

Dynamic Detection of Mobile Malware Using Smartphone Data and Machine Learning

J. S. PANMAN DE WIT and D. BUCUR, University of Twente, Netherlands

J. VAN DER HAM, University of Twente, Netherlands and NCSC-NL, Netherlands

Mobile malware are malicious programs that target mobile devices. They are an increasing problem, as seen with the rise of detected mobile malware samples per year. The number of active smartphone users is expected to grow, stressing the importance of research on the detection of mobile malware. Detection methods for mobile malware exist but are still limited.

In this article, we provide an overview of the performance of machine learning (ML) techniques to detect malware on Android, without using privileged access. The ML-classifiers use device information such as the CPU usage, battery usage, and memory usage for the detection of 10 subtypes of Mobile Trojans on the Android Operating System.

We use a real-life dataset containing device and malware data from 47 users for a year (2016). We examine which features, i.e., aspects, of a device, are most important to monitor to detect (subtypes of) Mobile Trojans. The focus of this article is on dynamic hardware features. Using these dynamic features we apply state-of-the-art machine learning classifiers: Random Forest, K-Nearest Neighbour, and AdaBoost. We show classification results on different feature sets, making a distinction between global device features, and specific app features. None of the measured feature sets require privileged access.

Our results show that the Random Forest classifier performs best as a general malware classifier: across 10 subtypes of Mobile Trojans, it achieves an F1 score of 0.73 with a False Positive Rate (FPR) of 0.009 and a False Negative Rate (FNR) of 0.380. The Random Forest, K-Nearest Neighbours, and AdaBoost classifiers achieve F1 scores above 0.72, an FPR below 0.02 and, an FNR below 0.33, when trained separately to detect each subtype of Mobile Trojans.

CCS Concepts: • **Security and privacy** → **Mobile platform security**; *Malware and its mitigation*;

Additional Key Words and Phrases: Machine learning, sensor data, smartphones, classification

ACM Reference format:

J. S. Panman de Wit, D. Bucur, and J. van der Ham. 2022. Dynamic Detection of Mobile Malware Using Smartphone Data and Machine Learning. *Digit. Threat. Res. Pract.* 3, 2, Article 9 (February 2022), 24 pages.

<https://doi.org/10.1145/3484246>

1 INTRODUCTION

Nowadays smartphones have become an integral part of life, with people using their phones in both their private and professional life. The number of active smartphone users globally is expected to be 7.3 billion by 2025 [39]. The rise in smartphone users has also led to an increase in malicious programs targeting mobile devices, i.e., mobile malware. Criminals try to exploit vulnerabilities on smartphones of other people for their own purposes. Additionally, over the past years, malware authors have become less recreational-driven and more profit-driven as they are actively searching for sensitive, personal, and enterprise information [40].

Authors' addresses: J. S. Panman de Wit and D. Bucur, University of Twente, Drienerlolaan 5, 7522 NB Enschede, The Netherlands; emails: j.s.panmandewit@alumnus.utwente.nl, d.bucur@utwente.nl; J. van der Ham, University of Twente, Drienerlolaan 5, 7522 NB Enschede, The Netherlands and NCSC-NL, The Hague, Netherlands; email: j.vanderham@utwente.nl.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2022 Copyright held by the owner/author(s).

2576-5337/2022/02-ART9 \$15.00

<https://doi.org/10.1145/3484246>

Detection of mobile malware is becoming increasingly relevant, with machine learning showing the most promise. In this article, we present an overview of the performance of state-of-the-art machine learning classifiers using sensor data on the Android mobile operating system. We show different training approaches for each of the classifiers, showing their performance on individual and general cases. These machine learning classifiers can be trained on a cluster on a large dataset, resulting in a classifier. This classifier can then be deployed on mobile phones to detect future malware based on the dynamic sensor data similar to the dataset, i.e., dynamic analysis.

Academic work is mainly divided into dynamic analysis and static analysis of mobile malware. Dynamic analysis refers to the analysis of malware during runtime, i.e., while the application is running. Static analysis refers to the analysis of malware outside runtime, e.g., by analysing the installation package of a malware app. Dynamic analysis has advantages over static analysis but methods are still imperfect, ineffective, and incomprehensive [42].

An important limitation is that in most studies of malware detection, virtual environments are used, e.g., analysis on a PC, instead of real mobile devices. An increasing trend is seen in malware that use techniques to avoid detection in virtual environments, thereby making methods based on analysis in virtual environments less effective than methods based on analysis on real devices [40]. Moreover, we found that most methods are assessed with (i) malware running isolated in an emulator and (ii) malware running for a brief period. This kind of assessment does not reflect the circumstances of a real device with for example different applications running at the same time. Therefore, most research does not provide a realistic assessment of detection performances of their detection methods due to their unrealistic circumstances.

Using a dataset [34] containing data of 47 users throughout the year 2016, this research provides valuable knowledge on detecting mobile malware on real devices and a realistic assessment of the research's detection methods. This article focuses on dynamic hardware features as our initial literature exploration showed relatively few research compared to other types of dynamic features (e.g., software and firmware). In this work, we compare the detection performance of the Random Forest, K-Nearest Neighbour, and AdaBoost classifiers. The classifier performance is assessed on 10 subtypes of Mobile Trojans, i.e., malware showing benign behaviour but performing malicious behaviour, as these are the most prevalent mobile malware type [27, 29]. Furthermore, we examine which features are important in detecting different subtypes of Mobile Trojans. We find that the Random Forest and AdaBoost classifiers have similar and better performances than the K-Nearest Neighbour. Additionally, our initial results show that AdaBoost may be more suitable than the other classifiers, for use cases where detection models are built with data from one subset of users and subsequently used for protecting other users. Additionally, we show that application-specific features are important in detecting mobile malware. In contrast to other literature, which often assessed detection models on devices running one or few applications in virtual environments, we find low predictability using only global device features.

In summary, we make the following contributions:

- We use a recent real-life sensor dataset to create multiple mobile malware detection models. Most literature is based on environments with models based on emulator data, simulated user input, and/or short testing durations. Our research takes a realistic approach, contributing to the body of scientific literature on malware detection that is based on real-life data rather than data derived from emulators and virtual environments.
- Our results are based on a dataset containing data from 47 real users over a time span of 1 year. To the authors' best knowledge, this is one of the first research on data from this many users over a long time period. The large user pool and time span increases the confidence in the generalisability of the results.
- We evaluated our detection methods on a Random Forest, K-Nearest Neighbour, and AdaBoost classifiers with several performance metrics. AdaBoost has been researched little in literature. We provide the performance of the classifiers with several training and testing methods.

- Performance of the classifiers is shown for separate mobile malware types. The results show a large difference in performance for several mobile malware types, indicating the difficulty of detecting some mobile malware types. Earlier research, to the best of our knowledge, did not include a comparison of this many mobile malware types.
- We briefly describe a testing method where models are trained on data of one set of users and tested on data from another set of users. This testing method may be useful when detection models are built with one user pool and are subsequently used to protect another, perhaps larger, user pool.
- All features included in this research do not require any privileged access. This increases the usability of mobile applications based on our detection methods. Additionally, we show which feature categories and which specific features are most important in detecting different malware types.

2 RELATED WORKS

In an early, highly cited study (Shabtai [38] from 2011), the authors designed a behavioural malware detection framework for Android devices called *Andromaly*. This used device data: touch screen, keyboard, scheduler, CPU load, messaging, power, memory, applications, calls, processes, network, hardware, binder, and LED. Forty benign applications and 4 self-developed malware applications (a DOS Trojan, SMS Trojan, Spyware Trojan, and Spyware malware) were executed. When training the detection model on data from the same mobile device, the J48 decision tree classifier performed best; however, a Naive Bayes classifier detected malware best in experiments in which the testing was done on data from a new device (see the overview in Table 1). Despite relatively high FPR with all their models, *Andromaly* showed the potential of detecting malware based on dynamic features using machine learning, compared different classifiers, trained and tested on data collected from real devices. However, the study is relatively old, and the malware ecosystem has since changed.

The *STREAM* framework [11], aims to enable rapid large-scale validation of machine-learning classifiers for mobile malware. *STREAM* used 41 features (listed in Table 1) collected every 5 s from emulators running in a so-called *ATAACK* cloud. The emulator used the Android Monkey application to simulate pseudo-random user behaviour such as input on the touchscreen. To evaluate their detection model, the authors used six classifiers, with good accuracy and **True Positive Rate (TPR)** for the Bayesian Network. The training set had 408 popular applications from the Google Play Store and 1,330 malware applications from the Malware Genome Project database [43] and the VirusTotal database [3]. The testing set had 24 benign and 23 malware applications from the same databases. This showed the potential of using dynamic features, although the FPR for all classifiers was relatively high. Additionally, this research ran applications separately for 90 s and made use of a virtual environment in the form of an emulator with userlike behaviour created by the Android Monkey tool. This lowers the confidence that the model would perform the same in real life. The same dataset from Amos [11] was used in Alam [8] for anomaly detection with application behaviour features: only the Binder, CPU, and memory features. The dataset was then balanced. The Random Forest classifier had an excellent accuracy of 99.9857% and a root MSE of 0.0183%. Only 2 false positives were measured during this experiment. However, it is sensitive to the same limitations as Amos [11].

