARM and MIPS cases.

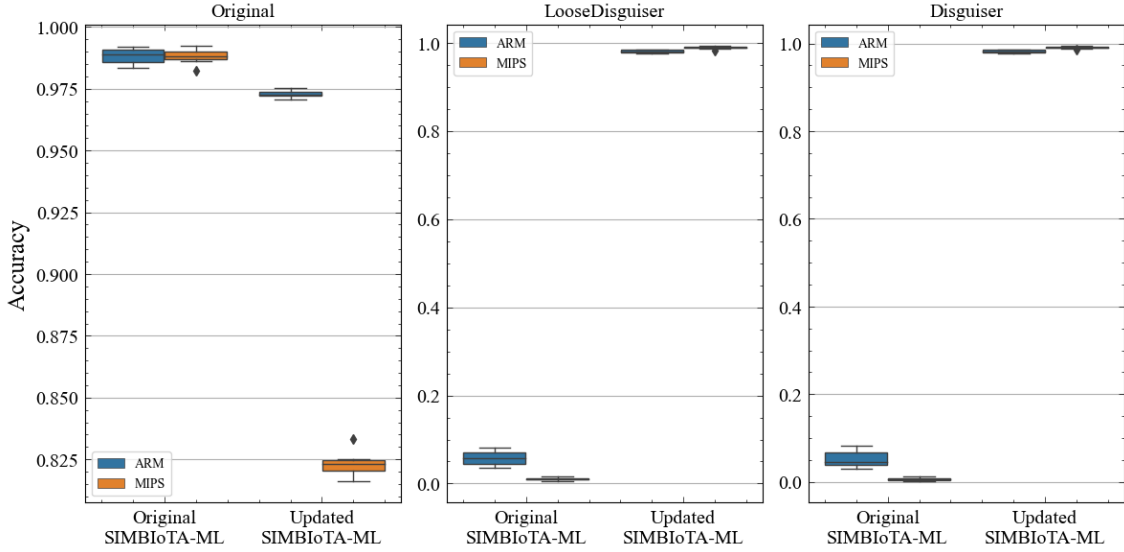
First, we present the results of adversarial training with samples created with the Chunker strategy. On the left side of Figure 4.2 we see that both the original and the updated SIMBIO-TA-ML have ca. 99% accuracy on the original test set. On the right side we notice that the test adversarial examples of Chunker somewhat mislead the original SIMBIO-TA-ML, its accuracy decreases to ca. 93%, while accuracy of the updated SIMBIO-TA-ML remains high at ca. 99%.

In Subsection 3.3.3 we showed that the original SIMBIO-TA-ML can be completely misled by the Disguiser strategy. In Figure 4.3 we see that the original SIMBIO-TA-ML can be completely misled by the LooseDisguiser strategy too. After adversarial training with the samples of LooseDisguiser, the updated SIMBIO-TA-ML has a significantly increased accuracy on the adversarial test set constructed with LooseDisguiser (ca. 97%). Moreover, the updated SIMBIO-TA-ML, which was trained with the adversarial examples of LooseDisguiser, remains surprisingly robust against adversarial examples of Disguiser too. While the accuracy of SIMBIO-TA-ML remarkably increases on adversarial examples after adversarial training, the accuracy of the updated SIMBIO-TA-ML on the original test sample set is only slightly lower than the original SIMBIO-TA-ML’s accuracy.

### 4.3 Discussion

In this chapter we showed that, by using adversarial training, antivirus companies can make SIMBIO-TA-ML more robust against previously presented adversarial evasion techniques. SIMBIO-TA-ML is not only a theoretical solution, but can also be used in industry. We strived for realistic adversarial training methodology that can be used in real-life situations. Therefore, we constructed adversarial examples for adversarial training from



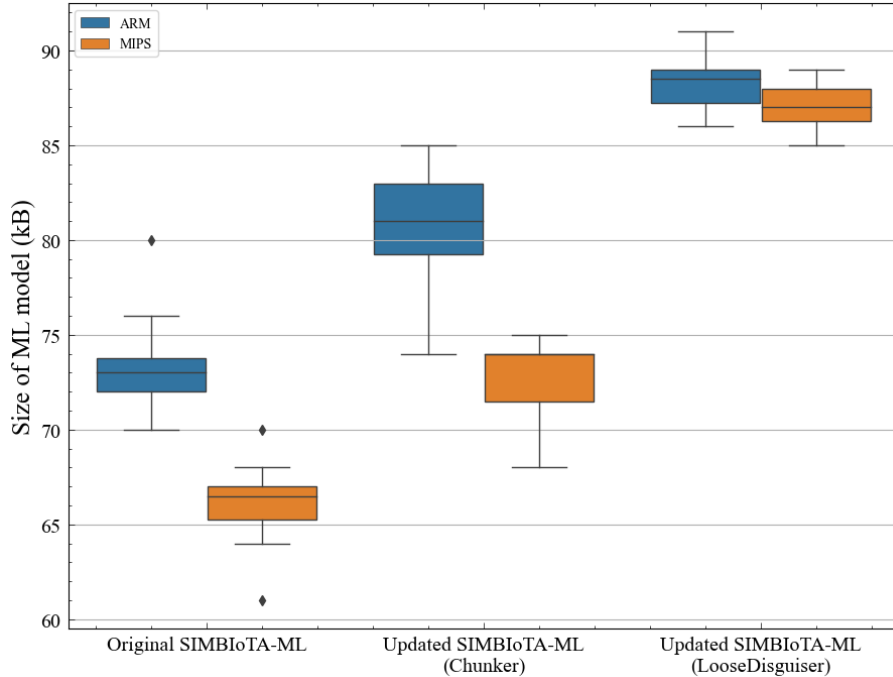


**Figure 4.3:** Comparison of the accuracy of the original and the updated SIMBIoTA-ML on the original test sample set, and on the adversarial test sample set constructed with LooseDisguiser strategy, and on the adversarial sample set constructed with Disguiser strategy, in the ARM and MIPS cases.

malware samples only from the original training set, because the antivirus company has knowledge only about these files. One may notice that in previous chapters we did not consider the test/train set when we selected malware samples for creating adversarial examples. This is because in the previous chapters we looked at the malware detection evasion scenario from the attacker’s perspective. The attacker does not know the malware database of the antivirus company, he can construct adversarial examples from all malware samples that he knows about.

The antivirus company does not necessarily know the exact algorithm of the attacker, in fact, most of the time this is the case. A good example of this is the presented scenario, where the antivirus company uses LooseDisguiser for training, but the attacker uses the more powerful Disguiser strategy. Nonetheless, the updated SIMBIoTA-ML, which was trained with the adversarial examples of LooseDisguiser, is surprisingly robust against adversarial examples of Disguiser too.

The price that we have to pay for this remarkable robustness is the increased training time and the increased size of the detection model, however, we argue that both are bearable in practice. In Section 4.1 we saw that adversarial training requires an extended training sample set. In machine learning, usually, an increased training set comes with increased training time and increased model size. This time, the increased training time is not critical, because in a real-life situation, similar to the original training of SIMBIoTA-ML, the adversarial training would be performed on the backend (see Subsection 2.3) and we can assume that the backend has practically unlimited resources compared to the IoT devices. Regarding the model size, in Figure 4.4 we indeed observe an increase of 10%, in case of Chunker, and 20%, in case of Disguiser, due to the 1.5 and 3 times, respectively, increase of the training set size. This translates to a few kilobytes of extra memory needed, which we believe to be still acceptable even on the resource constrained IoT devices.



