Are Your Training Datasets Yet Relevant?

An Investigation into the Importance of Timeline in Machine Learning-based Malware Detection

Kevin Allix, Tegawendé F. Bissyandé, Jacques Klein, and Yves Le Traon

SnT - University of Luxembourg

Abstract
In this paper, we consider the *relevance of timeline* in the construction of datasets, to highlight its impact on the performance of a machine learning-based malware detection scheme. Typically, we show that simply picking a random set of known malware to train a malware detector, as it is done in many assessment scenarios from the literature, yields *significantly biased* results. In the process of assessing the extent of this impact through various experiments, we were also able to confirm a number of intuitive assumptions about Android malware. For instance, we discuss the existence of Android malware lineages and how they could impact the performance of malware detection in the wild.

1 Introduction

Malware detection is a challenging endeavor in mobile computing, where thousands of applications are uploaded everyday on application markets [1] and often made available for free to end-users. Market maintainers then require efficient techniques and tools to continuously analyze, detect and triage malicious applications in order to keep the market as clean as possible and maintain user confidence. For example, Google has put in place a number of tools and processes in the Google Play official market for Android applications. However, using antivirus software on large datasets from Google reveals that hundreds of suspicious apps are still distributed incognito through this market [2].

Unfortunately, malware pose various threats that cannot be ignored by users, developers and retailers. These threats range from simple user tracking and leakage of personal information [3], to unwarranted premium-rate subscription of SMS services, advanced fraud, and even damaging participation to botnets [4]. To address such threats, researchers and practitioners increasingly turn to new techniques that have been assessed in the literature for malware detection in the wild. Research work have indeed yielded promising approaches for malware detection. A comprehensive survey of various techniques can be found in [5]. Approaches for large-scale detection are often based on Machine learning techniques, which allow to sift through large sets of applications to detect anomalies based on measures of similarity of features [6–14].

To assess malware detection in the wild, the literature resorts to the 10-Fold Cross validation scheme with datasets that we claim are biased and yield biased results. Indeed, various aspects of construction of training datasets are usually overlooked. Among such aspects is the *history aspect* which assumes that the training dataset, which is used for building classifiers, and the test dataset, which is used to assess the performance of the technique, should be *historically coherent*: the former must be historically anterior to the latter. This aspect is indeed a highly relevant constraint for real-world use cases and we feel that evaluation and practical use of state-of-the-art malware detection approaches must follow a process that mimics the history of creation/arrival of applications in markets as well as the history of appearance of malware: *detecting malware before they are publicly distributed in markets is probably more useful than identifying them several months after they have been made available.*

Nevertheless, in the state-of-the-art literature, the datasets of evaluation are borrowed from well-known labelled repositories of apps, such as the Genome project, or constructed randomly, using market-downloaded apps, with the help of Antivirus products. However, the history of creation of the various apps that form the datasets are rarely, if ever, considered, leading to situations where **some items in the training datasets are "*from the future*", i.e., posterior, in the timeline, to items in the tested dataset**. Thus, different research questions are systematically eluded in the discussion of malware detector performance:

RQ-1. Is a randomly sampled training dataset equivalent to a dataset that is historically coherent to the test dataset?

RQ-2. What is the impact of using malware knowledge "from the future" to detect malware in the present?

RQ-3. How can the potential existence of families of malware impact the features that are considered by machine learning classifiers?

RQ-4. How *fresh* must be the apps from the training dataset to yield the best classification results?

RQ-5. Is it sound/wise to account for all known malware to build a training dataset?

This paper. We propose in this paper to investigate the effect of ignoring/considering historical coherence in the selection of training and test datasets for malware detection processes that are built on top of Machine learning techniques. Indeed we note from literature reviews that most authors do not take this into account. Our ultimate aim is thus to provide insights for building approaches that are consistent with the practice of application –including malware– development and registration into markets. To this end, we have devised several typical machine learning classifiers and built a set of features which are textual representations of basic blocks extracted from the Control-Flow Graph of applications' byte-code. Our experiments are also based on a sizeable dataset of about 200,000 Android applications collected from sources that are used by authors of contributions on machine learning-based malware detection.

The contributions of this paper are:

- We propose a thorough study of the history aspect in the selection of training datasets. Our discussions highlight different biases that may be introduced if this aspect is ignored or misused.
- Through extensive experiments with tens of thousands of Android apps, we show the variations that the choice of datasets age can have on the malware

detection output. To the best of our knowledge, we are the first to raise this issue and to evaluate its importance in practice.
- We confirm, or show how our experiments support, various intuitions on Android malware, including the existence of so-called lineages.
- Finally, based on our findings, we discuss (1) the assessment protocols of machine learning-based malware detection techniques, and (2) the design of datasets for training real-world malware detectors.

The remainder of this paper is organized as follows. Section 2 provides some background on machine learning-based malware detection and highlights the associated assumptions on dataset constructions. We also briefly describe our own example of machine-learning based malware detection. Section 3 presents related work to support the ground for our work. Section 4 describes the experiments that we have carried out to answer the research questions, and presents the take-home messages derived from our empirical study. We propose a final discussion on our findings in Section 5 and conclude in Section 6.

2 Preliminaries

The Android mobile platform has now become the most popular platform with estimated hundreds of thousands of apps in the official Google Play market alone and downloads in excess of billions. Unfortunately, as this popularity has been growing, so is malicious software, i.e., malware, targeting this platform. Studies have shown that, on average, Android malware remain unnoticed up to 3 months before a security researcher stumbles on it [15], leaving users vulnerable in the mean time. Security researchers are constantly working to propose new malware detection techniques, including machine learning-based approaches, to reduce this 3-months gap.