Saracino [37] describes an improvement of an earlier anomaly detection model (Dini [22]). This version includes system calls, user presence, and SMS features, plus a static analysis of application packages. Many classifiers were trained and tested. Their performance is not reported, but the authors state that 1-Nearest Neighbour achieved the best classification results. To evaluate the TPR, the detection framework has been tested against 2,800 malware applications from the Malware Genome Project, Contagio [6], and VirusShare [2], representing the Botnet, Installer, Trojan, Ransomware, Rootkit, SMS Trojan, and Spyware malware families. How long the different malware applications were run is not mentioned. Also, it is unknown whether the same device was used for the training and testing of the classifiers. The framework achieved very good TPR and FPR, and also detected some zero-day attacks, undetected by multiple Antivirus Software at that time. The performance was

Table 1. Related Work

Article		Features		Training and testing			Performance			
Ref	Year	Dynamic	Static	Benign	Malware	Platform	Classifiers	Acc	TPR	FPR
[38]	2012	various (14)	–	40	4 ^{cust}	2 devices	BN, histo, J48 , Kmeans, LR, NB	0.81	0.79	0.48
[11]	2013	bat, Binder, CPU, mem, netw	perm	408 ^{PS}	1330 ^{Ge,VT}	VE, Monkey	BN , J48, LR, MLP, NB, RF	0.81	0.97	0.31
[8]	2013	Binder, CPU, mem	–	408 ^{PS}	1130 ^{Ge,VT}	VE, Monkey	RF , BN, NB, MLP, J48, DS, LR	1.00	–	few
[28]	2013	CPU, net, mem, SMS	–	30 ^{PS}	5 ^{cust}	1 device	NB, RF , LR, SVM	–	0.99	0.00
[13]	2014	bat, CPU, mem, netw	–	>2 ^{PS}	3 ^{Cr}	12 devices	GM-LDC	≈1	≈1	≈0
[23]	2014	bat, time, loc	–	–	2 ^{cust}	11 devices	stdev	≈1	–	≈0
[37]	2016	SC, SMS, UP	md	9804 ^{PS}	2800	3 devices	1-KNN, LDA, QDC, MLP, PC, RBF	–	0.97	0.00
[19]	2016	bat	–	–	7 ^{cust}	1 device	NN , DT	–	>0.85	–
[33]	2016	CPU, mem	–	940 ^{PS}	1120 ^{Ge}	VE, Monkey	LR	–	0.86	0.17
[25]	2016	CPU, mem, SC	–	1709 ^{PS}	1523 ^{Dr}	VE, Monkey	Kmeans-RF	0.67	0.61	0.28
[18]	2016	CPU, mem, net, sto	–	1059 ^{PS}	1047 ^{Dr}	VE, Monkey	RF	1.00	0.82	0.01
[31]	2017	CPU, mem, net	–	–	<5560 ^{Dr}	VE, Monkey	SVM	0.82	–	–
[41]	2018	CPU, net	–	>10	3 ^{ShD}	47 devices	DT	1	1	<1%
[16]	2018	Method calls, Intents	–	17k ^{PS, AZ}	13.9k ^{mult}	VE, Monkey	RF	0.98	0.98	–
[30]	2018	Network, SMS, Files, more	–	–	4442 ^{Dr}	VE, Monkey	RF , DT, KNN, SVM, NN	0.813	0.813	–
[20]	2019	CPU, mem, api	–	–	27000 ^{Om}	VE	SVM, RF, NN	0.97	–	–
[32]	2019	Bat, CPU, net, proc	–	3 ^{ShD}	3 ^{ShD}	47 devices	IR, RF, DT, GBT , MLP, SVM, LR	1	1	<1%
[10]	2020	API calls	perm	11505	19620	8 devices, DynaLog	NN , NB, SVM, DT	0.95	0.98	0.09
this		CPU, bat, mem, net, sto	–	10 ^{cust}	10 ^{cust}	47 devices	RF , NB, KNN, MLP, AdaBoost	0.96	0.65	0.01

When multiple malware classifiers are studied, the best one is emphasized.

Legend: Acc, accuracy; AZ, Androzo [9]; bat, battery; BN, Bayesian Network; Cr, Crowdroid [14]; cust, Custom; Dr, Drebin [12]; DS, Decision Stump; DT, Decision Tree; GBT, Gradient Boosted Tree; Ge, Malware Genome Project [43]; GM, Gaussian Mixture; histo, Histogram; IR, Isotonic Regression; Kmeans, k -means clustering; KNN, K-Nearest Neighbour; LDA, Linear Discriminant Analysis; LDC, Local Density Cluster-Based Outlier Factor; loc, Location; LR, Logistic Regression; md, metadata; MLP, Multi-Layer Perceptron; mult, Multiple datasets; NB, Naive Bayes; netw, network; NN, Neural Network; Om, Open Malware [7]; perm, permissions; PC, Parzen classifier; proc, process; PS, Play Store; QDA, Quadratic Discriminant Analysis; RBF, Radial Basis Function; RF, Random Forest; sens. act, Sensitive activities; ShD, Sherlock dataset; stdev, standard deviation; sto, storage; SVM, Support Vector Machine; SC, system calls; UP, user presence; VE, Virtual Environment; VS, VirusShare [2]; VT, VirusTotal [3].

assessed on real devices, varying the usage intensity over a period of one week, to reflect real-life circumstances. A limitation is a requirement for root permissions, as system calls were used as features.

In Milosevic [33], memory and CPU usage are predictive features, tracked by running every application separately for 10 minutes in an Android emulator. The emulator was fed with userlike input by the Monkey application. Their optimized feature set only contained 7 features. The classification algorithm used was linear Logistic Regression with the use of a sliding window technique. This was validated on a set of 94 benign and 89 malware applications. The classifier with the highest TPR of 95.7% of malware had a relatively high FPR of 25%; the classifier with the highest F-measure had a TPR 85.5% and an FPR of 17.2%. This study showed the potential of using the memory and CPU usage as features for the dynamic detection of mobile malware. However, the detection had relatively high FPR, besides the disadvantage of using emulator data.

More studies were based on emulation. In Ferrante [25], the features were related to system calls, memory, and CPU usage. Benign apps were downloaded from the Google Play store and malicious applications were collected

from the Drebin dataset [12]. The applications were run for 10 minutes in an Android emulator and features were collected every 2 s. The emulator was fed with user-input from the Monkey application. As detection models, they first used a k -means clustering algorithm to cluster apps based on the similarity of memory and CPU usage, then a Random Forest classifier on every cluster that classified the applications based on their system calls. The best performing model was the classifier with 7-means clustering and a Random Forest classifier of 50 trees. This showed the potential of using system calls, memory, and CPU features for malware detection. The performance is relatively low compared to other detection methods, and the method needs root permission as it uses system calls as features. In Canfora [18], the features are related to CPU, memory, storage and network. The applications were run for 60 s in an Android emulator. The features were used both raw and after a Discrete Cosine Transformation. Furthermore, the authors varied the granularity of the detection method by either taking all the features, only the global features, or only the features of the application under analysis. The Random-Forest classifier using global features had very good accuracy and FPR.

In Massarelli [31], the authors used the Drebin dataset for malicious applications (but it is unknown how many malware applications were included from the 5,560 in the dataset). The features were both systemwide and application-specific. An emulator was used for training the model, fed with simulated user input events from the Monkey application. The classifier was an SVM with a Radial Basis Function kernel, which had good accuracy. The exact FPR is unknown, but the precision of the model ranges from 10% to 90% depending on the malware family, which is low in comparison with other methods. This also used an emulator, leading to limitations in its applicability.

In Martin [30], multiple classifiers, including Random Forest, K-Nearest Neighbour, and SVMs are assessed. They achieve an F1 score of 0.803 on a dataset containing 4,442 malicious malware samples. They use hardware features related to network, SMS, file reads/writes. Other features included are: services started, classes loaded, information leaks via the network, file and SMS services, circumvented permissions, cryptographic operations, SMSs sent, and phone calls. Their detection models are also assessed using a virtual environment rather than devices. This is also the case for Dai [20]. They use features related to CPU, Memory, and API calls (but also software features such as API calls, which this article does not). Their best classifier is well performing with an F1 score of 0.98. Additionally, they assess their detection model on a large dataset of more than 27.000 malicious applications.

During the writing of this article, other authors used the SherLock dataset but to a much more limited extent than this research. In Memon [32], seven classifiers were trained for malware detection: Isotonic Regression, Random Forest, Decision Trees, Gradient-Boosted Trees, MultiLayer Perceptron, SVM, and Logistic Regression. The data were balanced so that 50% of the labels were benign and 50% were malicious. Only data from Q3 2016 were studied, meaning that only three malware types (Madware, Ransomware, and Click-jacking) were included for training and testing, with CPU, network traffic, battery, and process features. The Gradient-Boosted Trees had the best results with an F1 score of 0.91% and an FPR of 0.09%. Their Random Forest classifier also worked almost equally well. These make for better results than prior work, which may be due to the limited number of malware types included.

In Wassermann [41], Decision Trees were trained as classifiers, only on SherLock data from Q2 2016, so only three malware types (Spyware, Phishing, Adware). The features used were network traffic and CPU related. The label to predict was whether the malicious app (Moriarty) was running. Their model results in a recall of almost 100% with less than 1% FPR. This research showed good app-detection results, although on a limited set of malware; also, the model cannot distinguish between individual (benign or malicious) actions; it only predicts the presence of the malware on the system.

Recent literature on dynamic malware detection (Droidcat [16] and DI-Droid [10]) have high performance on a large dataset containing more than 10,000 malware applications. However, DroidCat focuses on software features such as Method Calls and trained and tested their method within a virtual environment with simulated

pseudo-random user behaviour. DL-Droid achieves a TPR of 0.95 and an FPR of 0.09 but also only including software features such as API calls. This article focuses on hardware features only.