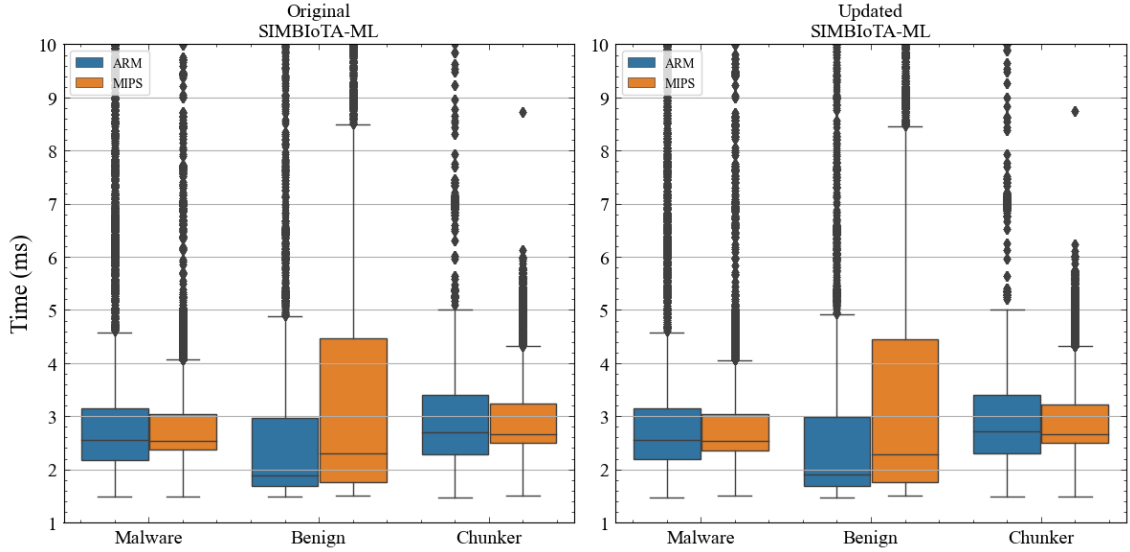
**Figure 4.4:** Comparison of the size of ML models trained without adversarial examples and trained with adversarial examples, in the ARM and MIPS cases.

In Subsection 2.3.1 we mentioned that SIMBioTA-ML uses a random forest classifier [9], which also need to be configured. Specifically, the number of decision trees that make up the random forest has to be specified. Similar to the original SIMBioTA-ML measurements in [28], we set the number of decision trees to 10, which give a good trade-off between the detection capability of the machine learning model and the memory required to apply the model on the embedded IoT device. Another important random forest parameter is the maximum depth of the decision trees. In the case of the original SIMBioTA-ML measurements in [28], the maximum depth of the decision trees of the random forest was not limited. However, the updated SIMBioTA-ML requires more memory on the embedded IoT device due to the extended training data set. Therefore, we set the maximum depth of random forest’s decision trees to 6, thus reducing the size of the model to half of the original in [28], while its detection capability practically remained the same.

For adversarial training of SIMBioTA-ML, similar to [28], our implementation for the random forest classifier uses the scikit-learn<sup>4</sup> Python module. In order to measure the amount of storage necessary to hold the model, similar to [28], we used the pickle<sup>5</sup> module to transform the Python object into a byte string that could be written to disk and later reloaded into memory. We then calculated the length of the byte string to get the number of bytes necessary to represent the object. We understand that there are more efficient ways to serialize a random forest model than using the pickle module. So, in practice, the representation of the model may require even smaller amount of memory on the IoT device than the amount we observed in our measurements (which is already acceptably small anyway).

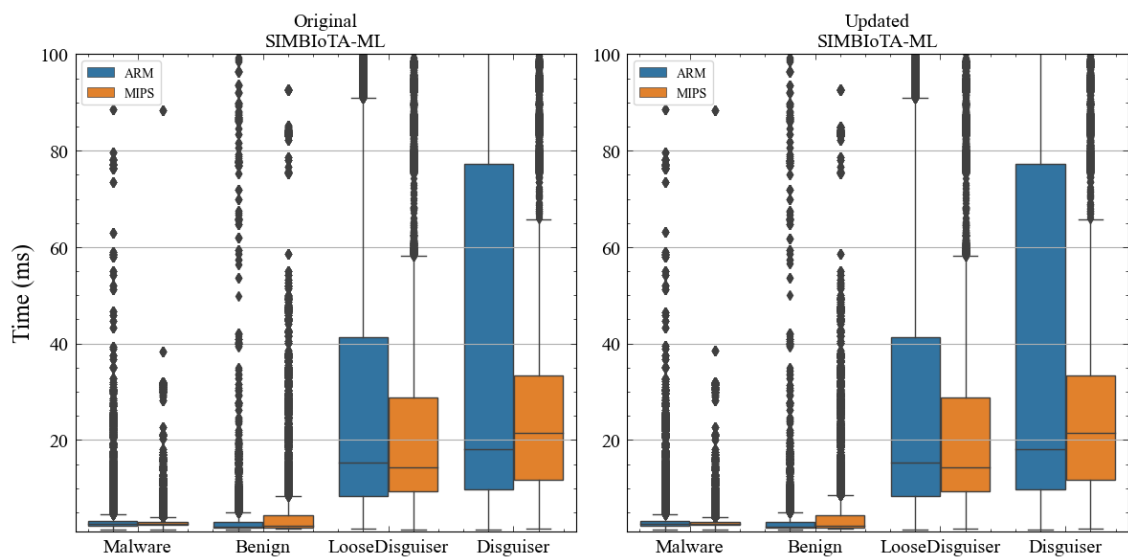
<sup>4</sup><https://scikit-learn.org/stable/> (accessed: October 16, 2023)

<sup>5</sup><https://docs.python.org/3/library/pickle.html> (accessed: October 16, 2023)



**Figure 4.5:** Comparison of the required detection times of SIMBIO TA-ML before and after adversarial training, separately for malware samples, benign samples, and adversarial samples of Chunker strategy, in the ARM and MIPS cases.

Finally, in Figure 4.5 and Figure 4.6 we measure the required processing time (i.e., detection time) of a given file on the IoT device. By detection time, we mean the time that elapses from the beginning of the binary scan of any file to the decision whether it is malicious or not. In the case of SIMBIO TA-ML, this time consists of the TLSH hash computation time of the binary and the decision time of the model. The decision time of the model is practically constant at 1.5 - 2 ms, due to the properties and operation of the random forest classifier [9]. However, the time that is required to compute a TLSH hash for a given binary depends strongly on the size of the binary itself [27]. Therefore, the total detection time is more dependent on the size of the processed file, which can be observed also in Figure 4.5, where samples of Chunker strategy requires more processing time than malware samples, because the size of an adversarial example created with the Chunker strategy is 20% larger than the original malware from which it was crafted. Moreover, this phenomenon is even more noticeable in Figure 4.6, in which case samples created with the Disguiser strategy require much more processing time than the malware samples, since the size of an adversarial example created with the Disguiser strategy is at least 6 times larger, than the original malware from which it was created. Nevertheless, there is no significant difference between the required detection times of the original and updated SIMBIO TA-ML for the malware, benign and adversarial samples, separately.



**Figure 4.6:** Comparison of the required detection times of SIMBioTA-ML before and after adversarial training, separately for malware samples, benign samples, and adversarial samples of LooseDisguiser and Disguiser strategy, in the ARM and MIPS cases.