Machine Learning: Features & Algorithms: As summarized by Alpaydin, "Machine Learning is programming computers to optimize a performance criterion using example data or past experience" [16]. A common method of learning is known as *supervised* learning, a scheme where the computer is helped through a first step of *training*. The training data consists of Feature Vectors, each associated with a label, e.g., in our case, apps that are already known to be malicious (*malware* class) or benign (*goodware* class). After a run of the learning algorithm, the output is compared to the target output and learning parameters may be corrected according to the magnitude of the error. Consequently, to perform a learning that will allow a *classification* of apps into the malware and goodware classes, the approach must define a correlation measure and a discriminative function. The literature of Android malware detection includes diverse examples of features, such as n-grams of bytecode, API usages, application permission uses, etc. There also exist a variety of classification algorithms, including Support Vector Machine (SVM) [17], the RandomForest ensemble decision-trees algorithm [18], the RIPPER rule-learning algorithm [19] and the tree-based *C4.5* algorithm [20]. In our work, because we focus exclusively on the history aspect, we constrain all aforementioned variables to values that are widely used in the literature, or based on our own experiments which have allowed us to select the most appropriate settings. Furthermore, it is noteworthy that we do not aim for absolute performance, but rather measure performance delta between several approaches of constructing training datasets.

Working Example: We now provide details on the machine-learning approach that will be used as a working example to investigate the importance of history in the selection of training and test datasets. Practically, to obtain the features for our machine-learning processes, we perform static analysis of Android applications' bytecode to extract an abstract representation of the program's control-flow graph (CFG). We obtain a CFG that is expressed as character strings using a method devised by Pouik *et al.* in their work on establishing similarity between Android applications [21], and that is based on a grammar proposed by Cesare and Xiang [22]. The string representation of a CFG is an abstraction of the application's code; it retains information about the *structure* of the code, but discards low-level details such as variable names or register numbers. This property is desirable in the context of malware detection as two variants of a malware may share the same abstract CFG while having different bytecode. Given an application's abstract CFG, we collect all basic blocks that compose it and refer to them as the features of the application. A basic block is a sequence of instructions in the CFG with only one entry point and one exit point. It thus represents the smallest piece of the program that is always executed altogether. By learning from the training dataset, it is possible to expose, if any, the basic blocks that appear statistically more in malware.

The basic block representation used in our approach is a high-level abstraction of the atomic parts of an Android application. A more complete description of this feature set can be found in [23]. For reproducibility purposes, and to allow the research community to build on our experience, the data we used (full feature matrix and labels) is available on request.

Methodology: This study is carried out as a large scale experiment that aims at investigating the extent of the relevance of history in assessing machine learning-based malware detection. This study is important for paving the road to a true success story of trending approaches to Android malware detection. To this end, our work must rely on an extensive dataset that is representative of real-world Android apps and of datasets used in the state-of-the-art literature.

Dataset: To perform this study we collect a large dataset of android apps from various markets: 78,460 (38.04%) apps from `Google Play`, 72,093 (34.96%) from `appchina`, and 55,685 (27.00%) from `Other markets`[1]. A large majority of our dataset comes from `Google Play`, the official market, and `appchina`.

An Android application is distributed as an `.apk` file which is actually a ZIP archive containing all the resources an application needs to run, such as the application binary code and images. An interesting side-effect of this package format is that all the files that makes an application go from the developer's computer to end-users' devices without any modification. In particular, all metadata of the files contained in the .apk package, such as the last modification date, are

[1] *Other markets* include anzhi, 1mobile, fdroid, genome, etc.

Approach	Year	Sources	Historical Coherence
DREBIN [6]	2014	"Genome, Google Play, Chinese and russian markets, VirusTotal	No
[24]	2013	"common Android Markets" for goodware, "public databases of antivirus companies" for malware	No
[13]	2012	Undisclosed	No
DROIDMAT [25]	2012	Contagio mobile for malware, Google Play for goodware	No
[26]	2013	Genome, VirusTotal, Google Play	No
[27]	2013	Contagio mobile and Genome for malware, Undisclosed for goodware	No
[28]	2013	"from official and third party Android markets" for Goodware, Genome for malware	No
[29]	2013	Google Play (labels from 10 commercial Anti virus scanners)	No

Table 1. A selection of Android malware detection approaches

preserved. All bytecode, representing the application binary code, is assembled into a *classes.dex* file that is produced at packaging-time. Thus the last modification date of this file represents the packaging time. In the remainder of this paper, packaging date and compilation date will refer to this date.

To infer the historical distribution of the dataset applications, we rely on compilation date at which the Dalvik[2] bytecode (*classes.dex* file) was produced. We then sent all the app packages to be scanned by virus scanners hosted by VirusTotal [3] . VirusTotal is a web portal which hosts about 40 products from renown anti virus vendors, including McAfee®, Symantec® or Avast®. In this study, an application is labelled as malware if at least one scanner flags it as such.

Machine learning Parameters: In all our experiments, we have used the parameters that provided the best results in a previous large-scale study [23]. Thus, we fixed the number of features to 5,000 and selected the 5,000 features with highest Information Gain values as measured on the training sets. The RandomForest algorithm, as implemented in the Weka[4] Framework, was used for all our experiments.