Recent work by Cai et. al. [17] analysed 17.664 Android applications developed throughout 2010–2017. The paper describes differences in method calls, ICC calls and source/sink calls during static-code analysis and dynamic analysis. One of the findings of the paper is that access to sensitive data was almost 10× more frequent during dynamic analysis than during static analysis (their Figure 17 and Finding 8). This finding indicates that dynamic analysis may support in capturing all sensitive data calls and could therefore help in detecting mobile malware.

Other research assessed have also used machine learning for the dynamic detection of mobile malware, but this literature is unclear about the feature collection method (Ham and Choi [28]), include few malware samples (Attar [13] and Dixon [23]), or experiment on an idle phone (Caviglione [19]).

While we have seen a body of prior work using dynamic hardware features to detect malware on mobile devices, important issues remain unaddressed. Most recent papers have trained their detection model on emulators with simulated input; this makes it difficult to extrapolate their performances with real users and to have high confidence in how realistic the results are. Additionally, most studies have only executed the benign and malicious applications for a few minutes, meaning that any delayed malicious behaviour would not be detected. Also, a limited amount of research has shown excellent results, particularly in what regards the FPR metric. In this article, we address these issues and provide a more realistic picture by learning from data gathered with real users, from a study that executed malware from different categories, over a large time span.

3 METHOD

3.1 Malware Data

We use the SherLock dataset [34] provided by Ben-Gurion University. This dataset contains 10 billion data records in system logs collected from 47 Samsung Galaxy S5 devices used by 47 different volunteer users, over a period of over one year (2016). During this year, the volunteers executed both benign and malicious applications on their devices. The data collection agent on the phone is based on the FunF Open Sensing framework, developed by MIT Media Lab [1]. This framework allows for the collection of sensor data (e.g., memory consumption or CPU usage). The authors adjusted the source code to improve stability and reliability. Additionally they added features such as the collection of statistics on all running applications of a smartphone. According to the SherLock paper [34] all features can be tracked without any root permissions. The source code of the data collection agent can be found on GitHub [5].

The malicious applications were written by the researchers based on wild malware. Every month a different subtype of Mobile Trojan was installed on the devices. Each malware version resembled a subtype of Mobile Trojan with both benign and malicious actions. In total, 11 malware versions were included, listed in Table 2 with their benign behaviour, malicious behaviour (as described in Reference [34]), malware type, a description of their actions, and the wild malware on which the implementation was based. Important to note is that each time data were transmitted by the malware application, it was scrambled prior to sending to protect the privacy of the users.

The dataset is divided into 13 probes, namely groups of multiple sensors that shared the same sample interval. This research uses the following probes:

- the Moriarty (*Malware*) probe for malware data,
- the T4 (*System*) probe for global device data, and
- the Applications (*Apps*) probe for app-specific device data.

The Malware probe sensed data from the malicious application installed on the devices. Each malware probe's record is a log of the *action* taken by the user's malware application. Each action is part of a malware *session*:

Table 2. Malware Types

	Benign behaviour	Malicious behaviour	Type	Description	Wild malware
1	Game	Contact theft	Spyware	Steals, encrypts, and transmits all contact stored on device.	SaveMe, SocialPath
2	Web browser	Spyware	Spyware	(i) Spies on location and audio, or (ii) spies on web traffic and web history.	Code4hk, xRAT
3	Utiliz. Widget	Photo theft	Spyware	Steals photos that are taken and in storage, and takes candid photos of the user.	Photsy, Phopsy
4	Sports app	SMS bank thief	SMS Fraud	Captures and reports immediately on SMSs that contain codes and various keywords.	Spy.Agent.SI
5	Game	Phishing	Phishing	Makes fake shortcuts and notifications to login to Facebook, Gmail, and Skype.	Xbot
6	Game	Adware	Adware	Gathers information and places ads, popups, and banners.	—
7	Game	Madware	Hostile downloader	Gathers private information, places shortcuts, notifications, and attempts to install new applications.	—
8	Lock-screen	Ransomware	Ransomware	Performs either: (1) lock screen ransomware, or (2) crypto ransomware.	Simplocker.A, SLocker
9	File Manager	Click-jacking	Privilege escalation	Tricks the user to activate accessibility services to then hijack the user interface.	Shedun (GhostPush)
10	—	Device theft	—	—	—
11	Music Player	Botnet	DOS	Either performs: (1) DDoS attacks on command, or (2) SMS botnet activities	Tascudap.A, Nitmo.A...
12	Web media player	Recon. Infiltr.	Other	Maps the connected local network and searches for files and vulnerabilities.	—

Table 3. Malware Probe Data Overview

Column	Datatype	Description
UserId	String	User ID
UUID	Numeric	Timestamp of action
Details	String	Details of action
Action	String	Action performed
ActionType	String	Whether action is benign or malicious
SessionType	String	Whether session is benign or malicious
Version	Decimal	Version number
SessionID	Decimal	Session ID
Behavior	String	Behaviour of current session

the malware application started in either a benign or malicious session. Within a benign session, only benign actions were performed; in a malicious session, the malware performed both benign and malicious actions. An overview of Malware probe's columns, their data type, and a description is shown in Table 3. For this research, we excluded version 10 and 12 as no malicious data were available for these versions.

The System probe tracked global device data every 5 s. Each record of the system probe is a log of the user's global device data at a given time, taken from the /proc/ folder. The feature categories that were tracked are battery, CPU, network, memory, I/O interrupts, and storage. Forty-one features are used from this probe.

The App probe recorded app data every 5 s for each application installed on the device. For this research, the only relevant data are the app data of the malware application. Each record is a log of the malware app data at a given time, taken from the /proc/ folder. 35 features are used from this probe. For a complete overview of the features used, refer to Table 12.

Table 4. Pros and Cons of Classifiers

Classifier	Pros	Cons
RF	Can overcome overfitting	A greedy algorithm
AdaBoost	Nonlinear classification Efficient with high-dimensional data	Sensitive to imbalanced dataset
KNN	High precision and accuracy Nonlinear classification No assumption of features	Computationally expensive Sensitive to imbalanced dataset

The data from the three probes are integrated by joining on the *userid* and *timestamp* fields. Not all probes had the same tracking frequency so joining on the timestamp column cannot be done with a precise join. Therefore, probe data with a maximum of 5 s after the timestamp of the Malware data are used in the joining of the data. This way we can match 83% of the Malware data with probe data. Additionally, we apply random undersampling to rebalance the original dataset. The original dataset contains 90% malicious records and only 10% benign records. The dataset is balanced by downsizing the number of malicious data points per malware type until 90% of the rows are benign actions and 10% are malicious. The malicious data points that are removed are chosen uniformly at random. The undersampling method is applied to both the training and testing data. We apply the resampling method on the training data to prevent the classifiers from a bias toward the (original) majority class (i.e., malware data). We apply the resampling method on the testing data to better reflect real-life circumstances and to provide more realistic classification results. In the real world, the number of malicious actions on a mobile phone would be relatively low compared to the number of benign records. Last, we normalize any features in the dataset by scaling it between zero and 1 for the KNN classifier. No data normalization is used for AdaBoost and RF, as these are tree-based classifiers. The dataset now consists of 28,821 rows and 102 columns and is 21.4 MB.

3.2 Supervised Learning Algorithms

We train a statistical classifier able to recognize malware signatures in any log data collected on a smartphone. The classifier is trained, cross-validated, and tested using the dataset described above. Three training algorithms yielded well-performing classifiers: *Random Forest* and *K-Nearest Neighbour* classifiers had a good performance in our related work (see Section 2), so we include them in our study. We add the *AdaBoost* classifier, which is designed specifically to improve the performance of a Random Forest. All are nonlinear classifiers, able to capture complex relationships between the variables in the dataset. Table 4 (based on Reference [24]) lists advantages and disadvantages for these algorithms.

3.3 Training, Cross-validating, and Testing Malware Classifiers

All classifiers are trained to identify a malicious action (the *ActionType* column of the Malware dataset). The classifiers that are trained differ in their input data (feature set), classification target, and test mode. Three feature sets are assessed: Global, Apps, and a combination of both. The global feature set includes global device features from the System probe. The apps feature set include features from the Apps probe.

Two classification targets are tried. The first classification target (all malware) trains one general malware classifier to identify malicious actions by any of the 10 malware types in the dataset. The second classification target (single malware) trains one single-purpose classifier per malware type, resulting in 10 different trained models.

Two test modes are run. Both use a percentage of the data for training, and the remaining for testing. In the *normal holdout test mode*, the training and testing data are sampled at random. In the *unknown device test mode*, the testing data contain data of devices that are not in the training data. The unknown device test mode is used to

Table 5. Different Configurations for the Classifiers Trained

	Type	Description
Feature set (complete listing in Table 12)	Global	Device features (the System probe)
	Apps	App-specific features (the Apps probe)
	Combined	Both Global and Apps features
Classification target	All malware	Generic classifier across malware types
	Single malware	One classifier for each malware type
Test mode	Normal holdout	Training data includes all devices
	Unknown device	Test set includes data from new devices

simulate a situation in which no data are yet available for a certain device. In that case, a model trained on data from other devices is a solution. Testing the models on other devices also helps to detect the potential bias of the model on a specific device. Note that only the testing on the test set is adjusted; the cross-validation remains the same. For a complete overview of the different configurations, please refer to Table 5.