## Chapter 5

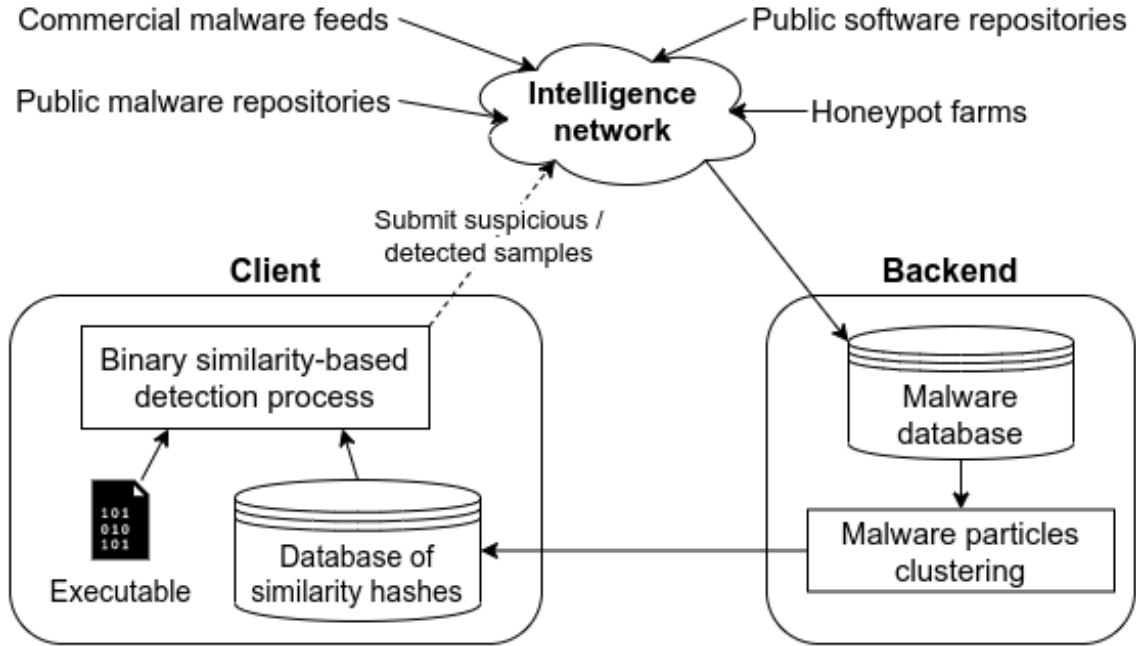
# PATRIoTA inspired by SIMBIoTA

In this chapter, we propose PATRIoTA (PArTicle TRained IoT Antivirus), a similarity-based IoT malware detection method inspired by SIMBIoTA, but being significantly more robust than SIMBIoTA is. The main idea of PATRIoTA is to split malware samples known to the antivirus provider into multiple fixed-size parts, referred to as *particles*, and to perform the same operations on those particles as the operations performed by SIMBIoTA on entire samples. This means that the antivirus provider builds a similarity graph of known particles, computes its dominating set, and distributes similarity preserving hash values (in our case, TLSH [27] values) corresponding to the particles in the dominating set to the clients. The clients also split any file to be scanned (e.g., a binary extracted from network traffic) into particles, compute the similarity preserving hash values of them, and if a threshold number of those computed hashes are similar to the hashes in the dominating set, then the file is detected as malware.

PATRIoTA is robust against adversarial sample creating strategies that add extra bytes to an existing malware binary, because the sample created in this way will always contain in it the original binary, and, hence, all of its particles, which can be recognized by PATRIoTA despite the presence of the added extra bytes. In addition, PATRIoTA may be robust against even more sophisticated adversarial strategies that keep sizable chunks of the original malware binary intact within the created adversarial sample, as those chunks may result in particles that are similar to the particles of the original sample. Moreover, our measurement results indicate that, besides increased robustness, PATRIoTA also has better malware detection capabilities than SIMBIoTA has.

### 5.1 Design

In Subsection 3.3.3, we mentioned that SIMBIoTA is not robust against the adversarial samples created with the Chunker and Disguiser strategies. Both attack strategies append some bytes at the end of an existing malware binary in such a way that those bytes are never executed, while the TLSH value of the modified sample becomes dissimilar to that of the original malware, and therefore, SIMBIoTA has a good chance of misclassifying it. In the case of these, and similar, append attacks, the original malware binary can be found in the adversarial sample. Hence, in order to detect such an adversarial sample as malware, we need a method that identifies the original malware inside the adversarial sample. PATRIoTA, the method we propose and describe in this section, will do exactly this: it identifies parts of known malware samples inside any file being checked with it.



**Figure 5.1:** High-level overview of PATRIoTA.

PATRIoTA can be viewed as a general method of defense against adversarial samples created with append attack strategies.

### 5.1.1 Overview

The design of PATRIoTA was inspired by SIMBioTA (and their similarity is also reflected in their names). Basically, PATRIoTA is a modified version of SIMBioTA where the difference is that PATRIoTA works with fixed-size parts of malware samples instead of entire malware binaries. Not surprisingly, the architecture of PATRIoTA is also almost the same as that of SIMBioTA, as it is illustrated in Figure 5.1. Malware samples are continuously collected from the intelligence network of the antivirus provider, and the PATRIoTA backend splits them into fixed-size parts, which we call *particles* in the sequel. Similar to SIMBioTA, a similarity graph is built by the backend, but in this case, this graph is built from the malware particles. In addition, PATRIoTA uses a different similarity threshold to build the similarity graph. In Subsection 5.1.4, we explain how to determine the optimal values for the particle size and the similarity threshold used by PATRIoTA. Again similarly to SIMBioTA, the backend computes a dominating set of the current similarity graph and makes the list of TLSH hash values of the dominating vertices available to clients.

The detection method on the client side is somewhat different in the case of PATRIoTA: the client splits the file to be checked into particles (of the same size used by the backend); calculates the TLSH hashes of the particles; and compares these TLSH hashes with those of the current dominating set. A file is considered malicious if it contains a threshold number of particles that are similar to known malicious particles. The selection of this threshold is discussed in Subsection 5.1.3.

### 5.1.2 Particle size and similarity threshold

PATRIoTA uses some special parameters, which we have already mentioned in the previous subsection, including the size of the particles and the similarity threshold used during the graph construction from the TLSH hashes. Finding the optimal configuration of these parameters is not a trivial task. We can state that the optimal configuration (if it exists at all) is highly context dependent; for example, we can imagine a situation where the low latency of the detection process is more critical than its memory usage.

Despite all this, we developed an iterative methodology to determine a recommended parameter configuration. We performed measurements to determine the optimal parameters on a smaller data set, not the one presented in Subsection 5.2.1. This data set consisted of 2000 malware and 2000 benign samples for both the ARM and the MIPS architectures.

When we were designing PATRIoTA, the first question was the size of the particles. We first considered the values of 1 kB, 2 kB, 4 kB, 8 kB, 12 kB and 16 kB, but later excluded 1 kB and 2 kB, because the number of graph nodes built from particles of those sizes grew unmanageably large.

PATRIoTA builds a graph from the TLSH values of malware particles, where the TLSH hash values are the nodes and there is an edge between two nodes if the TLSH difference of the two hash belonging to the nodes is below a certain similarity threshold value. SIMBIOta uses 40 as TLSH similarity threshold, because the average clustering coefficient of the built graph is the highest in that case [10]. The same value cannot be used for PATRIoTA, because it does not build the graph from the TLSH hashes of entire malware samples, but from its particles. To determine the optimal value of the TLSH similarity threshold for different particle sizes, we used the same technique as for SIMBIOta. In Figure 5.2, we measure the average clustering coefficient of the graph built from the particles of the 2000 malware samples using different TLSH similarity thresholds. We select the TLSH similarity threshold that gives the highest average clustering coefficient for each particle sizes (e.g., for particles of size 4 kB, the selected TLSH similarity threshold is 65 in the case of ARM samples).


### 5.1.3 Detection threshold

A suspicious sample is considered malicious if it contains at least a threshold number of particles that are similar to known malware particles. If this threshold is set to 1, the true positive detection rate (TPR) of malware will be as high as possible, but the unwanted effect may occur that even benign files are considered malicious (e.g., a benign and a malicious program use the same statically compiled library, therefore, both contain the same sequence of bytes). The consequence is that the larger the detection threshold is, the lower the false positive rate (FPR) and, unfortunately, the lower the TPR will be. So, we choose the smallest possible value where the FPR is still below 1%, which is 2 in the case of ARM samples and 4 in the case of MIPS samples (see Section 5.2).