3 Related Work

In this section, we propose to revisit related work to highlight the importance of our contributions in this paper. We briefly present previous empirical studies and their significance for the malware detection field. Then we go over the literature of malware detection to discuss the assessment protocols.

Empirical studies: Empirical studies have seen a growing interest over the years in the field of computer science. The weight of empirical findings indeed help ensure that research directions and results are in line with practices. This is especially important when assessing the performance of a research approach. A large body of the literature has resorted to extensive empirical studies for devising a reliable experimental protocol [30–32]. Recently, Allix *et al.* have proposed a large-scale empirical studies on the dataset sizes used in the assessment of machine learning-based malware detection approaches [23]. In their work, the authors already questioned the assessment protocols used in the state-of-the-art literature. Guidelines for conducting sound Malware Detection experiments were

[2] Dalvik is the virtual machine running Android apps.
[3] `https://www.virustotal.com`
[4] `http://www.cs.waikato.ac.nz/ml/weka/`

proposed by Rossow *et al* [33]. Our work follows the same objectives, aiming to highlight the importance of building a reliable assessment protocol for research approaches, in order to make them more useful for real-world problems.

In the field of computer security, empirical studies present distinct challenges including the scarcity of data about cybercrimes. We refer the reader to a report by Böhme and Moore [34]. Recently, Visaggio *et al.* empirically assessed different methods used in the literature for detecting obfuscated code [35]. Our work is in the same spirit as theirs, since we also compare different methods of selecting training datasets and draw insights for the research community.

With regards to state-of-the-art literature tackled in this work, a significant number of Machine Learning approaches for malware detection [6, 29, 36–39] have been presented to the research community. The feature set that we use in this paper was evaluated in [23] and achieved better performance than those approaches. Thus, our experiments are based on a sound feature set for malware detection. We further note that in the assessment protocol of all these state-of-the-art approaches, the history aspect was eluded when selecting training sets.

Malware Detection & Assessments: We now review the assessment of malware detection techniques that are based on machine learning. For comparing performances with our own approach, we focus only on techniques that have been applied to the Android ecosystem. In Table 1, we list recent "successful" approaches from the literature of malware detection, and describe the origin of the dataset used for the assessment of each approach. For many of them, the applications are borrowed from known collections of malware samples or from markets such as Google Play. They also often use scanners from VirusTotal to construct the ground truth. In our approach, we have obtained our datasets in the same ways. Unfortunately, to the best of our knowledge and according to their protocol descriptions from the literature, none of the authors has considered clearly ordering the data to take into account the history aspect. It is therefore unfortunate that the high performances recorded by these approaches may never affect the fight against malware in markets.

In the remainder of this section we list significant related work examples, provide details on the size of their dataset and compare them to our history-unaware 10-Fold experiments. None of them has indeed taken into account the history aspect in their assessment protocol. In 2012, Sahs & Khan [13] built an Android malware detector with features based on a combination of Android-specific permissions and a Control-Flow Graph representation. Their classifier was tested with k-Fold [5] cross validation on a dataset of 91 malware and 2 081 goodware. Using permissions and API calls as features, Wu et al. [25] performed their experiments on a dataset of 1 500 goodware and 238 malware. In 2013, Amos et al. [26] leveraged dynamic application profiling in their malware detector. Demme et al. [27] also used dynamic application analysis to perform malware detection with a dataset of 210 goodware and 503 malware. Yerima et al. [28] built malware classifiers based on API calls, external program execution

[5] The value of k used by Sahs & Khan was not disclosed.

and permissions. Their dataset consists of 1 000 goodware and 1 000 malware. Canfora et al. [24] experimented feature sets based on SysCalls and permissions.

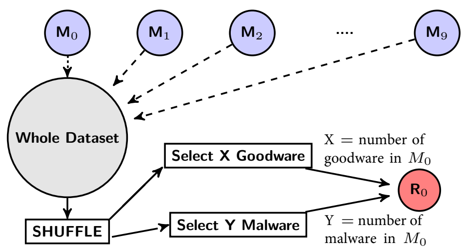

Fig. 1. Process of constructing a random training dataset R_0 for comparison with the training dataset constituted of all data from month M_0

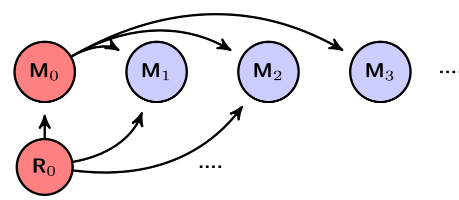

Fig. 2. Classification process: the training dataset is either the dataset of a given month (e.g., M_0) or a random dataset constructing as in Figure 1

4 Experimental Findings

In this section, we report on the experiments that we have conducted, and highlight the findings. First we discuss to what extent it is important that datasets remain historically coherent, as opposed to being selected at random (cf. Section 4.1). This discussion is based on qualitative aspects as well as quantitative evaluation. Second, we conduct experiments that attempt to provide a hint to the existence of lineages in Android malware in Section 4.2. Subsequently, we investigate in Section 4.3 the bias in training with new data for testing with old data, and inversely. Finally, we investigate the limitations of a naive approach which would consist in accumulating information on malware samples as time goes, in the hope of being more inclusive in the detection of malware in the future (cf. Section 4.4).

4.1 History-aware Construction of datasets

As described in Section 2, a key step of machine-learning approaches is the training of classifiers. The construction of the corresponding training dataset is consequently of importance, yet details about how it is achieved are largely missing from the literature, as was shown in Section 3.