To train the dataset, we use fourfold cross-validation with holdout. The hyperparameters of each classifier are tuned using a grid search. The search space for the hyperparameter settings are as follows. For AdaBoost, the number of estimators is in the interval [5, 400]; for Random Forest, the number of trees in the interval [5, 320], the maximum tree depth in [3, 320], the maximum number of features in [3, 102]; for K-Nearest Neighbours, the number of neighbours in [1, 61]. We also use Recursive Feature Elimination together with cross-validation (RFCV) to find the optimal number of features for all classifiers. RFCV requires a feature ranking, as the method recursively eliminates the least important feature. The feature rankings are based on the Random Forest classifier and are calculated per different experiment setup (i.e., per classification target, per test mode, per feature set).

To evaluate a classifier, we use the standard *F1 score* [42], namely the harmonic mean of the Precision and Recall scores. These two metrics are normally a tradeoff of each other, and a high F1 score indicates a good balance between these two metrics. Additionally, the training and testing times of the classifiers are assessed.

The data are processed on a computing cluster with a Hadoop File System suitable for big data. The cluster consists of 55 nodes with each 32 GB RAM and 8–16 cores running Apache PySpark 2.2.0. The training and testing is done with the package Spark Sklearn [4], a distributed implementation of the machine-learning classifiers in the popular package machine-learning library Scikit-learn [36]. Spark Sklearn is adjusted to provide and log more information during the execution of experiments.

4 RESULTS

4.1 Test Mode *Normal Holdout*

An overview of the performances of the classifiers for classification target *all malware* and test mode *normal holdout* is shown in Table 6. This table shows the feature categories present in the best model, i.e., with the best hyperparameter settings. From all models with similar performance as the best model of that classifier (performance results that do not differ statistically significantly (McNemar test $a < 0.05$ [21]), the classifier with the least number of features is chosen, because this shows which features are relevant for detecting specific malware types. The number of features is based on the RFCV method described in Section 3. Table 6 shows the following:

- (1) RF has the highest F1 score of 0.73 with feature set Apps and 0.72 with feature set Combined.
- (2) All classifiers have an F1 score below 0.42 with feature set Global. This suggests that global features are less discriminative when detecting malicious actions.
- (3) RF with feature set Combined uses 10 app features (related to App CPU, App Memory, App Process) to detect malicious actions of malware, suggesting that 10 app features are sufficient in detecting malicious actions of mobile malware. These 10 features are shown in Table 13. For the features included in the best

Table 6. Comparison Performance per Classifier for Classification Target *all Malware*

Classifier Feature set Best trained model	AdaBoost			Random Forest			KNN		
	Global	Apps	Comb.	Global	Apps	Comb.	Global	Apps	Comb.
Nr. features	26	25	43	24	29	10	26	13	9
Accuracy	0.922	0.943	0.945	0.922	0.958	0.959	0.897	0.944	0.946
F1 score	0.226	0.595	0.610	0.422	0.730	0.722	0.356	0.674	0.688
FPR	0.003	0.012	0.013	0.022	0.013	0.009	0.048	0.029	0.028
FNR	0.868	0.522	0.502	0.670	0.346	0.380	0.674	0.336	0.320
Feature categories used									
Battery	✓		✓	✓			✓		
CPU	✓		✓	✓			✓		
IO Interrupts									
Memory	✓		✓	✓			✓		
Metadata									
Network	✓		✓	✓			✓		
Storage									
Wifi									
App CPU		✓	✓		✓	✓		✓	✓
App Info									
App Memory		✓	✓		✓	✓		✓	✓
App Network traffic		✓	✓		✓	✓		✓	✓
App Process		✓	✓		✓	✓		✓	✓

The best performing classifier (based on f1-score) is emphasized in bold.

performing AdaBoost model and KNN model, please refer to Tables 14 and 15 respectively. All features and their corresponding feature category are shown in Tables 16 and 17.

- (4) All classifiers have a higher FNR than FPR. This indicates that most models incorrectly classify a malicious action as benign more often than vice versa.

Table 7 shows an overview of the performance per classifier per malware type and feature set. Following the same reason as for Table 6, the performances are shown for the least number of features and the best hyperparameter settings per classifier. Table 7 shows the following:

- (1) RF has the highest F1 score for 5 of 10 malware types.
- (2) KNN has the highest F1 score for 5 of 10 malware types.
- (3) AdaBoost has the highest F1 score for 4 of 10 malware types.¹
- (4) AdaBoost has a high F1 score (>0.7) for 8 of 10 malware types.
- (5) RF has a high F1 score for 7 of 10 malware types.
- (6) KNN has a high F1 score for 7 of 10 malware types.
- (7) All classifiers have a low F1 score for malware type 4 (Spyware SMS) and 6 (Adware).

The feature categories included in the feature sets of the best models of classifiers per malware type are shown in Table 8. Following the same reason as for Table 6, the performances are shown for the least number of features and the best hyperparameter settings per classifier. This shows the following:

- (1) Six of 10 malware versions are best detected using only app features.
- (2) Malware version 4 (Spyware SMS) is best detected using only global features. However, note that the F1 score for Spyware SMS is the lowest of all malware versions.

¹Some performance scores are the same, leading to a draw, hence to more than expected highest F1 scores.

Table 7. Comparison Performance per Classifier for Classification Target *Single Malware* and Test Mode *Normal Holdout*

Malware	AdaBoost			Random Forest			KNN		
	F1	FPR	FNR	F1	FPR	FNR	F1	FPR	FNR
1	0.704	0.025	0.321	0.571	0.025	0.500	0.536	0.047	0.464
2	0.944	0.003	0.085	0.956	0.001	0.077	0.945	0.003	0.077
3	0.764	0.000	0.382	0.793	0.003	0.324	0.727	0.003	0.412
4	0.400	0.029	0.667	0.500	0.000	0.667	0.400	0.029	0.667
5	0.927	0.005	0.089	0.927	0.005	0.089	0.946	0.005	0.054
6	0.194	0.011	0.880	0.271	0.039	0.782	0.352	0.057	0.662
7	0.813	0.015	0.216	0.857	0.003	0.227	0.793	0.012	0.268
8	1.000	0.000	0.000	0.800	0.000	0.333	0.857	0.032	0.000
9	0.913	0.000	0.160	0.936	0.000	0.120	0.936	0.000	0.120
11	0.909	0.009	0.118	0.938	0.000	0.118	0.938	0.000	0.118
avg.	0.757	0.010	0.292	0.755	0.008	0.324	0.743	0.019	0.284
stdev	0.263	0.010	0.283	0.232	0.013	0.251	0.233	0.021	0.252

The numerical malware types in the first column are those listed in Table 2.
 The best performing classifier (based on f1-score) per malware type is emphasized in bold.

Table 8. Best Performing Classifiers per Malware Type

Malware	1	2	3	4	5	6	7	8	9	11
Best trained model										
Classifier	AdaBoost	KNN	RF	RF	KNN	KNN	RF	AdaBoost	RF	KNN
Feature set	Comb.	Apps	Apps	Global	Apps	Apps	Apps	Comb.	Comb.	Global
Nr. features	49	8	5	28	7	11	9	14	10	1
F1 score	0.70	0.94	0.73	0.50	0.95	0.35	0.86	1.00	0.94	0.94
Feature categories used										
Battery	✓			✓					✓	
CPU	✓			✓					✓	
I/O Interrupts										
Memory	✓			✓						
Network traffic	✓			✓					✓	✓
Storage										
Wifi										
App CPU	✓	✓	✓		✓	✓	✓	✓	✓	
App memory	✓	✓	✓		✓	✓	✓	✓	✓	
App network traffic	✓					✓				
App process	✓	✓	✓					✓		

The numerical malware types in the first column are those listed in Table 2.

- (3) Malware version 11 (DOS) is best detected using only one global feature regarding network traffic.
- (4) Malware version 1 and 9 (Spyware contacts theft and Privilege Escalation / Hostile downloader) is best detected using both the Apps and Global features.
- (5) None of the classifier includes I/O interrupts, Storage, or Wifi features for any of the malware versions.

4.2 Test Mode *Unknown Device*

An overview of the performances of the classifiers for classification target *all malware* and test mode *unknown device* is shown in Table 9. This table shows high FNR for all classifiers. Therefore, we do not include the feature categories as in the previous sections, as little insight can be drawn from these poorly performing models.

Table 9. Comparison Performance Per Classifier for Classification Target *All Malware* and Test Mode *Unknown Device*

Classifier Feature set	AdaBoost			Random Forest			KNN		
	Global	Apps	Comb.	Global	Apps	Comb.	Global	Apps	Comb.
Best model									
Nr. features	29	28	44	28	6	8	30	4	11
Accuracy	0.923	0.944	0.926	0.888	0.918	0.925	0.824	0.917	0.910
F1 score	0.207	0.583	0.563	0.249	0.562	0.561	0.202	0.537	0.522
FPR	0.005	0.013	0.042	0.052	0.057	0.043	0.126	0.053	0.061
FNR	0.878	0.528	0.423	0.776	0.362	0.423	0.732	0.415	0.407

The best performing classifier (based on f1-score) is emphasized in bold.

