

Survey of Machine Learning Techniques for Malware Analysis

Daniele Ucci^a, Leonardo Aniello^b, Roberto Baldoni^a

^a*Research Center of Cyber Intelligence and Information Security, “La Sapienza” University of Rome*

^b*Cyber Security Research Group, University of Southampton*

Abstract

Coping with malware is getting more and more challenging, given their relentless growth in complexity and volume. One of the most common approaches in literature is using machine learning techniques, to automatically learn models and patterns behind such complexity, and to develop technologies to keep pace with malware evolution. This survey aims at providing an overview on the way machine learning has been used so far in the context of malware analysis in Windows environments, i.e. for the analysis of Portable Executables. We systematize surveyed papers according to their objectives (i.e., the *expected output*), what information about malware they specifically use (i.e., the *features*), and what machine learning techniques they employ (i.e., what *algorithm* is used to process the input and produce the output). We also outline a number of issues and challenges, including those concerning the used *datasets*, and identify the main current topical trends and how to possibly advance them. In particular, we introduce the novel concept of *malware analysis economics*, regarding the study of existing trade-offs among key metrics, such as analysis accuracy and economical costs.

Keywords: portable executable, malware analysis, machine learning, benchmark, malware analysis economics

Email addresses: ucci@diag.uniroma1.it (Daniele Ucci), l.aniello@soton.ac.uk (Leonardo Aniello), baldoni@diag.uniroma1.it (Roberto Baldoni)

1. Introduction

Despite the significant improvement of cyber security mechanisms and their continuous evolution, malware are still among the most effective threats in the cyber space. Malware analysis applies techniques from several different fields, such as program analysis and network analysis, for the study of malicious samples to develop a deeper understanding on several aspects, including their behaviour and how they evolve over time. Within the unceasing arms race between malware developers and analysts, each advance in security technology is usually promptly followed by a corresponding evasion. Part of the effectiveness of novel defensive measures depends on what properties they leverage on. For example, a detection rule based on the MD5 hash of a known malware can be easily eluded by applying standard techniques like obfuscation, or more advanced approaches such as *polymorphism* or *metamorphism*. For a comprehensive review of these techniques, refer to Ye *et al.* [1]. These methods change the binary of the malware, and thus its hash, but leave its behaviour unmodified. On the other side, developing detection rules that capture the semantics of a malicious sample is much more difficult to circumvent, because malware developers should apply more complex modifications. A major goal of malware analysis is to capture additional properties to be used to improve security measures and make evasion as hard as possible. Machine learning is a natural choice to support such a process of knowledge extraction. Indeed, many works in literature have taken this direction, with a variety of approaches, objectives and results.

This survey aims at reviewing and systematising existing literature where machine learning is used to support malware analysis of Windows executables, i.e. Portable Executables (PEs). The intended audience of this survey includes any security analysts, i.e. security-minded reverse engineer or software developer, who may benefit from applying machine learning to automate part of malware analysis operations and make the workload more tractable. Although mobile malware represents an ever growing threat, Windows largely remains the preferred target [2] among all the existing platforms. Malware analysis

techniques for PEs are slightly different from those for Android apps because there are significant dissimilarities on how operating system and applications work. As a matter of fact, literature papers on malware analysis commonly point out what specific platform they target, so we specifically focus on works that consider the analysis of PEs. 64 recent papers have been selected on the basis of their bibliographic significance, reviewed and systematised according to a taxonomy with three fundamental dimensions: (i) the specific *objective* of the analysis, (ii) what types of *features* extracted from PEs they consider and (iii) what machine learning *algorithms* they use. We distinguish three main objectives: *malware detection*, *malware similarity analysis* and *malware category detection*. PE features have been grouped in eight types: *byte sequences*, *APIs/System calls*, *opcodes*, *network*, *file system*, *CPU registers*, *PE file characteristics* and *strings*. Machine learning algorithms have been categorized depending on whether the learning is *supervised*, *unsupervised* or *semi-supervised*. The characterisation of surveyed papers according to such taxonomy allows to spot research directions that have not been investigated yet, such as the impact of particular combination of features on analysis accuracy. The analysis of such a large literature leads to single out three main issues to address. The first concerns overcoming modern anti-analysis techniques such as encryption. The second regards the inaccuracy of malware behaviour modelling due to the choice of what operations of the sample are considered for the analysis. The third is about the obsolescence and unavailability of the datasets used in the evaluation, which affect the significance of obtained results and their reproducibility. In this respect, we propose a few guidelines to prepare suitable benchmarks for malware analysis through machine learning. We also identify a number of topical trends that we consider worth to be investigated more in detail, such as malware attribution and triage. Furthermore, we introduce the novel concept of *malware analysis economics*, regarding the existing trade-offs between analysis accuracy, time and cost, which should be taken into account when designing a malware analysis environment.