### 5.1.4 Optimal configurations

At this point we have 4 possible particle size and TLSH similarity threshold configurations, but which of them is the best?

Before answering this question let's take a look at Figure 5.3, where we examine the number of similar particles between malware samples with the configuration of 4k particle

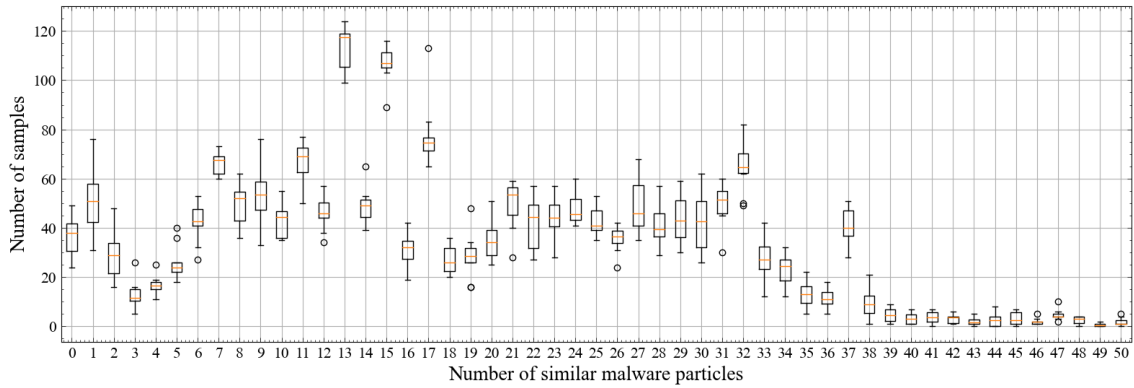


**Figure 5.2:** Average clustering coefficient as a function of the TLSH similarity threshold in the case of ARM (left) and MIPS (right) architectures. Different curves belong to different particle sizes, including the case where the particle size and the file size are equal (full).

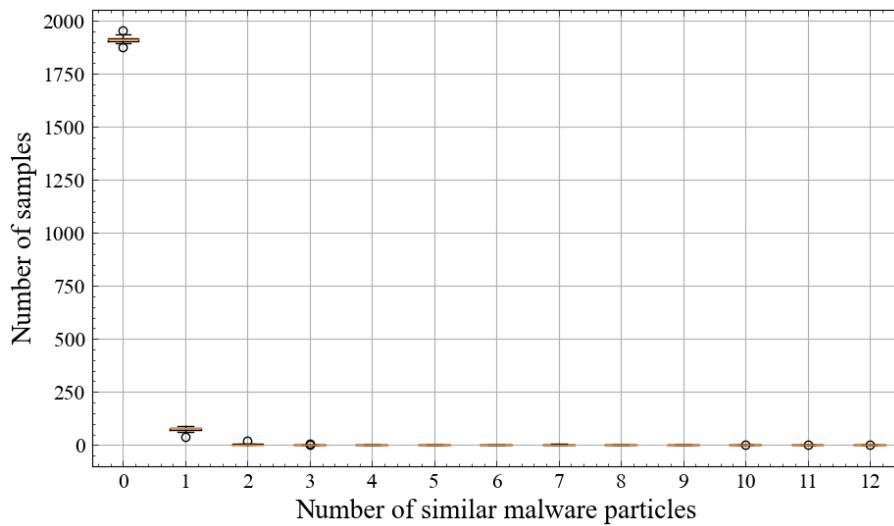
size and 65 similarity threshold for ARM samples. For that, we divide the 2000 malware samples to 10% train and 90% test set, we split to particles the items in the train set and we build the graph from their TLSH hashes. Finally, we iterate over the elements of the test set, split each file, and count how many similar particles there are between them and the particles in the graph. In other words, we simulate the operation of PATRIoTA on a small data set and repeat this simulation ten times (just like in Section 5.2, in the case of the large data set). There are ca. 40 samples in Figure 5.3 that do not contain any similar malware particles to the particles in the train set, therefore, we cannot detect these. Furthermore, in Figure 5.3 we see how many malware samples would not be detected depending on the selected detection threshold. For instance, if we required that at least 3 particles of the file should be similar to some known malware particle to detect the file as malicious (i.e., detection threshold 3), then ca.  $40+50+30=120$  malware samples would not be detected. In Figure 5.4 we examine the number of similar particles between the train malware samples and the test benign samples, we iterate all over the 2000 benign samples, split each file, and count how many similar particles there are between them and the particles in the graph built from the particles of the 200 malware train samples. This figure shows that the vast majority of benign samples do not contain particles that resemble malware particles, however there are a few samples that unfortunately do, which would cause the false positive decisions of PATRIoTA.

To select the best particle size and TLSH similarity threshold configuration, we examine how it changes the TPR and FPR values depending on the detection threshold. Figure 5.5 shows the ROC (Receiver Operating Characteristic) curves of the different configurations, where each jump in the step function corresponds to a certain detection threshold value between 1 and 10. According to our expectations, as the detection threshold increases, the FPR decreases, but so does the TPR. We choose the configuration with the highest AUC (Area Under the ROC Curve) value. This is 4k particle size and 65 TLSH similarity threshold for ARM samples and 8k particle size and 60 TLSH similarity threshold for MIPS samples.






**Figure 5.3:** The  $x$  axis shows the number of particles in a sample from the **malware** test set that are similar to items in the train set (from 0 to 50), while the  $y$  axis shows the number of these samples in the case of 4k particle size and 65 TLSH similarity threshold configuration for ARM samples.



**Figure 5.4:** The  $x$  axis shows the number of particles in a sample from the **benign** test set that are similar to items in the train set (from 0 to 12), while the  $y$  axis shows the number of these samples in the case of 4k particle size and 65 TLSH similarity threshold configuration for ARM samples.



**Figure 5.5:** ROC curves of different configurations, where each jump in the step function corresponds to a certain detection threshold value between 1 and 10 in the ARM and MIPS cases.

## 5.2 Evaluation

In Section 5.1, we presented PATRiOTA, including its architecture and operating principles, as well as the selection of its parameters (particle size, similarity threshold, and detection threshold). It is time to evaluate PATRiOTA’s performance, especially its ability to detect adversarial samples. However, before doing that, we present the data set and methodology used for the performance measurements.

### 5.2.1 Experiment design

In this work, we perform all experiments using the same dataset as used for the adversarial training of SIMBiOTA-ML in Chapter 4, which is described in Subsection 3.3.1. As a first step for testing PATRiOTA, we split the malware samples, into a 10% train set and a 90% test set. To do this, we use K-folds cross-validation [30], which is a reliable and frequently used model checking technique. We use K-folds with 10 folds and repeat each measurement 10 times, where the samples of each fold belong to the train set once, and the test set consists of the samples of the other 9 folds. This ensures that each malware appears exactly once in the train set and 9 times in the test set<sup>1</sup>.

PATRiOTA does not need benign samples for training, so we add benign samples only to the test set. Moreover, we extend the test set with adversarial samples for measuring the robustness of the system. These adversarial samples are created using the two strategies mentioned in Subsection 3.2: Chunker and Disguiser. We create these adversarial samples from the malware binaries in the test set, simulating that an attacker has malware samples unknown to the antivirus company and can create adversarial samples from them. Table 5.1 shows the exact number of samples in the train and test set.

<sup>1</sup>We use exactly the same train and test sets in case of PATRiOTA as used for adversarial training on SIMBiOTA-ML.