There are two common selection patterns for training datasets: (1) use a collected and published dataset of malware, such as Genome, to which one adds a subset of confirmed goodware; (2) build the dataset by randomly picking a subset of goodware and malware from a dataset collected from either an online market or an open repository. Both patterns lead to the same situations: i.e. that *some items in the training dataset may be historically posterior to items in the tested dataset*. In other words, (1) the construction of the training set is equivalent to a random history-unaware selection from a mix of known malware and goodware; and (2) the history of creation/apparition of android applications is not considered as a parameter in assessment experiments, although the practice of malware detection will face this constraint.

Following industry practices, when a newly uploaded set of applications must be analyzed for malware identification, the training datasets that are used are, necessarily, historically anterior to the new set. This constraint is however eluded in the validation of malware detection techniques in the research literature. To clearly highlight the bias introduced by current assessment protocols, we have devised an experiment that compares the performance of the machine learning detectors in different scenarios. The malware detectors are based on classifiers that are built in two distinct settings: either with randomly-constructed training datasets using a process described in Figure 1 or with datasets that respect the history constraint. To reduce the bias between these comparisons, we ensure that the datasets are of identical sizes and with the same class imbalance between goodware and malware. Thus to build a history-unaware dataset R_0 for comparing with training dataset constituted of data from month M_0, we randomly pick within the whole dataset the same numbers of goodware and malware as in M_0. We perform the experiments by training first on M_0 and testing on all following months, then by training on R_0 and testing on all months (cf. Figure 2).

Figure 3 illustrates the results of our experiments. When we randomly select the training dataset from the entire dataset and build classifiers for testing applications regrouped by month, the precision and recall values of the malware detector range between 0.5 and 0.85. The obtained F-Measure is also relatively high and roughly stable. This performance is in line with the performances of state-of-the-art approaches reported in the literature.

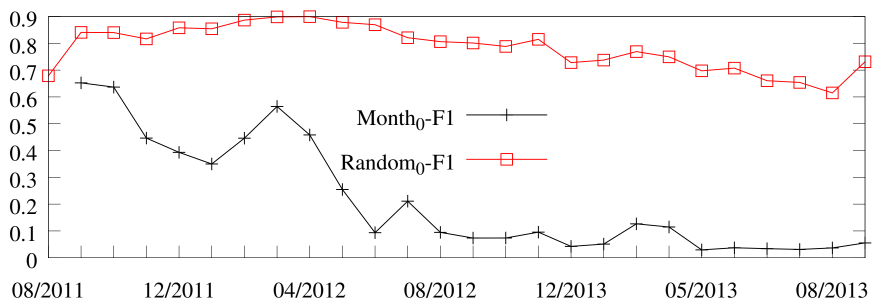

Reading: The $Month_0$ curve shows the F-Measure for a classfier trained on the month 0, while the $Random_0$ curve presents the F-Measure for a classifier built with a training set of same size and same goodware/malware ratio as month 0, but drawn randomly from the whole dataset.

Fig. 3. Performance of malware detectors with history-aware and with history-unaware selection of training datasets

We then proceed to constrain the training dataset to be historically coherent to the test dataset. We select malware and benign apps in the set of apps from a given month, e.g., M_0, as the source of data for building the training dataset for the classification. The tests sets remain the same as in the previous experiments, i.e., the datasets of applications regrouped by month. We observe that as we move away from M_0 to select test data, the performance considerably drops.

We have repeated this experiment, alternatively selecting each different month from our time-line as the month from which we draw the training dataset. Using a training set that is not historically coherent always led to significantly higher performance than using a historically coherent training set.

Finding RQ-1: *Constructing a training dataset that is consistent with the history of apparition of applications yields performances that are significantly worst than what is obtained when simply randomly collecting applications in markets and repositories. Thus,* **without further assessment, state-of-the-art approaches cannot be said to be powerful in real-world settings**.

Finding RQ-2: *With random selections, we allow malware "from the future" to be part of the training sets. This however leads to biased results since the performance metrics are artificially improved.*

4.2 Lineages in Android Malware

Our second round of experiments has consisted in investigating the capabilities of a training dataset to help build classifiers that will remain performant over time. In this step of the study we aim at discovering how the variety of malware is distributed across time. To this end, we consider building training datasets with applications in each month and test the yielded classifiers with the data of each following months.

Figures 4 and 5 provide graphs of the evolution over time of, on the one hand, F-Measure and, on the other hand, Precision of malware detectors that have been built with a training dataset at month M_i and applied on months $M_{k,k>i}$. Disregarding outliers which lead to the numerous abrupt rise and breaks in the curves, the yielded classifiers have, on average, a stable and high Precision, with values around 0.8. This finding suggests that *whatever the combination of training and test dataset months, the built classifiers still allow to identify with good precision the malware whose features have been learnt during training.*

On the other hand, the F-measure performance degrades over time: for a given month M_i whose applications have been used for the training datasets, the obtained classifier is less and less performant in identifying malware in the following months $M_{k,k>i}$. This finding, correlated to the previous one, suggests that, over time, the features that are learnt in the training dataset correspond less and less to all malware when we are in the presence of **lineages** in the Android malware. We define a `lineage` as a set of malware that share the same traits, whether in terms of behavior or of coding attributes. Note that we differentiate the term `lineage` from the term `family` which, in the literature, concern a set of malware that exploit the same vulnerability. *Lineage* is a more general term.