The novel contributions of this work are

- the definition of a taxonomy to synthesise the state of the art on machine learning for malware analysis of PEs;
- a detailed comparative analysis of existing literature on that topic, structured according to the proposed taxonomy, which highlights possible new research directions;
- the determination of present main issues and challenges on that subject, and the proposal of high-level directions to investigate to overcome them;
- the identification of a number of topical trends on machine learning for malware analysis of PEs, with general guidelines on how to advance them;
- the definition of the novel concept of malware analysis economics.

The rest of the paper is structured as follows. Related work are described in Section 2. Section 3 presents the taxonomy we propose to organise reviewed malware analysis approaches based on machine learning, which are then characterised according to such a taxonomy in Section 4. From this characterisation, current issues and challenges are pointed out in Section 5. Section 6 highlights topical trends and how to advance them. Malware analysis economics is introduced in Section 7. Finally, conclusions and future works are presented in Section 8.

2. Related Work

Other academic works have already addressed the problem of surveying contributions on the usage of machine learning techniques for malware analysis. The survey written by Shabtai *et al.* [3] is the first one on this topic. It specifically deals with how classifiers are used on static features to detect malware. As most of the other surveys mentioned in this subsection, the main difference with our work is that our scope is wider as we target other objectives besides malware detection, such as similarities analysis and category detection. Furthermore, a novel contribution we provide is the idea of malware economics, which is not

mentioned by any related work. Also in [4], the authors provide a comparative study on papers using pattern matching to detect malware, by reporting their advantages, disadvantages and problems. Souri and Hosseini [5] proposes a taxonomy of malware detection approaches based on machine learning. In addition to consider detection only, their work differs from ours because they do not investigate what features are taken into account. LeDoux and Lakhotia [6] describe how machine learning is used for malware analysis, whose end goal is defined there as “automatically detect malware as soon as possible, remove it, and repair any damage it has done”.

Bazrafshan *et al.* [7] focus on malware detection and identify three main methods for detecting malicious software, i.e. based on signatures, behaviours and heuristics, the latter using also machine learning techniques. They also identify what classes of features are used by reviewed heuristics for malware detection, i.e. API calls, control flow graphs, n-grams, opcodes and hybrid features. In addition to going beyond malware detection, we propose a larger number of feature types, which reflects the wider breadth of our research.

Basu *et al.* [8] examine different works relying on data mining and machine learning techniques for the detection of malware. They identify five types of features: API call graph, byte sequence, PE header and sections, opcode sequence frequency and kernel, i.e. system calls. In our survey we establish more feature types, such as strings, file system and CPU registers. They also compare surveyed papers by used features, used dataset and mining method.

Ye *et al.* [1] examine different aspects of malware detection processes, focusing on feature extraction/selection and classification/clustering algorithms. Also in this case, our survey looks at a larger range of papers by also including many works on similarity analysis and category detection. They also highlight a number of issues, mainly dealing with machine learning aspects (i.e. incremental learning, active learning and adversarial learning). We instead look at current issues and limitations from a distinct angle, indeed coming to a different set of identified problems that complement theirs. Furthermore, they outline several trends on malware development, while we rather report on trends about machine

learning for malware analysis, again complementing their contributions.

Barriga and Yoo [9] briefly survey literature on malware detection and malware evasion techniques, to discuss how machine learning can be used by malware to bypass current detection mechanisms. Our survey focuses instead on how machine learning can support malware analysis, even when evasion techniques are used. Gardiner and Nagaraja [10] concentrate their survey on the detection of command and control centres through machine learning.

3. Taxonomy of Machine Learning Techniques for Malware Analysis

This section introduces the taxonomy on how machine learning is used for malware analysis in the reviewed papers. We identify three major dimensions along which surveyed works can be conveniently organised. The first one characterises the final *objective* of the analysis, e.g. malware detection. The second dimension describes the *features* that the analysis is based on in terms of how they are extracted, e.g. through dynamic analysis, and what features are considered, e.g. CPU registers. Finally, the third dimension defines what type of *machine learning algorithm* is used for the analysis, e.g. supervised learning.