**Table 5.1:** Number of samples in the train and test set, in the ARM and MIPS cases.

<b>ARM</b>				
	Malware	Benign	Chunker	Disguiser
Train	2,921	–	–	–
Test	26,288	4,727	24,285	26,288
<b>MIPS</b>				
	Malware	Benign	Chunker	Disguiser
Train	1,872	–	–	–
Test	16,843	9,392	13,862 - 13,916	16,813 - 16,819

With the presented construction, we simulate the operation of PATRIoTA: we build the model properly from the train samples, and then give samples from the test set to the model for detection (see Section 5.1). Furthermore, since PATRIoTA was inspired by SIMBIOtA, we compare their performances in all aspects. Indeed, we train and test the two systems on the same samples and measure the same performance metrics. In the next 3 subsections, we present the results of the performed simulation.

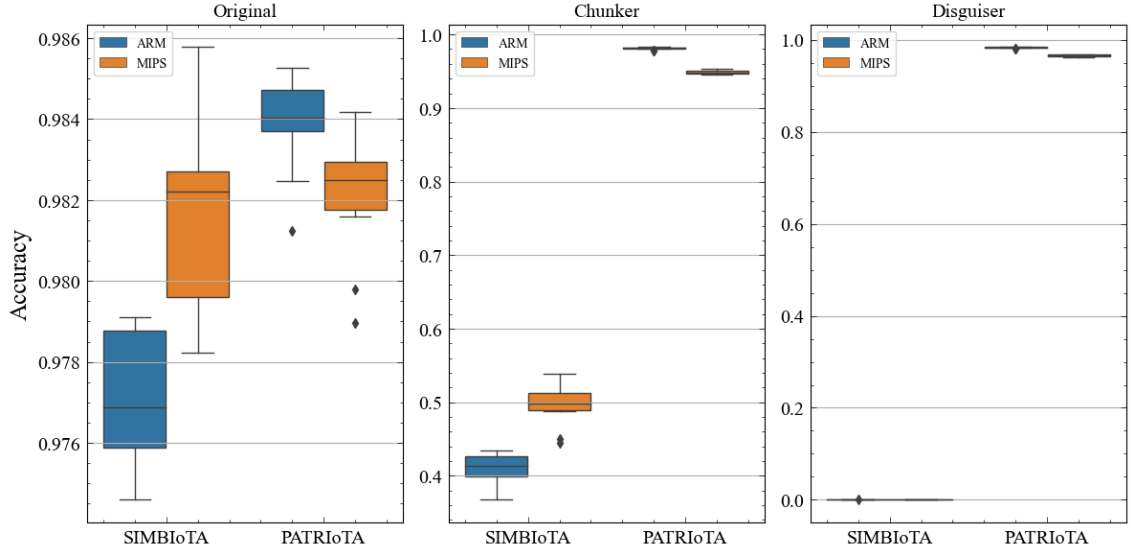
### 5.2.2 Detection capability

Using the experimental setup presented in the previous subsection, we measure the detection accuracy of SIMBIOtA and PATRIoTA in 3 different cases: on unmodified malware and benign files, on adversarial samples of the Chunker strategy, and on adversarial samples of the Disguiser strategy. Our goal is to achieve the highest possible accuracy in all 3 cases, while keeping the FPR of benign files below 1%. To do this, we try PATRIoTA with different detection thresholds (i.e., minimum number of particles of a file that need to be similar to known malware particles in order for the file to be classified as malware). The smaller the detection threshold is, the higher the TPR will be, but the FPR will increase too. According to our measurements, the optimal value for the detection threshold is 2 for ARM samples and 4 for MIPS samples, as for smaller values, the FPR exceeds 1%.

In Figure 5.6, we compare the detection accuracy of the two system. PATRIoTA drastically outperforms SIMBIOtA in terms of accuracy in all test cases! On the sample set of unmodified benign and malicious programs, PATRIoTA has an impressive 98.5% accuracy in the case of ARM samples and 98.2% in the case of MIPS samples. Moreover, it performs extremely well even on adversarial samples of both the Chunker and Disguiser strategies, with 98% accuracy on ARM samples and 95% on MIPS samples.

### 5.2.3 Storage requirement

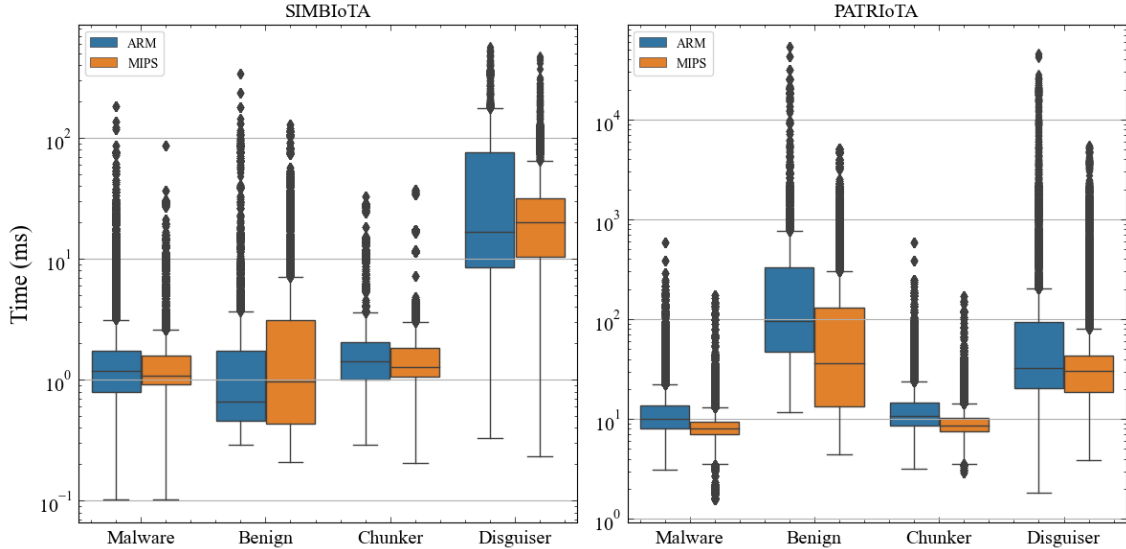
IoT devices are usually limited by resources, including available memory and storage capacity. Therefore, we measure the storage space requirement of PATRIoTA on the client side (i.e., on the IoT device), which in our case is the size of the dominating set multiplied by the size of the TLSh hash. Compared to SIMBIOtA, unfortunately, the higher accuracy and robustness of PATRIoTA comes with a higher storage requirement, due to the increased number of nodes in the dominating set. In Table 5.2, we present the required memory sizes of the two system.



**Figure 5.6:** Comparison of the detection accuracy of SIMBIOtA and PATRIoTA on unmodified malicious and benign samples (Original), adversarial samples created with the Chunker strategy, and adversarial samples created with the Disguiser strategy, in the ARM and MIPS cases.

**Table 5.2:** Storage requirement of SIMBIOtA and PATRIoTA on the client side in the ARM and MIPS cases.

	SIMBIOtA	PATRIoTA
ARM	8,365 - 9,030 B	333,585 - 371,805 B
MIPS	5,775 - 6,440 B	111,965 - 139,370 B



**Figure 5.7:** Comparison of the required detection times (on logarithmic y-axis scale) of SIMBIOtA and PATRIoTA, separately for malware samples, benign samples, adversarial samples of Chunker strategy, and adversarial samples of Disguiser strategy, in the ARM and MIPS cases.