The experiments also highlight the bias introduced when training classifiers with a specific and un-renewed set of malware, such as the Genome dataset, which is widely used. It also confirms why the random selection of malware in the entire time-line as presented in Section 4.1, provides good performances: many lineages are indeed represented in such training datasets, including lineages that should have appeared for the first time in the test dataset.

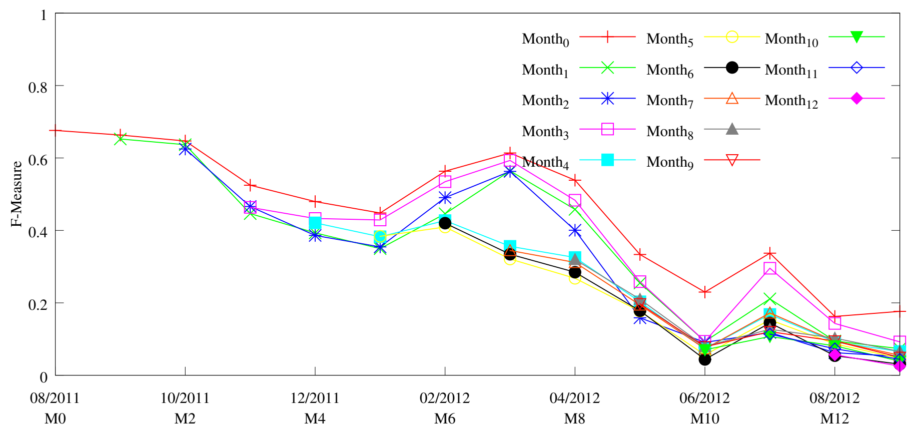

Fig. 4. Performance Evolution of malware detectors over time

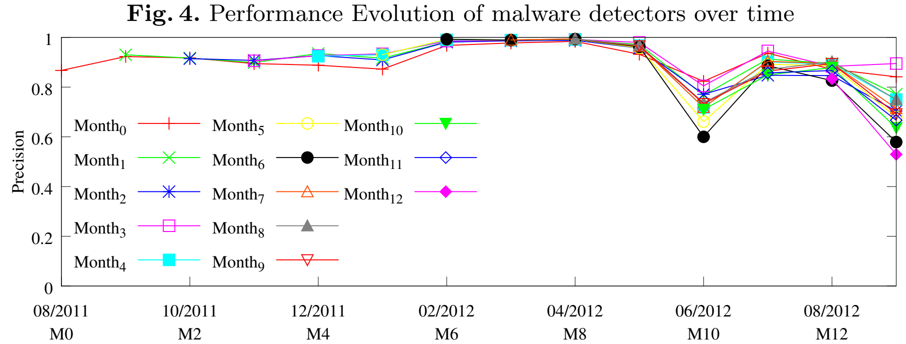

Fig. 5. Evolution of Precision of malware detectors over time

Finding-RQ3: *Android malware is diversified. The existence of lineages complicates malware detection, since training datasets must be regularly updated to include a larger variety of malware lineages representatives.*

4.3 Is knowledge "from the future" the Grail?

Previous experiments have shown that using applications from the entire timeline, without any historical constraint, favorably impacts the performance of malware detectors. We have then proceeded to show that, when the training dataset is too old compared to the test dataset, this performance drops significantly. We now investigate whether training data that are strictly posterior to the test dataset could yield better performance than using data that are historically anterior (coherent). Such a biased construction of datasets is not fair when the objective is to actively keep malicious apps from reaching the public domain. However, such a construction can be justified by the assumption that the present might always contain representatives of malware lineages that have appeared in the past.

In the Android ecosystem, thousands of applications are created weekly by developers. Most of them, including malware from new lineages, cannot be thoroughly checked. Nevertheless, after some time, antivirus vendors may identify

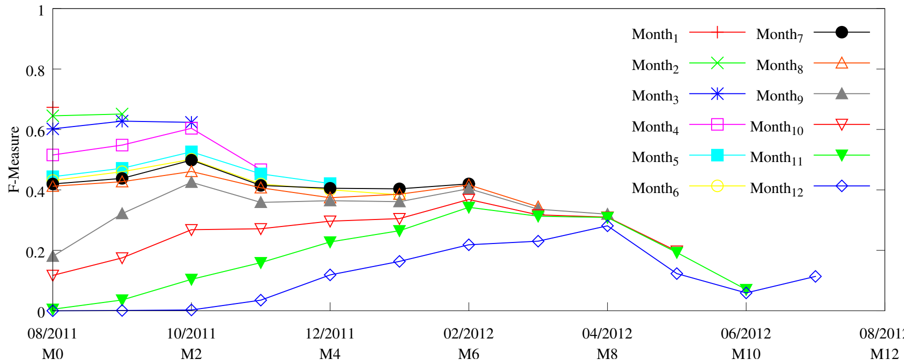

Fig. 6. Performance of malware detectors when using recent data to test on old datasets

the new malware. Machine-learning processes can thus be used to automate a large-scale identification of malware in applications that have been made available for some time. Figure 6 depicts the F-Measure performance evolution of the malware detectors: for each month M_i, that is used for training, the obtained classifiers are used to predict malware in the previous months $\mathrm{M}_{k,k<i}$. Overall, the performance is dropping significantly with the time difference between test and training datasets.