Figure 1 shows a graphical representation of the taxonomy. The rest of this section is structured according to the taxonomy. Subsection 3.1 describes in details the *objective* dimension, *features* are pointed out in subsection 3.2 and machine learning algorithms are reported in subsection 3.3.

3.1. Malware Analysis Objectives

Malware analysis, in general, demands for strong detection capabilities to find matches with the knowledge developed by investigating past samples. Anyway, the final goal of searching for those matches differs. For example, a malware analyst may be specifically interested in determining whether new suspicious samples are malicious or not, while another may be rather inspecting new malware looking for what family they likely belong to. This subsection details the analysis goals of the surveyed papers, organized in three main objectives:

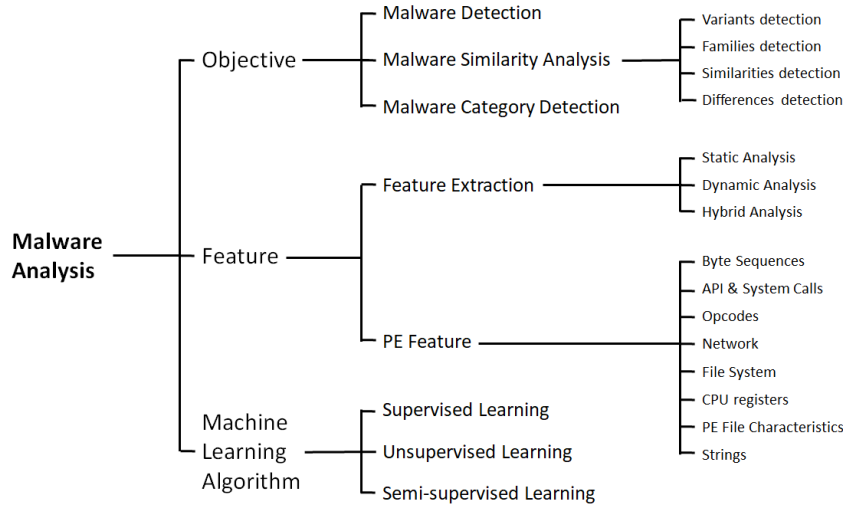


Figure 1: Taxonomy of machine learning techniques for malware analysis

malware detection (§ 3.1.1), malware similarity analysis (§ 3.1.2) and malware category detection (§ 3.1.3).

3.1.1. Malware Detection

The most common objective in the context of malware analysis is detecting whether a given sample is malicious. This objective is also the most important because knowing in advance that a sample is dangerous allows to block it before it becomes harmful. Indeed, the majority of reviewed works has this as main goal [11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36]. Depending on what machine learning technique is used, the generated output can be provided with a confidence value that can be used by analysts to understand if a sample needs further inspection.

3.1.2. Malware Similarity Analysis

Another relevant objective is spotting similarities among malware, for example to understand how novel samples differ from previous, known ones. We find four slightly different versions of this objective: *variants detection*, *families detection*, *similarities detection* and *differences detection*.

Variants Detection. Developing variants is one of the most effective and cheapest strategies for an attacker to evade detection mechanisms, while reusing as much as possible already available codes and resources. Recognizing that a sample is actually a variant of a known malware prevents such strategy to succeed, and paves the way to understand how malware evolve over time through the development of new variants. Also this objective has been deeply studied in literature, and several reviewed papers target the detection of variants. Given a malicious sample m , variants detection consists in selecting from the available knowledge base the samples that are variants of m [37, 30, 38, 39, 40, 41]. Considering the huge number of malicious samples received daily from major security firms, recognising variants of already known malware is crucial to reduce the workload for human analysts.

Families Detection. Given a malicious sample m , families detection consists in selecting from the available knowledge base the families that m likely belongs to [42, 43, 44, 45, 46, 47, 48, 49, 50, 31, 51, 52, 53, 54, 36]. In this way, it is possible to associate unknown samples to already known families and, by consequence, provide an added-value information for further analyses.

Similarities Detection. Analysts can be interested in identifying the specific similarities and differences of the binaries to analyse with respect to those already analysed. Similarities detection consists in discovering what parts and aspects of a sample are similar to something that has been already examined in the past. It enables to focus on what is really new, and hence to discard the rest as it does not deserve further investigation [55, 56, 57, 58, 59].