Table 10. Comparison Performance per Classifier for Classification Target *Single Malware* and Test Mode *unknown device*

Malware	AdaBoost			Random Forest			KNN		
	F1	FPR	FNR	F1	FPR	FNR	F1	FPR	FNR
1	0.444	0.028	0.600	0.364	0.056	0.600	0.316	0.153	0.400
2	0.817	0.015	0.230	0.861	0.014	0.164	0.847	0.027	0.115
3	0.618	0.013	0.514	0.712	0.013	0.400	0.597	0.080	0.343
4	0.000	0.028	1.000	0.000	0.000	1.000	0.000	0.000	1.000
5	0.967	0.000	0.064	0.978	0.000	0.043	0.932	0.000	0.128
6	0.200	0.056	0.817	0.245	0.037	0.802	0.207	0.134	0.706
7	0.480	0.027	0.500	0.583	0.019	0.417	0.476	0.015	0.583
8	1.000	0.000	0.000	0.000	0.000	1.000	0.667	0.032	0.000
9	0.901	0.037	0.089	0.889	0.037	0.111	0.871	0.022	0.178
11	0.977	0.000	0.045	0.930	0.010	0.091	0.952	0.000	0.091
avg.	0.640	0.020	0.386	0.556	0.018	0.463	0.586	0.046	0.354
stdev	0.352	0.018	0.353	0.380	0.019	0.372	0.330	0.057	0.322

The numerical malware types in the first column are those listed in Table 2.

The best performing classifier (based on f1-score) per malware type is emphasized in bold.

Table 10 shows an overview of the performance per classifier per malware type and feature set. Comparing this table with Table 7, we see the following:

- (1) All classifiers show lower performance for 7 of 10 malware versions (versions 1–4 and 7–9) for test mode *unknown device*, compared to test mode *normal holdout*.
- (2) AdaBoost and RF show better performance for malware version 5 (Phishing) for test mode *unknown device*, compared to test mode *normal holdout*.
- (3) AdaBoost shows a perfect score (f1 score of 1) for detecting malware version 8 (Ransomware) for test mode *unknown device*. For test mode *normal holdout*, AdaBoost showed a perfect score as well.
- (4) AdaBoost and KNN shows a better performance for malware version 11 (DOS) for test mode *unknown device*, compared to test mode *normal holdout*.

4.3 Training and Testing Times

Table 11 shows the training and testing times per classifier. This shows that training times of both AdaBoost and RF are significantly higher than for KNN (an algorithm that does no training). In contrast, testing times for KNN are significantly higher in comparison.

Table 11. Training and Testing Times per Classifier for Classification Target *all Malware* and Test Mode *Normal Holdout*

Classifier Feature set	AdaBoost			Random Forest			KNN		
	Global	Apps	Comb.	Global	Apps	Comb.	Global	Apps	Comb.
Train (s)	17	18	31	<1	30	10.5	<1	<1	<1
Test (s)	<1	<1	<1	<1	<1	<1	2	3	8

5 DISCUSSION

Below a discussion on the main findings of this research is presented. Most related papers use a different evaluation method making direct comparison difficult. Nevertheless, we tried ordering the findings from good to bad, compared to existing detection methods.

High-performance RF, KNN, AdaBoost for classification target single malware. In this research, RF and KNN showed good results (F1 scores on average above 0.7) when separate models were trained per malware type. These findings are in line with earlier studies that found good performances for RF [8, 18, 28] and KNN [37] on the detection of mobile malware. None of the literature we found examined the AdaBoost classifier, hence the good performance of AdaBoost is a new finding.

AdaBoost is consistent with different test modes. All classifiers, except AdaBoost, show worse performances when tested on new devices (test mode *unknown device*), than if tested on the same devices (test mode *normal holdout*). Most studies do not provide information on the test mode, thus test modes of other research cannot be compared to our study. Our results suggest that AdaBoost may be more suitable when its purpose is to first build a detection method with a subset of users and subsequently use this detection method to protect other users. This approach is scalable: new users can be protected with an already existing detection method. Otherwise, the model needs to be retrained constantly for every new user. AdaBoost with the Apps feature set and test mode *unknown device* had an F1 score of 0.583, an FPR of 0.013, and an FNR of 0.528. The high FNR calls for improvements in detecting malicious actions of new devices, given the better scalability of this method. Some suggestions are given in Section 8.

Perfect score on Ransomware, but few data. AdaBoost showed a perfect score for malware version 8 (Ransomware). The dataset used for training and testing contained only 170 records for malware version 8. Given the low number of testing instances, it is hard to estimate whether this performance is the same on larger test sets.

High-performance RF for classification target all malware. The best RF model with classification target *all malware* used the Apps feature set, containing 29 app features, and had an F1 score of 0.730, an FPR of 0.013, and an FNR of 0.346 (i.e., TPR of 0.654). Similar results (with no statistically significant difference to the best RF model), are achieved with the Combined feature set (F1 score of 0.722, an FPR of 0.009, and TPR of 0.620). A TPR of 0.62 is relatively low compared to other studies on the dynamic detecting of mobile malware, which showed TPRs between 0.61 and 1 (see Section 2 for a complete comparison). However, in contrast to References [11, 18, 33], with TPRs between 0.82 and 0.97, our study used data from real-life users instead of virtual environments for the training and testing of detection methods. The use of real-life data for the training of models may have led to lower performance due to the additional noise included in our dataset (see paragraph *Overall high false-negative rate* below).

Other studies using real devices for training and testing of models, based their performance on apps running isolated for 10 minutes [11], devices in an idle state [19] or unknown circumstances [13, 28]. One study that used real devices under real-life circumstances for the assessment of their detection method is Reference [37]. In that study, the researchers created a multi-level framework (called MADAM) that showed high performance (TPR 0.97, FPR 0.005). MADAM requires root permissions as it used System Calls for the detection of malware.

In contrast to MADAM, we do not require any root permissions. Furthermore, MADAM used both dynamic and static features, while we used only dynamic features. Last, MADAM is a detection method consisting of multiple architectural blocks that monitor different aspects of the device, in contrast to this present study that does not consist of a complex architecture.

False-negative rate higher than false-positive rate. In this research, the FNR (undetected malicious actions) was overall higher than the FPR (benign actions labelled as malicious). In other studies that we examined during our literature research, the FPR is overall higher than the FNR. This may be due to the difference in the distribution of malicious and benign data points. In our research, the dataset contained 90% benign data points and 10% malicious data points. As most models are biased toward the majority class, the FNR is expected to be higher than the FPR in our case. In other studies the distribution of malicious and benign data points used are often 50/50 or a majority of malicious data points, resulting in equal FPRs and FNRs, or higher FPRs than FNRs. This stresses the importance of presenting both FPR and FNR values in research, which was not done in all examined studies.

Testing times long for KNN. The testing times for KNN were more than 3 s for both the Apps and Combined feature set. This is based on our experimental setup that contained more computing power than the average smartphone. Given that in real-life, malicious, or suspicious data should be tracked and noticed near real time, long testing times may not be practical or feasible. Additionally long testing times stress the battery more as long computations are needed. Therefore, based on testing times, AdaBoost and RF may be preferable when models are running on smartphones.

Low predictability on Spyware SMS and Adware. Performance scores on the detection of malware versions 4 (Spyware SMS) and 6 (Adware) were low (F1 score below 0.5). The dataset used for training and testing contained 190 records for malware version 4. The low performance on the detection of Spyware SMS may, therefore, be due to the small number of training instances. The training and testing dataset contained 7,940 records for malware version 6. Therefore, the low performance is most likely not a result of insufficient training instances, and more research is needed to find the cause for the low performance on detecting Adware.

Lower F1-scores compared to other state-of-the-art dynamic malware detectors. Our best classifier (RF) shows an F1 score of 0.73. This is lower compared to state-of-the-art dynamic malware detectors such as Cai [16] and Saracino [37]. In Reference [16], the performance is high, but the results are taken from malware running in virtual environments for only 10 minutes with pseudo-random user input from Monkey application. This may lower the confidence that the models would perform the same when evaluated on a real device with a real user. Additionally, Cai [16] uses non-hardware features such as Method Calls for their detection method. Our article focuses on hardware features only. Our lower performance may indicate that using hardware features only result in lower overall performance. Saracino [37] uses a combination of hardware features and non-hardware features (System Calls) for their detection methods. As mentioned in Section 2, System Calls require root permissions that our detection method does require. However, still the detection method by Saracino shows a higher performance, indicating again that a combination of hardware and non-hardware features may increase the performance of dynamic malware detectors.

Overall high false-negative rate. In this research, all detection methods showed relatively high FNR (F1 score above 0.3). This implies that a high number of malicious actions are undetected, and may have different causes. Many features in the feature set are influenced by many factors, as the features describe devices that are running multiple applications simultaneously. For instance, the priority, the CPU allocation, and the memory allocation of the malware app depend on other applications running parallel to it. Therefore, the features in the dataset do not solely reflect the (type of) action of the malware app, but also the state of the device at a given moment.

This may result in excessive noise in the sensed data. Another possible cause may be the fact that the influences of malicious and benign actions on the features sensed are similar, and indistinguishable to the classifier.

6 LIMITATIONS

All detection methods in this research are created using the same dataset. The limitations caused by using this particular dataset are listed below:

- (1) All devices in the dataset are Samsung Galaxy S5. Therefore, it is unknown how the models perform when used on other devices.
- (2) All malware types in the dataset were written for the purpose of this research. This limits the findings of the research due to two reasons. First, before the malware probe sent any data to a server, the data were scrambled to ensure the privacy of the volunteers. It is possible that this scrambling influenced the features analysed in the research. As a result, our detection models may be biased toward detecting scrambling actions, limiting its efficacy on wild malware. Second, it is unknown whether real malware executes the same way as custom malware. Although the behaviour of custom malware is based on wild malware, its implementation may differ. It is therefore unknown whether similar detection performance is to be expected on wild malware.
- (3) The dataset contained a low number of data points for malware versions 4 and 6. This limits the conclusions drawn from their results.
- (4) The dataset did not contain any information on the creation times of the malware apps. Therefore, the sustainability metric, as mentioned in recent works [15, 35], could not be included in the analysis. The papers argue that detection methods should be evaluated by taking creation time of apps into consideration. They argue that detection methods should be trained on older apps (e.g., 2018) and tested on newer apps (e.g., 2019). This way detection methods are evaluated on their ability to detect unseen behaviour from newer apps.