Finding-RQ4: *Apps, including malware, used for training in machine learning-based malware detection must be historically close to the target dataset that is tested. Older training datasets cannot account for all malware lineages, and newer datasets do not contain enough representatives of malware from the past.*

4.4 Naive Approaches to the Construction of Training Datasets

Given the findings of our study presented in previous sections, we investigate through extensive experiments the design of a potential research approach for malware detection which will be in line with the constraints of industry practices. At a given time t, one can only build classifiers using datasets that are anterior to t. Nevertheless, to improve our chances of maintaining performance, two protocols can be followed:

(1) *Keep renewing the training dataset entirely to stay historically close to the target dataset of test. This renewal process must however be automated to remain realistic*: In this scenario, we assume that a bootstrap step is achieved with antivirus products at month M_0 to provide a first reliable training dataset. The malware detection system is then on its own for the following months. Thus, the classification that is obtained on month M_1, using month M_0 for training, will be used "as is" to train the classifiers for testing applications data of month M_2. This system is iterated until month M_n as depicted in Figure 7, meaning that, once it is bootstrapped, the detection system is automated and only relies on its test results to keep training new classifiers. In practice, such an approach makes sense due to the high precision values recorded in previous experiments.

(2) *Include greedily the most knowledge one can collect on malware lineages*: This scenario is also automated and requires bootstrapping. However, instead of renewing the training dataset entirely each month, new classification results are added to the existing training dataset and used to build classifiers for the following month.

Figure 8 shows that the F-measure performance is slightly better for scenario 2. The detailed graphs show that, in the long run, the Recall in scenario 2 is indeed better while the Precision is lower than in scenario 1. In summary, these two scenarios exhibit different trade-offs between Precision and Recall in the long run: Scenario 1 manages to pinpoint a small number of malware with good precision while scenario 2 instead finds more malware at the cost of a higher false-positive rate.

While of little use in isolation, those scenarios provide insights through empirical evidence on how machine learning-based malware detection systems should consider the construction of training sets.

Finding-RQ5: *Maintaining performance of malware detectors cannot be achieved by simply adding/renewing information in training datasets based on the output of previous runs. However, these scenarios have shown interesting impact on performance evolution over time, and must be further investigated to identify the right balance.*

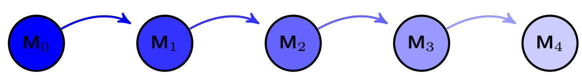

Fig. 7. Using classification results of M_{n-1} as training dataset for testing M_n

5 Insights and Future work

Findings (1) History constraints must not be eluded in the construction of training datasets of machine learning-based malware detectors. Indeed, they introduce significant bias in the interpretation of the performance of malware classifiers. (2) There is a need for building a reliable, and continuously updated, benchmark for machine learning-based malware detection approaches. We make available, upon request, the version we have built for this work. Our benchmark dataset contains about 200,000 Android applications spanning 2 years of historical data of Android malware.

Insights (1) Machine-learning cannot assure the identification of an entirely new lineage of malware that is not represented in the training dataset. Thus, there is need to regularly seed the process with outside information, such as from antivirus vendors, of new lineages of malware. (2) In real world settings, practitionners cannot be presented with a reliable dataset for training. Indeed, most malware are discovered, often manually, by antivirus vendors far later after they have been available to end-users [15]. Large-scale ML-based malware detection must therefore be used to automate the discovery of variants of malware which have been authenticated in a separate process.

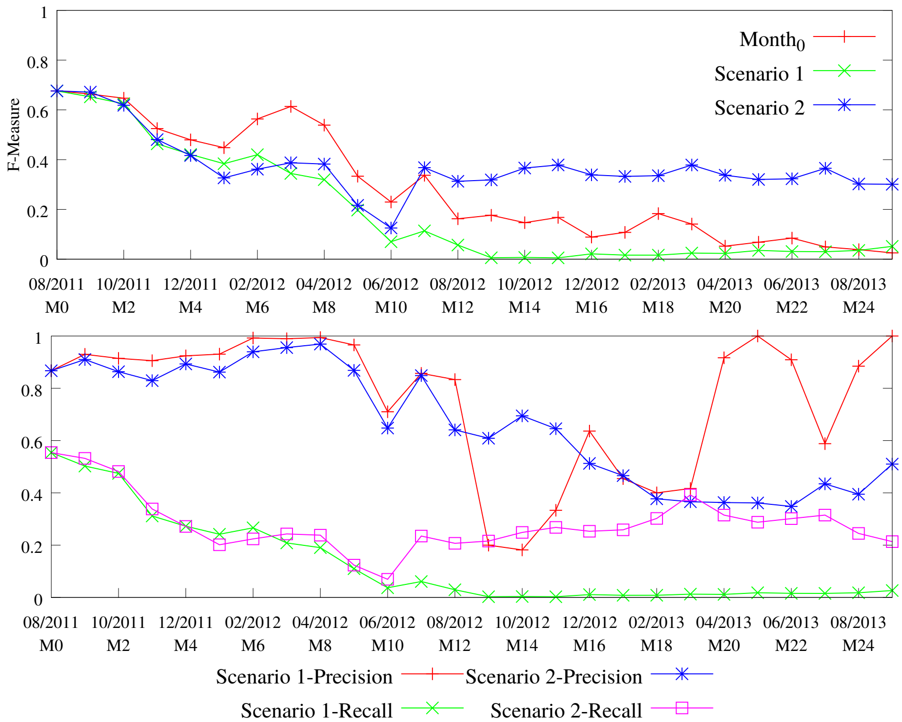

Fig. 8. Comparing two naive approaches

Threat to Validity To perform this study, we have considered a unique use-case scenario for using machine learning-based malware detection. This scenario consists in Actively preventing malware from reaching markets and is extremely relevant to most real-world constraints. Indeed, in practice, it is important to keep the window of opportunity very narrow. Thus, to limit the number of infected devices, Android malware must be detected as they arrive in the market. It is therefore important that state-of-the-art approaches be properly assessed, taking into account history constraints.