Differences Detection. As a complement, also identifying what is different from everything else already observed in the past results worthwhile. As a matter of fact, differences can guide towards discovering novel aspects that should be analysed more in depth [56, 60, 57, 58, 61, 62].

3.1.3. Malware Category Detection

Malware can be categorized according to their prominent behaviours and objectives. They can be interested in spying on users' activities and stealing their sensitive information (i.e., *spyware*), encrypting documents and asking for a ransom (i.e., *ransomware*), or gaining remote control of an infected machine (i.e., *remote access toolkits*). Using these categories is a coarse-grained yet significant way of describing malicious samples [63, 64, 65, 66, 32, 67]. Although cyber security firms have not still agreed upon a standardized taxonomy of malware categories, effectively recognising the categories of a sample can add valuable information for the analysis.

3.2. Malware Analysis Features

This subsection deals with the features of samples that are considered for the analysis. How features are extracted from executables is reported in subsection 3.2.1, while subsection 3.2.2 details which specific features are taken into account.

3.2.1. Feature Extraction

The information extraction process is performed through either static or dynamic analysis, or a combination of both, while examination and correlation are carried out by using machine learning techniques. Approaches based on static analysis look at the content of samples without requiring their execution, while dynamic analysis works by running samples to examine their behaviour. Several techniques can be used for dynamic malware analysis. *Debuggers* are used for instruction level analysis. *Simulators* model and show a behaviour similar to the environment expected by the malware, while *emulators* replicate the behaviour of a system with higher accuracy but require more resources. Sandboxes are virtualised operating systems providing an isolated and reliable environment where to detonate malware. Refer to Ye *et al.* [1] for a more detailed description of these techniques. Execution traces are commonly used to extract features when dynamic analysis is employed. Reviewed articles generate execution traces by

using either sandboxes [42, 56, 68, 15, 16, 60, 57, 58, 69, 51, 52, 33] or emulators [70, 40]. Also program analysis tools and techniques can be useful in the feature extraction process by providing, for example, disassembly code and control- and data-flow graphs. An accurate disassembly code is important for obtaining correct *Byte sequences* and *Opcodes* features (§ 3.2.2), while control- and data-flow graphs can be employed in the extraction of *API and System Calls* (§ 3.2.2). For an extensive dissertation on dynamic analyses, refer to [71].

Among reviewed works, the majority relies on dynamic analyses [42, 55, 56, 15, 44, 16, 60, 57, 66, 46, 50, 58, 24, 26, 28, 30, 51, 52, 53, 35, 40], while the others use, in equal proportions, either static analyses alone [11, 12, 63, 64, 72, 17, 65, 19, 47, 49, 61, 22, 23, 25, 31, 73, 27, 29, 37, 38, 54, 67, 74, 39] or a combination of static and dynamic techniques [75, 18, 21, 48, 20, 59, 69, 62, 41]. Depending on the specific features, extraction processes can be performed by applying either static, dynamic, or hybrid analysis.

3.2.2. Portable Executable Features

This section provides an overview on what features are used by reviewed papers to achieve the objectives outlined in section 3.1. In many cases, surveyed works only refer to macro-classes without mentioning the specific features they employed. As an example, when n-grams are used, only a minority of works mention the size of n .

Byte Sequences. A binary can be characterised by computing features on its byte-level content. Analysing the specific sequences of bytes in a PE is a widely employed static technique. A few works use chunks of bytes of specific sizes [11, 74, 36], while many others rely on n-grams [12, 16, 75, 57, 18, 46, 26, 31, 27, 29, 52, 67, 74, 39, 35].

An n-gram is a sequence of n bytes, and features correspond to the different combination of these n bytes, namely each feature represents how many times a specific combination of n bytes occurs in the binary. The majority of works that specified the size of used n-grams relies on sequences no longer than 3 (i.e.

trigrams) [74, 52, 31, 16, 18, 67, 46, 48]. Indeed, the number of features to consider grows exponentially with n .

Opcodes. Opcodes identify the machine-level operations executed by a PE, and can be extracted through static analyses by examining the assembly code [63, 64, 45, 16, 18, 47, 49, 61, 20, 31, 37, 38, 54, 67, 74]. Opcode frequency is one of the most commonly used feature. It measures the number of times each specific opcode appears within the assembly or is executed by a PE [45, 38]. Others [18, 38] count opcode occurrences by aggregating them by operation type, e.g., mathematical instructions, memory access instructions. Similarly to n-grams, also sequences of opcodes are used as features [45, 37, 38, 74].