The detection methods of this research share characteristics that may limit their performance or applicability. The limitations imposed by this are listed below:

- (1) All detection methods in this research are signature-based, i.e., identify malware based on a pattern of behaviour. In this research, the signature of a malicious action is a record of feature values from a some given feature set. This signature may be different for other (wild) malware types. Therefore, it is unknown whether similar performances are achieved on other (wild) malware types or other versions of the same malware type.
- (2) All detection methods in this research use RFCV as a feature-selection method. This method allows for feature reduction but is a greedy-search strategy, so may be sub-optimal.
- (3) All detection methods using Apps features require data collection of all applications running. In our research, we discarded data from other applications, as we knew which application was malicious. However, in real-life circumstances, the malware application is unknown. Therefore, although Apps features show better performance than Global features, they are more resource-intensive to monitor than Global features. A suggestion for this issue is described in Section 8.

7 CONCLUSION

This research showed (i) what machine-learning classifier is most suitable for detecting Mobile Trojans and (ii) which aspects of a smartphone (features), related to hardware, are most important in detecting Mobile Trojans. The Random Forest classifier showed the highest performance in F1 score, compared with AdaBoost and KNN, when one model is used that is trained on many types of Mobile Trojans. This classifier achieves an F1 score of 0.73 with an FPR of 0.009 and an FNR of 0.380. Random Forest, AdaBoost and K-Nearest Neighbour show high

performances when separate models are trained on each type of Mobile Trojan with an average F1 score above 0.72, FPR below 0.02 and FNR below 0.33. Additionally, on 5 malware types (Spyware, Phishing, Ransomware, Privilege Escalation, DOS) we show high performances (F1 score above 0.963 and FNR below 0.120) making these models relevant for future detection methods that need to detect these malware types.

Multiple dynamic hardware feature sets are examined in this research, making a distinction between global device features and features related to an application. When using one model for the detection of multiple types of Mobile Trojans, 10 app features related to the memory usage, process information, and CPU usage of the app, are sufficient for detecting malicious actions. Additionally, app features, in general, showed the best performance for detecting 7 of 10 malware versions, compared to global device features.

Recent related work performed their research in lablike environments with models based on emulators, simulated user input, and/or short testing durations. To create models that can detect mobile malware threats in real life, we need models assessed on data of real devices rather than data of lablike environments. This research tried to fill this gap by using data of real devices gathered over multiple years with malware based on real malware.

8 FUTURE WORK

Some suggestion for future research are listed below:

- (1) Focus on improving False Negative Rates as this research's detection methods show relatively high FNR. As Random Forest and AdaBoost showed relative high performance, other (boosted) ensemble classifiers can be examined such as GradientBoosting [26].
- (2) Improve the performance of detection methods on new devices and device models if new data become available.
- (3) Improve dynamic detection methods by analysing real wild malware samples to improve the efficacy of the methods on the detection of real malware to overcome the limitation imposed by the use of self-written malware.
- (4) Increase the sample size to improve the efficacy and applicability of the methods. This research analysed 10 different samples of Mobile Trojans.
- (5) Examine whether time-series analysis may improve the detection methods. This research's detection methods analyse dynamic features by individually assessing these values, without considering these value prior in time. Therefore, the absolute values of features are analysed. With time-series analysis, relative values can be used that may improve the detection methods, as is suggested by the results in Reference [33].
- (6) Optimize the resource requirements of detection methods using App features. As described in Section 5, monitoring App features of all applications installed on a device is resource-intensive. Research to decrease these resource requirements is needed. A suggestion may be to only monitor applications that require sensitive permissions, e.g., access to SD card or access to contacts, as was implemented in the framework of Reference [37].
- (7) Improve proposed detection methods by taking sustainability as an additional metric for detectors. Recent works [15, 35] have argued that detection methods should be evaluated by taking creation time of apps into considerations. They argue that detection methods should be trained on older apps (e.g., 2018) and tested on newer apps (e.g., 2019). This way detection methods are evaluated on their ability to detect unseen behaviour from newer apps. The malware applications of our research did not contain information on the creation time, hence this could not be evaluated.

APPENDIX

Table 12. All Features in Different Feature Sets

Global	Apps	Combined
userid	userid	Global plus Apps
uuid	uuid	
version	applicationname	
traffic_mobilerxbytes	cpu_usage	
traffic_mobilerxpackets	packagename	
traffic_mobiletxbytes	packageuid	
traffic_mobiletxpackets	uidrxbytes	
traffic_totalrxbytes	uidrxpackets	
traffic_totalrxpackets	uidtxbytes	
traffic_totaltxbytes	uidtxpackets	
traffic_totaltxpackets	cguest_time	
traffic_totalwifirxbytes	cmajflt	
traffic_totalwifirxpackets	cstime	
traffic_totalwifitxbytes	cutime	
traffic_totalwifitxpackets	dalvikprivatedirty	
traffic_timestamp	dalvikpss	
battery_charge_type	dalvikshareddirty	
battery_current_avg	guest_time	
battery_health	importance	
battery_icon_small	importancereasoncode	
battery_invalid_charger	importancereasonpid	
battery_level	lru	
battery_online	nativeprivatedirty	
battery_plugged	nativepss	
battery_present	nativeshareddirty	
battery_scale	num_threads	
battery_status	otherprivatedirty	
battery_technology	otherpss	
battery_temperature	othershareddirty	
battery_timestamp	pgid	
battery_voltage	pid	
cpuhertz	ppid	
cpu_0	priority	
cpu_1	rss	
cpu_2	rsslim	
cpu_3	sid	
total_cpu	start_time	
totalmemory_freesize	state	
totalmemory_max_size	stime	
totalmemory_total_size	tcomm	
totalmemory_used_size	utime	

(for the meaning of each feature, see the documentation of the original dataset [34]).

Table 13. Features Included in Best Performing Random Forest Models
(Classification Target: All Malware), Ordered by Decreasing Feature Importance

Global	Apps	Combined
traffic_mobilerxbytes	cpu_usage_mor_app	dalvikprivatedirty_mor_app
traffic_mobiletxbytes	uidrxbytes_mor_app	dalvikpss_mor_app
traffic_mobiletxpackets	uidrxpackets_mor_app	importance_mor_app
traffic_totalrxbytes	uidtxbytes_mor_app	num_threads_mor_app
traffic_totalrxpackets	uidtxpackets_mor_app	otherprivatedirty_mor_app
traffic_totaltxbytes	cmajflt_mor_app	otherpss_mor_app
traffic_totaltxpackets	cstime_mor_app	rss_mor_app
traffic_totalwifirxbytes	dalvikprivatedirty_mor_app	stime_mor_app
traffic_totalwifitxbytes	dalvikpss_mor_app	utime_mor_app
traffic_totalwifitxpackets	dalvikshareddirty_mor_app	vsize_mor_app
battery_current_avg	importance_mor_app	
battery_icon_small	importancereasonpid_mor_app	
battery_level	lru_mor_app	
battery_temperature	nativeprivatedirty_mor_app	
battery_voltage	nativepss_mor_app	
cpu_0	nativeshareddirty_mor_app	
cpu_1	num_threads_mor_app	
cpu_2	otherprivatedirty_mor_app	
cpu_3	otherpss_mor_app	
total_cpu	othershareddirty_mor_app	
totalmemory_freesize	pgid_mor_app	
totalmemory_max_size	pid_mor_app	
totalmemory_total_size	ppid_mor_app	
totalmemory_used_size	priority_mor_app	
	rss_mor_app	
	start_time_mor_app	
	stime_mor_app	
	utime_mor_app	
	vsize_mor_app	

(for the meaning of each feature, see the documentation of the original dataset [34]).

Table 14. Features Included in Best Performing AdaBoost Models (Classification Target = all Malware)
Ordered per Feature Importance

Global	Apps	Combined
totalmemory_total_size	rss_mor_app	rss_mor_app
totalmemory_used_size	utime_mor_app	utime_mor_app
battery_voltage	otherprivatedirty_mor_app	dalvikpss_mor_app
battery_temperature	dalvikprivatedirty_mor_app	otherprivatedirty_mor_app
traffic_totalrxbytes	importance_mor_app	importance_mor_app
total_cpu	otherpss_mor_app	otherpss_mor_app
totalmemory_freesize	dalvikpss_mor_app	dalvikprivatedirty_mor_app
traffic_totalrxbytes	num_threads_mor_app	num_threads_mor_app
battery_level	vsize_mor_app	vsize_mor_app
cpu_0	stime_mor_app	stime_mor_app
traffic_totaltxpackets	dalvikshareddirty_mor_app	traffic_totalrxbytes
cpu_2	cpu_usage_mor_app	dalvikshareddirty_mor_app
cpu_3	uidtxbytes_mor_app	totalmemory_used_size
traffic_totalwifitxbytes	othershareddirty_mor_app	cpu_usage_mor_app
battery_current_avg	start_time_mor_app	othershareddirty_mor_app
traffic_totalrxpackets	pid_mor_app	total_cpu
cpu_1	uidrxbytes_mor_app	uidtxbytes_mor_app
traffic_mobilerxbytes	pgid_mor_app	totalmemory_freesize
traffic_totalwifirxbytes	lru_mor_app	traffic_totaltxbytes
traffic_mobiletxbytes	uidtxpackets_mor_app	start_time_mor_app
totalmemory_max_size	ppid_mor_app	battery_voltage
traffic_totalwifitxpackets	nativeprivatedirty_mor_app	totalmemory_total_size
battery_icon_small	priority_mor_app	battery_temperature
traffic_mobiletxpackets	uidrxpackets_mor_app	pid_mor_app
traffic_totalwifirxpackets	cstime_mor_app	uidrxbytes_mor_app
traffic_mobilerxpackets		ppid_mor_app
		battery_level
		cpu_0
		totalmemory_max_size
		traffic_totalwifitxbytes
		traffic_totaltxpackets
		cpu_2
		cpu_3
		lru_mor_app
		uidtxpackets_mor_app
		nativeprivatedirty_mor_app
		priority_mor_app
		cpu_1
		traffic_totalrxpackets
		pgid_mor_app
		traffic_mobiletxbytes
		cstime_mor_app
		traffic_totalwifirxbytes

(for the meaning of each feature, see the documentation of the original dataset [34]).