There is however a second use-case scenario which concerns online repositories for research and would consist on *cleaning such repositories regularly*. In this scenario, repositories maintainers attempt to filter malicious apps once a new kind of malware has been discovered. In such a context, practitionners can afford to wait for a long time before building relevant classifiers for identifying malware that have been since in the repository. Nevertheless, such repositories are generally of reasonable size and can be scanned manually and with the help of anti virus products.

There is a possibility that the results obtained in this paper would not be reproduced with either a different feature set and/or a different dataset. Nonetheless, we have no reason to believe that the way the dataset was collected induced any bias.

Future work (1) Building on the insights of our experiments, we plan to investigate how to maintain the performance of machine learning-based malware detectors until antivirus vendors can detect a new strain of malware. This research direction could help bring research work to be applied on real-world processes, in conjunction with antivirus products which are still widely used, although they do not scale to the current rates of malware production. (2) To account for the evolution of representations of malware lineages in training datasets over time, we plan to investigate a multi-classifier approach, each classifier being trained with more or less outdated data and weighted accordingly. A first challenge will be on how to infer or automate the choice of weights for different months in the timeline to build the most representative training dataset.

6 Conclusion

Given the steady increase in the adoption of smartphones worldwide, and the growth of application development for such devices, it is becoming important to protect users from the damages of malicious apps. Malware detection has thus been in recent years the subject of renewed interest, and researchers are investigating scalable techniques to spot and filter out apps with malicious traits among thousands of benign apps.

However, more than in other fields, research in computer security must yield techniques and approaches that are truly usable in real-world settings. To that end, assessment protocols of malware detection research approaches must reflect the practice and constraints observed by market maintainers and users. Through this empirical study we aim to prevent security research from producing approaches and techniques that are not in line with reality. Furthermore, given the performances reported in state-of-the-art literature of malware detection, while market maintainers still struggle to keep malware out of markets, it is important to clear the research field by questioning current assessment protocols.

In this paper, we have investigated the relevance of history in the selection of assessment datasets. We have performed large-scale experiments to highlight the different bias that can be exhibited by different scenarios of dataset selection. Our main conclusion is that the assessment protocol used for current approaches in the state-of-the-art literature is far from the reality of a malware detection practice for keeping application markets clean. We have further investigated naive approaches to training dataset construction and drawn insights for future work by the research community.

References

1. AppBrain: Number of available android applications. `http://www.appbrain.com/stats/number-of-android-apps` Accessed: 2013-09-09.
2. Allix, K., Bissyandé, T.F., Klein, J., Le Traon, Y.: A forensic analysis of android malware - how is malware written and how it could be detected? In: Proceedings of the IEEE Computer Software and Applications Conference. COMPSAC (2014)

3. Enck, W., Octeau, D., McDaniel, P., Chaudhuri, S.: A study of android application security. In: Proceedings of the 20th USENIX conference on Security. SEC'11, San Francisco, CA (2011)
4. Pieterse, H., Olivier, M.: Android botnets on the rise: Trends and characteristics. In: Proceedings of the Conference on Information Security for South Africa. ISSA (2012)
5. Idika, M.: A survey of malware detection techniques. Technical report, Purdue University (Feb. 2007)
6. Arp, D., Spreitzenbarth, M., Hübner, M., Gascon, H., Rieck, K.: Drebin: Effective and explainable detection of android malware in your pocket. In: Proceedings of the Network and Distributed System Security Symposium. NDSS (2014)
7. Chau, D.H., Nachenberg, C., Wilhelm, J., Wright, A., Faloutsos, C.: Polonium: Tera-scale graph mining for malware detection. In: ACM SIGKDD Conference on Knowledge Discovery and Data Mining. (2010)
8. Boshmaf, Y., Ripeanu, M., Beznosov, K., Zeeuwen, K., Cornell, D., Samosseiko, D.: Augur: Aiding malware detection using large-scale machine learning. In: Proceedings of the 21st Usenix Security Symposium (Poster session). (Aug 2012)
9. Su, X., Chuah, M., Tan, G.: Smartphone dual defense protection framework: Detecting malicious applications in android markets. In: Eighth IEEE International Conference on Mobile Ad-hoc and Sensor Networks. MSN (2012)
10. Henchiri, O., Japkowicz, N.: A feature selection and evaluation scheme for computer virus detection. In: Proceedings of the Sixth International Conference on Data Mining. ICDM '06, Washington, DC, USA (2006)
11. Kolter, J.Z., Maloof, M.A.: Learning to detect and classify malicious executables in the wild. Journal of Machine Learning Research **7** (December 2006)
12. Zhang, B., Yin, J., Hao, J., Zhang, D., Wang, S.: Malicious codes detection based on ensemble learning. In: Proceedings of the 4th international conference on Autonomic and Trusted Computing. ATC, Hong Kong, China (2007)
13. Sahs, J., Khan, L.: A machine learning approach to android malware detection. In: Proceedings of the IEEE European Intelligence and Security Informatics Conference. EISIC (2012)
14. Perdisci, R., Lanzi, A., Lee, W.: Mcboost: Boosting scalability in malware collection and analysis using statistical classification of executables. In: Proceedings of the Annual Computer Security Applications Conference. ACSAC (2008)
15. Apvrille, A., Strazzere, T.: Reducing the window of opportunity for android malware gotta catch 'em all. Journal of Computer Virology **8**(1-2) (May 2012)
16. Alpaydin, E.: Introduction to Machine Learning. 2nd edn. The MIT Press (2010)
17. Cortes, C., Vapnik, V.: Support-vector networks. Machine Learning **20**(3) (1995)
18. Breiman, L.: Random forests. Machine learning **45**(1) (2001)
19. Cohen, W.W.: Fast effective rule induction. In: Proceedings of the International Machine Learning Conference, Morgan Kaufmann Publishers, Inc. (1995)
20. Quinlan, J.R.: C4. 5: programs for machine learning. Volume 1. Morgan kaufmann (1993)
21. Pouik, G0rfi3ld: Similarities for fun & profit. Phrack **14**(68) (April 2012) `http://www.phrack.org/issues.html?id=15&issue=68`.
22. Cesare, S., Xiang, Y.: Classification of malware using structured control flow. In: Proceedings of the Eighth Australasian Symposium on Parallel and Distributed Computing - Volume 107. AusPDC '10, Darlinghurst, Australia, Australia, Australian Computer Society, Inc. (2010)
23. Allix, K., Bissyandé, T.F., Jerome, Q., Klein, J., State, R., Le Traon, Y.: Empirical assessment of machine learning-based malware detectors for android: Measuring the gap between in-the-lab and in-the-wild validation scenarios. Empirical Software Engineering (2015) to be published.