API and System Calls. Similarly to opcodes, APIs and system calls enable the analysis of samples' behaviour, but at a higher level. They can be either extracted statically or dynamically by analysing the disassembly code (to get the list of all calls that can be potentially executed) or the execution traces (for the list of calls actually invoked). While APIs allow to characterise what actions are executed by a sample [13, 48, 49, 23, 59, 31, 51, 40], looking at system call invocations provides a view on the interaction of the PE with the operating system [42, 56, 44, 57, 18, 46, 58, 20, 59, 26, 70, 28, 33]. Data extracted by observing APIs and system calls can be really large, and many works carry out additional processing to reduce feature space by using convenient data structures. One of the most popular data structures to represent PE behaviour and extract program structure is the control flow graph. This data structure allows compilers to produce an optimized version of the program itself and model control flow relationships [76]. Several works employ control flow graphs and their extensions for sample analysis, in combination with other feature classes [18, 21, 69, 62, 35].

Network Activity. A huge number of key information can be obtained by observing how the PE interacts with the network. Contacted addresses and generated traffic can unveil valuable aspects, e.g. regarding the communication with a command and control centre. Relevant features include statistics

on used protocols, TCP/UDP ports, HTTP requests, DNS-level interactions. Many surveyed works require dynamic analysis to extract this kind of information [42, 55, 56, 60, 50, 69, 32, 53, 40]. Other papers extract network-related inputs by monitoring the network and analysing incoming and outgoing traffic [66, 24, 41]. A complementary approach consists in analysing download patterns of network users in a monitored network [22]. It does not require sample execution and focuses on network features related to the download of a sample, such as the website from which the file has been downloaded.

File System. What file operations are executed by samples is fundamental to grasp evidence about the interaction with the environment and possibly detect attempts to gain persistence. Features of interest mainly concern how many files are read or modified, what types of files and in what directories, and which files appear in not-infected/infected machines [42, 55, 14, 49, 69, 52, 33, 53]. Sandboxes and memory analysis toolkits include modules for monitoring interactions with the file system, usually modelled by counting the number of files created/deleted/modified by the PE. In [53], the size of these files is considered as well, while Lin *et al.* leverage the number of created hidden files [52].

A particularly relevant type of file system features are those extracted from the *Windows Registry*. The registry is one of the main sources of information for a PE about the environment, and also represents a fundamental tool to hook into the operating system, for example to gain persistence. Discovering what keys are queried, created, deleted and modified can shed light on many significant characteristics of a sample [42, 52, 33, 53]. Usually, works relying on file system inputs monitor also the Windows Registry.

CPU Registers. The way CPU registers are used can also be a valuable indication, including whether any hidden register is used, and what values are stored in the registers, especially in the FLAGS register [49, 31, 59, 30].

PE file characteristics. A static analysis of a PE can provide a large set of valuable information such as sections, imports, symbols, used compilers [42, 19,

77, 23, 70, 34].

Strings. A PE can be statically inspected to explicitly look for the strings it contains, such as code fragments, author signatures, file names, system resource information [11, 48, 34, 31].

3.3. Malware Analysis Algorithms

This subsection reports what machine learning algorithms are used in surveyed works by organising them on the basis of whether the learning is supervised (§ 3.3.1), unsupervised (§ 3.3.2) or semi-supervised (§ 3.3.3).

3.3.1. Supervised Learning

Supervised learning is the task of gaining knowledge by providing statistical models with correct instance examples, during a preliminary phase called training. The supervised algorithms used by reviewed papers are *rule-based classifier* [11, 29, 30, 67, 40, 78, 60, 13], *Bayes classifier* [61, 20, 26, 51, 35], *Naïve Bayes* [11, 12, 15, 26, 51, 67, 35], *Bayesian Network* [21, 61, 20], *Support Vector Machine (SVM)* [12, 13, 15, 16, 65, 66, 48, 49, 61, 20, 24, 26, 31, 29, 51, 52, 53, 67, 35], *Multiple Kernel Learning* [18], *Prototype-based Classification* [57], *Decision Tree* [12, 13, 15, 48, 50, 61, 20, 23, 26, 51, 38, 53, 74], *Random Forest* [72, 66, 48, 26, 31, 51, 38, 32, 33, 35], *Gradient Boosting Decision Tree* [65, 67], *Logistic Model Tree* [69, 46, 67, 58], *k-Nearest Neighbors (k-NN)* [42, 49, 53, 36, 13, 15, 48, 51, 61], *Artificial Neural Network* [46, 34], *Multilayer Perceptron Neural Network* [15].