#### 5.2.4 Run time performance

Another price we have to pay for the increased detection accuracy and robustness of PATRIoTA is the increased processing time compared to SIMBIOtA. By detection time, we mean the time that elapses from the beginning of the binary scan of any file to the decision whether it is malicious or not. In the case of SIMBIOtA, this time consists of the TLSH hash computation time of the binary and the decision time of the model. For PATRIoTA, the detection time consists of the sum of 3 components: the time required to split the binary into fixed-size particles, the sum of TLSH hash computation time of the particles, and the sum of decision times required for particles. Basically, PATRIoTA performs SIMBIOtA’s detection method multiple times, more precisely for each particle, until the number of particles considered malicious reaches the value of the detection threshold parameter. Therefore, PATRIoTA requires more time for detection than SIMBIOtA, as shown in Figure 5.7.

### 5.3 Discussion

In this chapter, we were concerned with increasing the robustness of binary similarity-based malware detection methods against adversarial samples that are crafted specifically to mislead a given malware detector. More specifically, we proposed PATRIoTA, a robust, similarity-based antivirus solution, which was inspired by SIMBIOtA [37]. In Chapter 4 we studied the robustness of SIMBIOtA-ML, where an adversarial training approach was proposed as a solution. So a natural question, at this point, is whether an adversarial training approach could have increased the robustness of SIMBIOtA as well. If so, then the need for a new approach, i.e., our PATRIoTA, would be much weaker.

	Similarity Graph	Machine Learning
Particle Searching	✓	?
Adversarial Training	✗	✓

**Table 5.3:** Evaluation of the combinations of the detection models with the proposed robustness enhancing techniques. The question mark represents that the combination of machine learning with particle searching, i.e., PATRIoTA-ML, requires further investigation.

It is clear what adversarial training means in case of a machine learning-based method: the training set is expanded with adversarial samples created by various known adversarial strategies. But SIMBIOta is not a machine learning-based method. Nevertheless, we can define adversarial training quite intuitively for SIMBIOta too: the antivirus provider extends the similarity graph of known malware samples with adversarial samples and computes the dominating set of this extended graph. One can then check the detection performance of this modified SIMBIOta on adversarial samples to determine how robust this approach is.

We performed adversarial training of SIMBIOta by extending the similarity graph of known malware samples with adversarial samples created from those known malware by the Chunker and Disguiser strategies introduced in Section 3.2, and computed the dominating set of the extended graph. We then measured the detection performance on adversarial samples created from malware unknown to the antivirus provider by the same Chunker and Disguiser strategies. The results we got were not so promising: adversarial ARM and MIPS samples created by the Chunker strategy were detected with 92% and 90% accuracy, respectively, while adversarial ARM and MIPS samples created by the Disguiser strategy were detected only with 86% and 67% accuracy, respectively. These results confirm the *raison d’être* of PATRIoTA.

In Table 5.3, we provide a summary of the discussion presented in this section. We list the types of our detection models (similarity graph and machine learning) and the proposed techniques designed to enhance adversarial robustness (particle searching and adversarial training). We placed a question mark in the cell that combines machine learning with particle searching, i.e., PATRIoTA-ML. This approach offers several advantages: Proper hyperparameter settings of the Random Forest Classifier allow us to regulate the model size; The Random Forest Classifier ensures constant and short decision times; The particle searching concept contributes to robustness against adversarial examples. Unfortunately, based on our preliminary measurements, the training time and model size grow unmanageably large, and the detection capability is not sufficient either. We believe that this might be due to excessive noise in the training set, where common particles could exist in both malware and benign files. Therefore, some form of prefiltering or preprocessing of the particles in the training set would be necessary. Nevertheless, PATRIoTA-ML does not show very good performance yet. However, we strongly believe that it has great potential and is worth further research.

In addition, while PATRIoTA was designed to be robust against adversarial samples that were created from existing malware samples by appending extra bytes to them, we have the intuition that it is also robust against other strategies that create adversarial samples that contain chunks of the original sample, as those chunks may result in particles that are similar to the particles of the original sample. In order to test this intuition, we measured the robustness of PATRIoTA against such a strategy. In particular, a very clever adversarial sample creation strategy against similarity-based malware detection was

proposed in [13] that consists in modifying a few unused portions of a malware binary (e.g., the section header tables were modified in [13]) such that the TLSH difference between the modified and the original files is maximized, while the functionality of the original binary is fully preserved, the size of the modified file remains the same as that of the original one, and even the binary content is only slightly changed. As reported in [13], it is rather easy to create adversarial samples in this way that are misclassified by SIMBIO<sub>TA</sub>: out of 2000 randomly chosen ARM malware samples, 1779 samples were suitable for such kind of modification, and 1465 samples could be created with a TLSH difference of at least 40 (the threshold used by SIMBIO<sub>TA</sub>) between the modified and original files. We tested both SIMBIO<sub>TA</sub> and PATRIO<sub>TA</sub> with those samples, and SIMBIO<sub>TA</sub> recognized only 17% of them as malware, while PATRIO<sub>TA</sub> detected a remarkable 98% of them as malware!

One may wonder whether statically linked libraries decrease the detection accuracy of PATRIO<sub>TA</sub>. Such libraries may be included in both malware and benign binaries, so actually, some portions of statically linked malware and benign samples that include the same libraries can be identical. This may lead to multiple similar particles in them, potentially above the threshold number used by PATRIO<sub>TA</sub>. In other words, benign files may contain particles resulting from linked libraries that are similar to particles seen in malware binaries using the same libraries. Such benign files may be classified as malware by PATRIO<sub>TA</sub>, which leads to an increased false positive rate. Indeed, we tested PATRIO<sub>TA</sub> on 118 statically linked ARM and 64 statically linked MIPS benign binaries and it misclassified 15% and 6% of them, respectively, as malware. This misclassification rate is not really acceptable, therefore, further research is needed to reduce it. We note that SIMBIO<sub>TA</sub> had a false positive detection rate of 0% throughout our experiments.

# Chapter 6

## Related work

In this chapter, we present a comprehensive overview of the state of the art in the related field. While SIMBIO TA-ML uses machine learning for malware detection, SIMBIO TA (and consequently PATRIO TA as well) does not belong to ML-based malware detectors. However, since adversarial examples and adversarial robustness, the main focus of this work, are closely tied to machine learning, we include an outlook on ML-based malware detection in this chapter.

### 6.1 ML-based (IoT) malware detection

As we mention in Chapter 1, traditional (i.e., signature-based and heuristic solutions) malware detection systems could have scalability problems. In addition, traditional systems use only static properties of malware files for detection, hence with special techniques (e.g., obfuscation) they can be deceived.

Unlike traditional solutions, ML-based malware detection can be highly automated [40, 38, 14]. Furthermore, they use static and dynamic program analysis for extracting the required feature vectors [32]. Hence, their detection capabilities are better than that of traditional malware detection approaches. Feature vectors can be extracted from different sources, including the samples' instructions [12, 36], their control-flow [2], invoked API functions and system calls [1, 31], grey scale images of binaries [22], strings [21], and messages sent over network [26, 16].

In addition, solutions that combine machine learning with cloud-based approach scale well and can be applied also in the IoT field [35, 20]. This construction is advantageous for resource constrained IoT devices, because resource heavy calculation and processing can be passed to cloud, and only a lightweight algorithm is needed on client side. They can use different ML models, including convolutional neural networks [33], recurrent neural networks [17], random forest classifiers [36], fuzzy and fast fuzzy pattern trees [12].

### 6.2 Adversarial examples and robustness analysis

For making machine learning models more accurate and reliable we have to prepare them against adversarial examples. Therefore, this is an actively researched area with a rich and diverse literature.



Originally, ML-based image recognition was the first scientific field where adversarial attacks were applied. As an example, consider Figure 3.1, where some perturbation is added to the original image for the purpose of deception of the given ML classifier. Practically, this concept can be applied to ML-based malware detection field too, because also in case of malware files we usually want some perturbation on the original binary. In addition, as the functionality of image remains the same (i.e., in Figure 3.1 we still recognize the panda), usually we want to preserve the functionality of malware too.