Table 15. Features Included in Best Performing KNN Models (Classification Target = all Malware) Ordered per Feature Importance

Global	Apps	Combined
totalmemory_total_size	rss_mor_app	rss_mor_app
totalmemory_used_size	utime_mor_app	utime_mor_app
battery_voltage	otherprivatedirty_mor_app	dalvikpss_mor_app
battery_temperature	dalvikprivatedirty_mor_app	otherprivatedirty_mor_app
traffic_totaltxbytes	importance_mor_app	importance_mor_app
total_cpu	otherpss_mor_app	otherpss_mor_app
totalmemory_freesize	dalvikpss_mor_app	dalvikprivatedirty_mor_app
traffic_totalrxbytes	num_threads_mor_app	num_threads_mor_app
battery_level	vsize_mor_app	vsize_mor_app
cpu_0	stime_mor_app	
traffic_totaltxpackets	dalvikshareddirty_mor_app	
cpu_2	cpu_usage_mor_app	
cpu_3	uidtxbytes_mor_app	
traffic_totalwifitxbytes		
battery_current_avg		
traffic_totalrxpackets		
cpu_1		
traffic_mobilerxbytes		
traffic_totalwifirxbytes		
traffic_mobiletbytes		
totalmemory_max_size		
traffic_totalwifitxpackets		
battery_icon_small		
traffic_mobiletbytes		
traffic_totalwifirxpackets		
traffic_mobilerxpackets		

(for the meaning of each feature, see the documentation of the original dataset [34]).

Table 16. All Features and Their Feature Category

Feature	Category	Feature	Category
cguest_time	App_CPU	ppid	App_Process
cpu_usage	App_CPU	sid	App_Process
cstime	App_CPU	start_time	App_Process
cutime	App_CPU	startcode	App_Process
guest_time	App_CPU	state	App_Process
ltrealvalue	App_CPU	tcomm	App_Process
nice	App_CPU	tgpip	App_Process
num_threads	App_CPU	Wchan	App_Process
priority	App_CPU	battery_charge_type	Battery
Processor	App_CPU	battery_current_avg	Battery
rt_priority	App_CPU	battery_health	Battery
stime	App_CPU	battery_icon_small	Battery
utime	App_CPU	battery_invalid_charger	Battery
packagename	App_Info	battery_level	Battery
packageuid	App_Info	battery_online	Battery
version_code	App_Info	battery_plugged	Battery
version_name	App_Info	battery_present	Battery
cmajflt	App_Memory	battery_scale	Battery
cminflt	App_Memory	battery_status	Battery
dalvikprivatedirty	App_Memory	battery_technology	Battery
dalvikpss	App_Memory	battery_temperature	Battery
dalvikshareddirty	App_Memory	battery_timestamp	Battery
lru	App_Memory	battery_voltage	Battery
majflt	App_Memory	btime	CPU
minflt	App_Memory	cpu_0	CPU
nativeprivatedirty	App_Memory	cpu_1	CPU
nativepss	App_Memory	cpu_2	CPU
nativeshareddirty	App_Memory	cpu_3	CPU
otherprivatedirty	App_Memory	cpuhertz	CPU
otherpss	App_Memory	ctxt	CPU
othershareddirty	App_Memory	processes	CPU
rss	App_Memory	procs_blocked	CPU
rsslim	App_Memory	procs_running	CPU
vsiz	App_Memory	tot_idle	CPU
uidrxbytes	App_Network	tot_iowait	CPU
uidrxpackets	App_Network	tot_irq	CPU
uidtxbytes	App_Network	tot_nice	CPU
uidtxpackets	App_Network	tot_softirq	CPU
endcode	App_Process	tot_system	CPU
exit_signal	App_Process	tot_user	CPU
Flags	App_Process	total_cpu	CPU
importance	App_Process	companion_cpu0	IO Interrupts
importancereasoncode	App_Process	companion_sum_cpu123	IO Interrupts
importancereasonpid	App_Process	cpu123_intr_prs	IO Interrupts
pgid	App_Process	cypress_touchkey_cpu0	IO Interrupts
pid	App_Process	cypress_touchkey_sum_cpu123	IO Interrupts

(Part 1 of 2; for the meaning of each feature, see the documentation of the original dataset [34]).

Table 17. All Features and their Feature Category

Feature	Category	Feature	Category
flip_cover_cpu0	IO Interrupts	sreclaimable	Memory
flip_cover_sum_cpu123	IO Interrupts	sunreclaim	Memory
function_call_interrupts_cpu0	IO Interrupts	swpcached	Memory
function_call_interrupts_sum_cpu123	IO Interrupts	swapfree	Memory
home_key_cpu0	IO Interrupts	swaptotal	Memory
home_key_sum_cpu123	IO Interrupts	totalmemory_freesize	Memory
mmsgpio_cpu0	IO Interrupts	totalmemory_max_size	Memory
mmsgpio_sum_cpu123	IO Interrupts	totalmemory_total_size	Memory
pn547_cpu0	IO Interrupts	totalmemory_used_size	Memory
pn547_sum_cpu123	IO Interrupts	unevictable	Memory
sec_headset_detect_cpu0	IO Interrupts	vmallocchunk	Memory
sec_headset_detect_sum_cpu123	IO Interrupts	vmmalloctotal	Memory
slimbus_cpu0	IO Interrupts	vmmallocused	Memory
slimbus_sum_cpu123	IO Interrupts	writeback	Memory
synaptics_rmi4_i2c_cpu0	IO Interrupts	applicationname	Metadata
synaptics_rmi4_i2c_sum_cpu123	IO Interrupts	sherlock_version	Metadata
volume_down_cpu0	IO Interrupts	userid	Metadata
volume_down_sum_cpu123	IO Interrupts	uuid	Metadata
volume_up_cpu0	IO Interrupts	traffic_mobilerxpackets	Network
volume_up_sum_cpu123	IO Interrupts	traffic_mobiletxbytes	Network
wcd9xxx_cpu0	IO Interrupts	traffic_mobiletxpackets	Network
wcd9xxx_sum_cpu123	IO Interrupts	traffic_timestamp	Network
active	Memory	traffic_totalrxbytes	Network
active_anon	Memory	traffic_totalrxpackets	Network
active_file	Memory	traffic_totaltxbytes	Network
anonpages	Memory	traffic_totaltxpackets	Network
buffers	Memory	traffic_totalwifirxbytes	Network
cached	Memory	traffic_totalwifirxpackets	Network
commitlimit	Memory	traffic_totalwifitxbytes	Network
committed_as	Memory	traffic_totalwifitxpackets	Network
dirty	Memory	external_availableblocks	Storage
highfree	Memory	external_availablebytes	Storage
hightotal	Memory	external_blockcount	Storage
inactive	Memory	external_blocksize	Storage
inactive_anon	Memory	external_freeblocks	Storage
inactive_file	Memory	external_freebytes	Storage
kernelstack	Memory	external_totalbytes	Storage
lowfree	Memory	internal_availableblocks	Storage
lowtotal	Memory	internal_availablebytes	Storage
mapped	Memory	internal_blockcount	Storage
Memfree	Memory	internal_blocksize	Storage
memtotal	Memory	internal_freeblocks	Storage
mlocked	Memory	internal_freebytes	Storage
pagetables	Memory	internal_totalbytes	Storage
shmem	Memory	connectedwifi_level	Wifi
slab	Memory	connectedwifi_ssid	Wifi

(Part 2 of 2; for the meaning of each feature, see the documentation of the original dataset [34]).

ACKNOWLEDGMENTS

The authors thank Ben-Gurion University for providing this research's dataset.