3.3.2. Unsupervised Learning

Unsupervised approaches do not rely on any training phase and learn directly from unlabeled data. Reviewed papers use these unsupervised learning algorithms: *Clustering with locality sensitive hashing* [56, 25, 39], *Clustering with Distance and Similarity Metrics* (using either Euclidean [57, 53] or Hamming distances [53], or cosine [53] or Jaccard similarities [53, 62]), *Expectation Maximization* [54], *k-Means Clustering* [43, 54], *k-Medoids* [45], *Density-based Spa-*

tial Clustering of Applications with Noise [41], *Hierarchical Clustering* [75, 53], *Prototype-based Clustering* [57], *Self-Organizing Maps* [65].

3.3.3. Semi-supervised Learning

Semi-supervised learning combines both labeled and unlabeled data for feeding statistical models to acquire knowledge. *Learning with Local and Global Consistency* is used in [17] while *Belief Propagation* in [14, 25, 27].

4. Characterization of Surveyed Papers

In this section we characterize each reviewed paper on the basis of analysis objective, used machine learning algorithm and features. Several details are also reported on the dataset used for the evaluation, including whether it is publicly available (*Public* column), where samples have been collected from (*Source* column) and whether the specific set of samples considered for the experiment is available (*Available* column). Indeed, many works declare they do not use all the executables in the dataset but they do not specify what samples they choose, which prevents to reproduce their results. The *Label* column states how samples have been labelled. Finally, *Benign*, *Malicious* and *Total* columns report benign executables count, malware count and their sum, respectively.

Malware detection. Table 1 lists all the reviewed works having malware detection as objective. Most used features are byte sequences and API/system call invocations, derived by executing the samples. Most of the works use more than one algorithm to find out the one guaranteeing more accurate results.

Malware similarity analysis. A table is provided for each version of this objective (§ 3.1.2). Tables 2 and 3 describe the works dealing with variants detection and families detection, respectively. For both, APIs and system calls are largely used, as well as malware interactions with the environment, i.e. memory, file system, and CPU registers. Tables 4 and 5 report the papers on similarities and differences detection, respectively. All the analysed papers but [61] rely on APIs and system calls collection. Works on differences detection,

in general, do not take into account the interactions with the hosting system, while those on similarities detection do.

Malware category detection. These articles focus on the identification of specific threats and, thus, on particular features such as byte sequences, opcodes, function lengths and network activity. Table 6 reports the works whose objective is the detection of malware category.

By reasoning on what algorithms and features have been used and what have not for specific objectives, the provided characterisation allows to easily identify gaps in the literature and, thus, possible research directions to investigate. For instance, all works on *differences detection* (see Table 5) but [61], rely on dynamically extracted APIs and system calls for building their machine learning models. Novel approaches can be explored by taking into account other features that capture malware interactions with the environment (e.g., memory, file system, CPU registers and Windows Registry).

Table 1: Characterization of surveyed papers having malware detection as objective.

Paper	Algorithms	Features	Limitations	Dataset samples					Total	
				Public	Source	Available	Labeling	Benign		Malicious
Schultz <i>et al.</i> [11]	Rule-based classifier, Naive Bayes	Strings and byte sequences	Proposed solutions are not effective against encrypted executables and the dataset used in the evaluations is small.	–	FTP sites	✓	Automated	1,001	3,265	4,266
Kolter and Maloof [12]	Decision Tree, Naive Bayes, SVM	Byte sequences	Payload classification fails in presence of binary obfuscation and the dataset used in the evaluations is very small.	✗	Internet, VX Heavens, and MITRE	✗	Automated	1,971	1,651	3,622
Ahmed <i>et al.</i> [13]	Decision Tree, Naive Bayes, SVM	APIs/System calls	The dataset used in the experimental evaluations is very small.	✓	Legitimate apps and VX Heavens	✗	–	100	416	516
Chau <i>et al.</i> [14]	Belief propagation	File system	Rare and new files cannot be accurately classified as benign or malicious.	✓	Symantec's Norton Community Watch	✗	–	?	?	$903 \cdot 10^6$
Firdausi <i>et al.</i> [15]	Decision Tree, Naive Bayes, SVM, k -NN, Multilayer Perceptron Neural Network	APIs/System calls, file system, and Windows Registry	The dataset used in the experimental evaluations is very small.	✓	Windows XP SP2	✗	–	250	220	470
Anderson <i>et al.</i> [16]	SVM	Byte sequences and APIs/system calls	The dataset used in the evaluations is small.	–	–	✗	–	615	1,615	2,230
Santos <i>et al.</i> [17]	Learning with Local and Global Consistency	Byte sequences	Proposed approach is not effective against packed malware and requires manual labeling of a portion of the small dataset. In particular, the dataset used in the experimental evaluations is small.	✓	Own machines and VX Heavens	✗	–	1,000	1,000	2,000
Anderson <i>et al.</i> [18]	Multiple Kernel Learning	Byte sequences, opcodes, and APIs/system calls	Instruction categorization is not optimal.	✗	Offensive Computing	✗	–	776	21,716	22,492
Younts [19]	Rule-based classifier	PE file characteristics	Only a subset of all the potential low level attributes is considered.	✗	SANS Institute	✗	–	65,000	$25 \cdot 10^5$	$25.65 \cdot 10^5$