Based on the attacker’s knowledge on the targeted ML detection system we can distinguish two type of attacker models [6]. In the white-box model, the attacker has a comprehensive idea of the ML model, he knows the exact training data and concrete model parameters. Moreover, there exists the black-box model, when the attacker has knowledge only about the input and output of the model. Our presented strategies (Chunker & Disguiser) rather follow a grey-box attacker model, because they do not know about concrete model parameters, but they take into account the fact that SIMBIO TA and SIMBIO TA-ML use similarity based hashes. So in our case, the attacker has partial information about the detection systems.

There are many different approaches for adversarial attacks also in the context of malware detection [6]. From these approaches we can highlight *append* and *slack* attacks [34] for their simplicity. Append attacks generate bytes and add them to the end of malware binary. Slack attacks add or modify bytes in slack regions of a binary, which are gaps between neighboring sections of an executable file. Our presented strategies (Chunker & Disguiser) resemble the previously mentioned append attack. There are other solutions for generating and appending bytes to the end of a binary, including gradient-based approach [23, 24].

Another more advanced technique is program obfuscation, which can change the binary representation of a program while preserving its functionality [29]. In order to do so, ML solutions can be used, including reinforcement learning-based approaches [4, 3], Generative Adversarial Networks (GAN)[19] and Recurrent Neural Networks (RNN) [18]. Obfuscating existing malware samples may be a successful strategy, but we do not use it, because from the perspective of SIMBIO TA and SIMBIO TA-ML, obfuscated samples appear to be new malware, as their binary representations can be completely different from those of the original samples from which they were created. In other words, obfuscated samples are considered new malware by SIMBIO TA and SIMBIO TA-ML, and their detection performance on them has already been measured in [28].

Adversarial training is an effective way to increase the robustness of ML-based systems against adversarial examples, and it can also be applied in the malware detection domain. Several existing solutions use this technique to improve their malware detection systems [25, 39]; however, we applied it first in the domain of ML-based IoT malware detection.

## Chapter 7

# Conclusion

To summarize this thesis, we presented two recently proposed IoT malware detection solutions: SIMBIO TA and SIMBIO TA-ML. We developed two adversarial strategies, Chunker and Disguiser, capable of misleading these solutions. To overcome this problem, we proposed two solutions that enhance the robustness of the systems against the devised adversarial strategies: adversarial training and PATRIO TA.

We were interested in how SIMBIO TA and SIMBIO TA-ML, which perform well initially, behave against adversarial examples. Therefore, we constructed two different strategies for creating adversarial examples from existing malware files: Chunker and Disguiser. Basically, both strategies append a few bytes to the end of the malware, so they are relatively simple methods. Our measurement study shows that in case of Chunker, SIMBIO TA-ML has higher detection rate than SIMBIO TA, while in case of Disguiser, both detection system have poor performance.

To overcome this problem we proposed two solutions: the first one uses SIMBIO TA-ML, the second uses SIMBIO TA. Firstly, we used the adversarial training concept to increase the robustness of SIMBIO TA-ML against the presented adversarial evasion strategies. For adversarial training, we extended the training sample set of SIMBIO TA-ML with adversarial examples constructed by the adversarial evasion strategies. After adversarial training, the updated SIMBIO TA-ML became much more robust against samples of Chunker and Disguiser. The price that we have to pay for this remarkable robustness was the increased training time and the increased size of the detection model, however, we showed that both are bearable in practice. Secondly, we introduced PATRIO TA, an IoT malware detection method inspired by the same principles as SIMBIO TA. We demonstrated its exceptional malware detection capabilities, coupled with its resilience against various adversarial sample creation strategies. More specifically, we compared the performance and robustness of PATRIO TA to that of SIMBIO TA. PATRIO TA has a higher true positive detection rate, but it also has a higher false positive rate, it requires more storage capacity, and it has longer detection time than SIMBIO TA has. Its true advantage is its strong robustness against adversarial samples: indeed, SIMBIO TA can be completely misled by adversarial samples created from existing malware by appending extra bytes to them, whereas PATRIO TA detects those samples with very high accuracy.

PATRIO TA exhibits an increased running time in comparison to SIMBIO TA, particularly when assessing benign files. Moreover, PATRIO TA might yield false positive decisions on statically linked benign binaries if they incorporate libraries that have also been utilized in malware. These challenges warrant further investigation and are part of our future research agenda. Could PATRIO TA-ML potentially resolve these issues?

# Acknowledgements

The research presented in this document was supported by the European Union project RRF-2.3.1-21-2022-00004 within the framework of the Artificial Intelligence National Laboratory. The presented work also builds on results of the SETIT Project (2018-1.2.1-NKP-2018-00004), which was implemented with the support provided from the National Research, Development and Innovation Fund of Hungary, financed under the 2018-1.2.1-NKP funding scheme.

The author would also like to acknowledge the help received from the members of the CrySyS Lab, especially from Dr. Dorottya Futóné Papp, in the interpretation of existing implementations of SIMBioTA and SIMBioTA-ML, which were previously created in the context of the SETIT Project. Furthermore, the author is deeply grateful to his supervisors, Dr. Levente Buttyán and Roland Nagy. They could be counted on, were always very helpful, and provided the right advice during all those memorable meetings. Special thanks are due to Dr. Levente Buttyán, who introduced the author to the world of scientific research and set him on that path, moreover, who has been not only an excellent advisor but also a mentor and a role model for the author. Finally, the author wants to recognize and thank his family members for their continuous support.

# Bibliography

- [1] Muhamed Fauzi Bin Abbas and Thambipillai Srikanthan. Low-complexity signature-based malware detection for iot devices. In Lynn Batten, Dong Seong Kim, Xuyun Zhang, and Gang Li, editors, *Applications and Techniques in Information Security*, pages 181–189, Singapore, 2017. Springer Singapore. ISBN 978-981-10-5421-1.
- [2] Hisham Alasmay, Aminollah Khormali, Afsah Anwar, Jeman Park, Jinchun Choi, Ahmed Abusnaina, Amro Awad, Daehun Nyang, and Aziz Mohaisen. Analyzing and detecting emerging internet of things malware: A graph-based approach. *IEEE Internet of Things Journal*, 6(5):8977–8988, 2019.
- [3] H. Anderson, Anant Kharkar, Bobby Filar, David Evans, and Phil Roth. Learning to evade static pe machine learning malware models via reinforcement learning. *ArXiv*, abs/1801.08917, 2018.
- [4] Hyrum S. Anderson, Anant Kharkar, Bobby Filar, and Phil Roth. Evading machine learning malware detection. *BlackHat USA*, 2017.
- [5] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J. Alex Halderman, Luca Invernizzi, Michalis Kallitsis, Deepak Kumar, Chaz Lever, Zane Ma, Joshua Mason, Damian Menscher, Chad Seaman, Nick Sullivan, Kurt Thomas, and Yi Zhou. Understanding the mirai botnet. In *26th USENIX Security Symposium (USENIX Security 17)*, pages 1093–1110, Vancouver, BC, August 2017. USENIX Association. ISBN 978-1-931971-40-9.
- [6] Kshitiz Aryal, Maanak Gupta, and Mahmoud Abdelsalam. A survey on adversarial attacks for malware analysis. *ArXiv*, abs/2111.08223, 2021.
- [7] Ömer Aslan and Refik Samet. A comprehensive review on malware detection approaches. *IEEE Access*, 8:6249–6271, 2020.
- [8] Marco Barreno, Blaine Nelson, Anthony D. Joseph, and J. D. Tygar. The security of machine learning. *Machine Learning*, 81(2):121–148, Nov 2010. ISSN 1573-0565.
- [9] Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, Oct 2001. ISSN 1573-0565.
- [10] Levente Buttyán, Roland Nagy, and Dorottya Papp. SIMBioTA++: Improved Similarity-based IoT Malware Detection. In *2022 IEEE 2nd Conference on Information Technology and Data Science (CITDS)*, pages 51–56. IEEE, 2022.
- [11] Emanuele Cozzi, Pierre-Antoine Vervier, Matteo Dell’Amico, Yun Shen, Leyla Bilge, and Davide Balzarotti. The tangled genealogy of iot malware. In *Annual Computer Security Applications Conference, ACSAC ’20*, page 1–16, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450388580.