REFERENCES

- [1] Google Funf framework. 2013. Retrieved August 30, 2020 from <https://www.funf.org/>.
- [2] VirusShare. 2015. Retrieved August 30, 2021 from <https://web.archive.org/web/20210827173428/https://virusshare.com/>.
- [3] VirusTotal Database. 2015. Retrieved August 30, 2021 from <http://www.virustotal.com>.
- [4] Auto-scaling Sci-kit learn with Apache Spark. 2016. Retrieved August 30, 2021 from <https://databricks.com/blog/2016/02/08/auto-scaling-scikit-learn-with-apache-spark.html>.
- [5] GitHub SherLock. 2017. Retrieved August 30, 2020 from <https://github.com/ymirsky/SherLock>.
- [6] Contagio. 2018. Retrieved August 30, 2021 from <http://contagiominidump.blogspot.com>.
- [7] Open Malware Dataset. 2019. Retrieved August 30, 2021 from <https://web.archive.org/web/20200216063919/http://malwarebenchmark.org/>.
- [8] M. S. Alam and S. T. Vuong. 2013. Random forest classification for detecting android malware. In *Proceedings of the IEEE International Conference on Green Computing and Communications and IEEE Internet of Things and IEEE Cyber, Physical and Social Computing*. 663–669. <https://doi.org/10.1109/GreenCom-iThings-CPSCoM.2013.122>
- [9] Kevin Allix, Tegawendé F. Bissyandé, Jacques Klein, and Yves Le Traon. 2016. Androzoo: Collecting millions of Android apps for the research community. In *Proceedings of the IEEE/ACM 13th Working Conference on Mining Software Repositories (MSR'16)*. IEEE, 468–471.
- [10] Mohammed K. Alzaylaee, Suleiman Y. Yerima, and Sakir Sezer. 2020. DL-droid: Deep learning based Android malware detection using real devices. *Comput. Secur.* 89 (2020), 101663.
- [11] B. Amos, H. Turner, and J. White. 2013. Applying machine learning classifiers to dynamic Android malware detection at scale. In *Proceedings of the 9th International Wireless Communications and Mobile Computing Conference (IWCMC'13)*. 1666–1671. <https://doi.org/10.1109/IWCMC.2013.6583806>
- [12] Daniel Arp, Michael Spreitzenbarth, Malte Hubner, Hugo Gascon, Konrad Rieck, and CERT Siemens. 2014. DREBIN: Effective and explainable detection of android malware in your pocket. In *Proceedings of the Network and Distributed System Security Symposium (NDSS'14)*.
- [13] A. E. Attar, R. Khatoun, and M. Lemercier. 2014. A gaussian mixture model for dynamic detection of abnormal behavior in smartphone applications. In *Proceedings of the Global Information Infrastructure and Networking Symposium (GIIS'14)*. 1–6. <https://doi.org/10.1109/GIIS.2014.6934278>
- [14] Iker Burguera, Urko Zurutuza, and Simin Nadjm-Tehrani. 2011. Crowdroid: Behavior-based malware detection system for Android. In *Proceedings of the 1st ACM Workshop on Security and Privacy in Smartphones and Mobile Devices*. ACM, 15–26.
- [15] Haipeng Cai. 2020. Assessing and improving malware detection sustainability through app evolution studies. *ACM Trans. Softw. Eng. Methodol.* 29, 2 (2020), 1–28.
- [16] Haipeng Cai, Na Meng, Barbara Ryder, and Daphne Yao. 2018. Droidcat: Effective Android malware detection and categorization via app-level profiling. *IEEE Trans. Inf. Forens. Secur.* 14, 6 (2018), 1455–1470.
- [17] Haipeng Cai and Barbara G. Ryder. 2020. A longitudinal study of application structure and behaviors in android. *IEEE Trans. Softw. Eng.* 1 (2020), 1–1.
- [18] Gerardo Canfora, Eric Medvet, Francesco Mercaldo, and Corrado Aaron Visaggio. 2016. Acquiring and analyzing app metrics for effective mobile malware detection. In *Proceedings of the ACM on International Workshop on Security And Privacy Analytics*. ACM, 50–57.
- [19] L. Cavaglione, M. Gaggero, J. F. Lalande, W. Mazurczyk, and M. Urbanski. 2016. Seeing the unseen: Revealing mobile malware hidden communications via energy consumption and artificial intelligence. *IEEE Trans. Inf. Forens. Secur.* 11, 4 (Apr. 2016), 799–810. <https://doi.org/10.1109/TIFS.2015.2510825>
- [20] Yusheng Dai, Hui Li, Yekui Qian, Ruipeng Yang, and Min Zheng. 2019. SMASH: A malware detection method based on multi-feature ensemble learning. *IEEE Access* 7 (2019), 112588–112597.
- [21] Thomas G. Dietterich. 1998. Approximate statistical tests for comparing supervised classification learning algorithms. *Neur. Comput.* 10, 7 (1998), 1895–1923.
- [22] Gianluca Dini, Fabio Martinelli, Andrea Saracino, and Daniele Sgandurra. 2012. MADAM: A multi-level anomaly detector for Android malware. *Comput. Netw. Secur.* (2012), 240–253.
- [23] Bryan Dixon, Shivakant Mishra, and Jeannette Pepin. 2014. Time and location power based malicious code detection techniques for smartphones. In *Proceedings of the IEEE 13th International Symposium on Network Computing and Applications (NCA'14)*. IEEE, 261–268.
- [24] Sumeet Dua and Xian Du. 2016. *Data Mining and Machine Learning in Cybersecurity*. CRC Press, Boca Raton, FL.
- [25] Alberto Ferrante, Eric Medvet, Francesco Mercaldo, Jelena Milosevic, and Corrado Aaron Visaggio. 2016. Spotting the malicious moment: Characterizing malware behavior using dynamic features. In *Proceedings of the 11th International Conference on Availability, Reliability and Security (ARES'16)*. IEEE, 372–381.

- [26] Jerome H. Friedman. 2001. Greedy function approximation: A gradient boosting machine. *Ann. Stat.* (2001), 1189–1232.
- [27] Google. 2017. *Android Security 2016 Year in Review*. Technical Report. Retrieved August 30, 2021 from https://source.android.com/security/reports/Google_Android_Security_2016_Report_Final.pdf.
- [28] Hyo-Sik Ham and Mi-Jung Choi. 2013. Analysis of Android malware detection performance using machine learning classifiers. In *Proceedings of the International Conference on ICT Convergence (ICTC'13)*. IEEE, 490–495.
- [29] Kaspersky. 2017. IT Threat Evolution Q1 2017. Retrieved August 30, 2021 from <https://securelist.com/it-threat-evolution-q1-2017-statistics/78475/>.
- [30] Alejandro Martín, Víctor Rodríguez-Fernández, and David Camacho. 2018. CANDYMAN: Classifying Android malware families by modelling dynamic traces with markov chains. *Eng. Appl. Artif. Intell.* 74 (2018), 121–133.
- [31] L. Massarelli, L. Aniello, C. Ciccotelli, L. Querzoni, D. Ucci, and R. Baldoni. 2017. Android malware family classification based on resource consumption over time. In *Proceedings of the 12th International Conference on Malicious and Unwanted Software (MALWARE'17)*. 31–38.
- [32] L. U. Memon, N. Z. Bawany, and J. A. Shamsi. 2019. A comparison of machine learning techniques for android malware detection using apache spark. *J. Eng. Sci. Technol.* 14, 3 (2019), 1572–1586.
- [33] Jelena Milosevic, Alberto Ferrante, and Mirosław Malek. 2016. MalAware: Effective and efficient run-time mobile malware detector. In *Proceedings of the Dependable, Autonomic and Secure Computing, 14th International Conference on Pervasive Intelligence and Computing, 2nd International Conference on Big Data Intelligence and Computing and Cyber Science and Technology Congress, and 14th International IEEE Cyber Science and Technology Congress (DASC/PiCom/DataCom/CyberSciTech'16)*. IEEE, 270–277.
- [34] Yisroel Mirsky, Asaf Shabtai, Lior Rokach, Bracha Shapira, and Yuval Elovici. 2016. Sherlock vs moriarty: A smartphone dataset for cybersecurity research. In *Proceedings of the ACM Workshop on Artificial Intelligence and Security*. ACM, 1–12.
- [35] Lucky Onwuzurike, Enrico Mariconti, Panagiotis Andriotis, Emiliano De Cristofaro, Gordon Ross, and Gianluca Stringhini. 2019. MaMaDroid: Detecting android malware by building markov chains of behavioral models (extended version). *ACM Trans. Priv. Secur.* 22, 2 (2019), 1–34.
- [36] 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. 2011. Scikit-learn: Machine learning in python. *J. Mach. Learn. Res.* 12 (2011), 2825–2830.
- [37] Andrea Saracino, Daniele Sgandurra, Gianluca Dini, and Fabio Martinelli. 2016. MADAM: Effective and efficient behavior-based Android malware detection and prevention. *IEEE Trans. Depend. Sec. Comput.* 15, 1 (2016), 83–97.
- [38] Asaf Shabtai, Uri Kanonov, Yuval Elovici, Chanan Glezer, and Yael Weiss. 2012. Andromaly: A behavioral malware detection framework for Android devices. *J. Intell. Inf. Syst.* 38, 1 (2012), 161–190.
- [39] S. O’Dea. 2021. Smartphone subscriptions worldwide 2016-2026. Retrieved August 30, 2021 from <https://www.statista.com/statistics/330695/number-of-smartphone-users-worldwide/>.
- [40] Kimberly Tam, Ali Feizollah, Nor Badrul Anuar, Rosli Salleh, and Lorenzo Cavallaro. 2017. The evolution of Android malware and Android analysis techniques. *ACM Comput. Surv.* 49, 4 (2017), 76.
- [41] Sarah Wassermann and Pedro Casas. 2018. BIGMOMAL: Big data analytics for mobile malware detection. In *Proceedings of the Workshop on Traffic Measurements for Cybersecurity*. ACM, 33–39.
- [42] Ping Yan and Zheng Yan. 2017. A survey on dynamic mobile malware detection. *Softw. Qual. J.* (2017), 1–29.
- [43] Yajin Zhou and Xuxian Jiang. 2012. Dissecting android malware: Characterization and evolution. In *Proceedings of the IEEE Symposium on Security and Privacy*. IEEE, 95–109.

Received October 2019; revised July 2021; accepted August 2021