Continue on the next page

Table 1: Characterization of surveyed papers having malware detection as objective. (Continued)

Paper	Algorithms	Features	Limitations	Dataset samples						
				Public	Source	Available	Labeling	Benign	Malicious	Total
Eskandari <i>et al.</i> [21]	Bayesian Network	APIs/System calls	Ignore specific instructions. Evasion/obfuscation techniques and samples requiring user interactions reduce the effectiveness of the proposed approach. The dataset is small.	✗	Research Laboratory at Shiraz University	✗	-	1,000	2,000	3,000
Santos <i>et al.</i> [20]	Bayesian Network, Decision Tree, k -NN, SVM	Opcodes, APIs/system calls, and raised exceptions	Packed malware, evasion techniques, and samples requiring user interactions reduce the accuracy of the proposed solution. The dataset is small.	✓	Own machines and VX Heavens	✗	-	1,000	1,000	2,000
Vádreu <i>et al.</i> [22]	Random Forest	PE file characteristics and network	Requires a huge number of samples labeled either as malicious or benign.	✗	Georgia Institute of Technology	✗	Automated	170,780	15,182	185,962
Bai <i>et al.</i> [23]	Decision Tree, Random Forest	PE file characteristics	Assume that samples are not packed and malware authors can properly modify PE header to remain undetected.	✓	Windows and Program Files folders and VXHeavens	✗	Automated	8,592	10,521	19,113
Kruczkowski and Szynkiewicz [24]	SVM	Network	The dataset used in the experimental evaluations is small.	✗	N6 Platform	✗	-	?	?	1,015
Tamersoy <i>et al.</i> [25]	Clustering with locality sensitive hashing	File system	Rare and new files cannot be accurately classified as benign or malicious.	✓	Symantec's Norton Community Watch	✗	-	1,663,506	47,956	4,970,865
Uppal <i>et al.</i> [26]	Decision Tree, Random Forest, Naive Bayes, SVM	Byte sequences and APIs/system calls	The dataset used in the experimental evaluations is very small.	✓	Legitimate apps and VX Heavens	✗	-	150	120	270
Chen <i>et al.</i> [27]	Belief propagation	File system	Rare and new files cannot be accurately classified as benign or malicious.	✗	Comodo Cloud Security Center	✗	-	19,142	2,883	69,165
Elhadi <i>et al.</i> [28]	Malicious graph matching	APIs/System calls	The dataset used in the experimental evaluations is extremely small.	✓	VX Heavens	✓	-	10	75	85

Table 1: Characterization of surveyed papers having malware detection as objective. (Continued)