- [12] Ensieh Modiri Dovom, Amin Azmoodeh, Ali Dehghantanha, David Ellis Newton, Reza M. Parizi, and Hadis Karimipour. Fuzzy pattern tree for edge malware detection and categorization in iot. *Journal of Systems Architecture*, 97:1–7, 2019. ISSN 1383-7621.
- [13] Gábor Fuchs. Adversarial example creation strategies for defeating similarity-based IoT malware detection algorithms. BSc Thesis, Budapest University of Technology and Economics, 2023.
- [14] Daniel Gibert, Carles Mateu, and Jordi Planes. The rise of machine learning for detection and classification of malware: Research developments, trends and challenges. *Journal of Network and Computer Applications*, 153:102526, 2020. ISSN 1084-8045.
- [15] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples, 2014.
- [16] Mohit Goyal, Ipsit Sahoo, and G. Geethakumari. Http botnet detection in iot devices using network traffic analysis. In *2019 International Conference on Recent Advances in Energy-efficient Computing and Communication (ICRAECC)*, pages 1–6, 2019.
- [17] Hamed HaddadPajouh, Ali Dehghantanha, Raouf Khayami, and Kim-Kwang Raymond Choo. A deep recurrent neural network based approach for internet of things malware threat hunting. *Future Generation Computer Systems*, 85:88–96, 2018. ISSN 0167-739X.
- [18] Weiwei Hu and Ying Tan. Black-box attacks against RNN based malware detection algorithms. *CoRR*, abs/1705.08131, 2017.
- [19] Weiwei Hu and Ying Tan. Generating adversarial malware examples for black-box attacks based on GAN. *CoRR*, abs/1702.05983, 2017.
- [20] Fatima Hussain, Rasheed Hussain, Syed Ali Hassan, and Ekram Hossain. Machine learning in iot security: Current solutions and future challenges. *IEEE Communications Surveys & Tutorials*, 22(3):1686–1721, 2020.
- [21] Chanwoong Hwang, Junho Hwang, Jin Kwak, and Taejin Lee. Platform-independent malware analysis applicable to windows and linux environments. *Electronics*, 9(5), 2020. ISSN 2079-9292.
- [22] Evanson Mwangi Karanja, Shedden Masupe, and Mandu Gasennelwe Jeffrey. Analysis of internet of things malware using image texture features and machine learning techniques. *Internet of Things*, 9:100153, 2020. ISSN 2542-6605.
- [23] Bojan Kolosnjaji, Ambra Demontis, Battista Biggio, Davide Maiorca, Giorgio Giacinto, Claudia Eckert, and Fabio Roli. Adversarial malware binaries: Evading deep learning for malware detection in executables. In *2018 26th European Signal Processing Conference (EUSIPCO)*, pages 533–537, 2018.
- [24] Felix Kreuk, Assi Barak, Shir Aviv-Reuven, Moran Baruch, Benny Pinkas, and Joseph Keshet. Adversarial examples on discrete sequences for beating whole-binary malware detection. *CoRR*, abs/1802.04528, 2018.
- [25] Deqiang Li, Qianmu Li, Yanfang (Fanny) Ye, and Shouhuai Xu. Arms race in adversarial malware detection: A survey. *ACM Comput. Surv.*, 55(1), nov 2021. ISSN 0360-0300. DOI: 10.1145/3484491. URL <https://doi.org/10.1145/3484491>.

- [26] Yair Meidan, Michael Bohadana, Yael Mathov, Yisroel Mirsky, Asaf Shabtai, Dominik Breitenbacher, and Yuval Elovici. N-baiot—network-based detection of iot botnet attacks using deep autoencoders. *IEEE Pervasive Computing*, 17:12–22, 07 2018.
- [27] Jonathan Oliver, Chun Cheng, and Yanggui Chen. Tlsh – a locality sensitive hash. In *2013 Fourth Cybercrime and Trustworthy Computing Workshop*, pages 7–13, 2013.
- [28] Dorottya Papp, Gergely Ács, Roland Nagy, and Levente Buttyán. SIMBioTA-ML: Light-weight, machine learning-based malware detection for embedded iot devices. In *Proceedings of the 7th International Conference on Internet of Things, Big Data and Security - IoTBDS*,, pages 55–66. INSTICC, SciTePress, 2022. ISBN 978-989-758-564-7.
- [29] Daniel Park, Haidar Khan, and Bülent Yener. Generation & evaluation of adversarial examples for malware obfuscation. In *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*, pages 1283–1290, 2019.
- [30] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [31] M. Shobana and S. Poonkuzhali. A novel approach to detect iot malware by system calls using deep learning techniques. In *2020 International Conference on Innovative Trends in Information Technology (ICITIIT)*, pages 1–5, 2020.
- [32] Silvia Wahballa Soliman, Mohammed Ali Sobh, and Ayman M. Bahaa-Eldin. Taxonomy of malware analysis in the iot. In *2017 12th International Conference on Computer Engineering and Systems (ICCES)*, pages 519–529, 2017.
- [33] Jiawei Su, Danilo Vargas Vasconcellos, Sanjiva Prasad, Daniele Sgandurra, Yaokai Feng, and Kouichi Sakurai. Lightweight classification of iot malware based on image recognition. In *2018 IEEE 42nd Annual Computer Software and Applications Conference (COMPSAC)*, volume 02, pages 664–669, 2018.
- [34] Octavian Suci, Scott E. Coull, and Jeffrey Johns. Exploring adversarial examples in malware detection. *2019 IEEE Security and Privacy Workshops (SPW)*, pages 8–14, 2019.
- [35] Hao Sun, Xiaofeng Wang, Rajkumar Buyya, and Jinshu Su. Cloudeyes: Cloud-based malware detection with reversible sketch for resource-constrained internet of things iot devices. *Softw. Pract. Exper.*, 47(3):421–441, mar 2017. ISSN 0038-0644.
- [36] Hayate Takase, Ryotaro Kobayashi, Masahiko Kato, and Ren Ohmura. A prototype implementation and evaluation of the malware detection mechanism for iot devices using the processor information. *International Journal of Information Security*, 19: 71–81, 2019.
- [37] Csongor Tamás, Dorottya Papp, and Levente Buttyán. SIMBioTA: Similarity-based malware detection on iot devices. In *IoTBDS*, pages 58–69. SCITEPRESS, 2021.
- [38] Daniele Ucci, Leonardo Aniello, and Roberto Baldoni. Survey of machine learning techniques for malware analysis. *Computers & Security*, 81:123–147, 2019. ISSN 0167-4048.

- [39] Chenyue Wang, Linlin Zhang, Kai Zhao, Xuhui Ding, and Xusheng Wang. Advandmal: Adversarial training for android malware detection and family classification. *Symmetry*, 13(6), 2021. ISSN 2073-8994. DOI: 10.3390/sym13061081. URL <https://www.mdpi.com/2073-8994/13/6/1081>.
- [40] Yanfang Ye, Tao Li, Donald Adjeroh, and S. Sitharama Iyengar. A survey on malware detection using data mining techniques. *ACM Comput. Surv.*, 50(3), jun 2017. ISSN 0360-0300.