24. Canfora, G., Mercaldo, F., Visaggio, C.A.: A classifier of malicious android applications. In: Proceedings on the 8th Conference on Availability, Reliability and Security (ARES). (2013)
25. Wu, D.J., Mao, C.H., Wei, T.E., Lee, H.M., Wu, K.P.: Droidmat: Android malware detection through manifest and api calls tracing. In: Proceedings of the 7th Asia Joint Conference on Information Security. AsiaJCIS (2012)
26. Amos, B., Turner, H., White, J.: Applying machine learning classifiers to dynamic android malware detection at scale. In: Proceedings of 9th International Wireless Communications and Mobile Computing Conference. IWCMC (2013)
27. Demme, J., Maycock, M., Schmitz, J., Tang, A., Waksman, A., Sethumadhavan, S., Stolfo, S.: On the feasibility of online malware detection with performance counters. In: Proceedings of the 40th Annual International Symposium on Computer Architecture. ISCA '13, New York, NY, USA, ACM (2013)
28. Yerima, S., Sezer, S., McWilliams, G., Muttik, I.: A new android malware detection approach using bayesian classification. In: Proceedings of the 27th IEEE International Conference on Advanced Information Networking and Applications. AINA (2013)
29. Gascon, H., Yamaguchi, F., Arp, D., Rieck, K.: Structural detection of android malware using embedded call graphs. In: Proceedings of the 2013 ACM workshop on Artificial intelligence and security. AISec (2013)
30. Bissyandé, T.F., Thung, F., Wang, S., Lo, D., Jiang, L., Réveillère, L.: Empirical Evaluation of Bug Linking. In: 17th European Conference on Software Maintenance and Reengineering (CSMR 2013), Genova, Italy (March 2013)
31. Jones, J.A., Harrold, M.J.: Empirical evaluation of the tarantula automatic fault-localization technique. In: Proceedings of the 20th IEEE/ACM international Conference on Automated software engineering. ASE, ACM (2005)
32. Hutchins, M., Foster, H., Goradia, T., Ostrand, T.: Experiments of the effectiveness of dataflow-and controlflow-based test adequacy criteria. In: Proceedings of the 16th international conference on Software engineering. ICSE (1994)
33. Rossow, C., Dietrich, C., Grier, C., Kreibich, C., Paxson, V., Pohlmann, N., Bos, H., van Steen, M.: Prudent practices for designing malware experiments: Status quo and outlook. In: Security and Privacy (SP), 2012 IEEE Symposium on. (May 2012)
34. Böhme, R., Moore, T.: Challenges in empirical security research. Technical report, Singapoore Management University (2012)
35. Visaggio, C.A., Pagin, G.A., Canfora, G.: An empirical study of metric-based methods to detect obfuscated code. International Journal of Security & Its Applications **7**(2) (2013)
36. Aafer, Y., Du, W., Yin, H.: Droidapiminer: Mining api-level features for robust malware detection in android. In: Proceedings of the International Conference on Security and Privacy in Communication Networks. SecureComm (2013)
37. Barrera, D., Kayacik, H., van Oorschot, P., Somayaji, A.: A methodology for empirical analysis of permission-based security models and its applications to android. In: Proceedings of ACM Conference on Computer and Communications Security. CCS (2010)
38. Chakradeo, S., Reaves, B., Traynor, P., Enck, W.: Mast: Triage for market-scale mobile malware analysis. In: Proceedings of ACM Conference on Security and Privacy in Wireless and Mobile Networks. WISEC (2013)
39. Peng, H., Gates, C.S., Sarma, B.P., Li, N., Qi, Y., Potharaju, R., Nita-Rotaru, C., Molloy, I.: Using probabilistic generative models for rangking risks of android apps. In: Proceedings of ACM Conference on Computer and Communications Security. CCS (2012)