Paper	Algorithms	Features	Limitations	Dataset samples						
				Public	Source	Available	Labeling	Benign	Malicious	Total
Feng <i>et al.</i> [29]	Rule-based classifier, SVM	Byte sequences	Only specific malware classes are considered for the approach evaluation.	✗	Windows system files and own AV platform	✗	–	100,000	135,064	235,064
Ghiasi <i>et al.</i> [30]	Rule-based classifier	APIs/System calls and CPU registers	APIs/System calls categorization could be not optimal and the dataset size is small.	✓	Windows XP system and Program Files folders and private repository	✗	–	390	850	1,240
Kwon <i>et al.</i> [32]	Random Forest	Network	Not able to detect bots with toolkit capabilities.	✓	Symantec's Worldwide Intelligence Network Environment	✗	–	?	?	24 * 10 ⁶
Mao <i>et al.</i> [33]	Random Forest	APIs/System calls, file system, and Windows Registry	Evasion techniques and samples requiring user interactions reduce the accuracy of the proposed approach. The dataset is small.	✓	Windows XP SP3 and VX Heavens	✗	–	534	7,257	7,791
Saxe and Berlin [34]	Neural Networks	Strings and PE file characteristics	Label assigned to training set may be inaccurate and the accuracy of the proposed approach decreases substantially when samples are obfuscated.	✗	Legitimate apps and own malware database	✗	Automated	81,910	350,016	431,926
Skraew <i>et al.</i> [74]	Decision Tree	Byte sequences and op-codes	Obfuscation techniques reduce detection accuracy.	✗	Legitimate files and apps and CWSandbox	✗	–	600	3,851	69,165
Wüchner <i>et al.</i> [35]	Naive Bayes, Random Forest, SVM	Byte sequences, APIs/system calls, file system, and Windows Registry	Obfuscation techniques applied by the authors may not reflect the ones of real-world samples. The dataset is small.	✗	Legitimate app downloads and Malicia	✗	–	513	6,994	7,507
Raff and Nicholas [36]	k-NN with Lempel-Ziv Jaccard distance	Byte sequences	Obfuscation techniques reduce detection accuracy.	✗	Industry partner	✗	–	240,000	237,349	477,349

Table 2: Characterization of surveyed papers having malware variants selection as objective. ¹Instead of using machine learning techniques, Gharacheh *et al.* rely on Hidden Markov Models to detect variants of the same malicious sample [37].

Paper	Algorithms	Features	Limitations	Dataset samples						
				Public	Source	Available	Labeling	Benign	Malicious	Total
Gharacheh <i>et al.</i> [37]	- ¹	Opcodes	Opcod sequence is not optimal and the dataset size is very small.	✓	Cywin and VX Heavens	✗	-	?	?	740
Ghiasi <i>et al.</i> [30]	Rule-based classifier	APIs/System calls and CPU registers	APIs/System calls categorization could be not optimal and the dataset size is small.	✓	Windows XP system and Program Files folders and private repository	✗	-	390	850	1,240
Khodamoradi <i>et al.</i> [38]	Decision Tree, Random Forest	Opcodes	Opcod sequence is not optimal and the dataset size is very small.	✓	Windows XP system and Program Files folders and self-generated metamorphic malware	✗	-	550	280	830
Upchurch and Zhou [39]	Clustering with locality sensitive hashing	Byte sequences	The dataset size is extremely small.	✗	Sampled from security incidents	✓	Manual	0	85	85
Liang <i>et al.</i> [40]	Rule-based classifier	APIs/System calls, file system, Windows Registry, and network	Monitored API/system call set could be not optimal and the dataset size is small.	✗	Anubis website	✗	-	0	330,248	330,248

Continue on the next page

Table 2: Characterization of surveyed papers having malware detection as objective. (Continued)

Paper	Algorithms	Features	Limitations	Dataset samples						
				Public	Source	Available	Labeling	Benign	Malicious	Total
Vadrevu and Perdisci [41]	DBSCAN clustering	APIs/System calls, PE file characteristics, and network	Evasion techniques and samples requiring user interactions reduce the accuracy of the proposed approach.	✗	Security company and large Research Institute	✗	–	0	1, 651, 906	1, 651, 906

Table 3: Characterization of surveyed papers having malware families selection as objective. ²Asquith describes aggregation overlay graphs for storing PE metadata, without further discussing any machine learning technique that could be applied on top of these new data structures.

Paper	Algorithms	Features	Limitations	Dataset samples						
				Public	Source	Available	Labeling	Benign	Malicious	Total
Huang <i>et al.</i> [43]	k -Means-like algorithm	Byte sequences	Instruction sequence categorization could be not optimal and the dataset size is small.	✗	Kingsoft Corporation	✗	–	0	2, 029	2, 029
Park <i>et al.</i> [44]	Malicious graph matching	APIs/System calls	Approach vulnerable to APIs/system calls injection and the dataset used in the experimental evaluations is very small.	✗	Legitimate apps and Anubis Sandbox	✗	Automated	80	300	380
Ye <i>et al.</i> [45]	k -Medoids variants	Opcodes	Instruction categorization could be not optimal.	✗	Kingsoft Corporation	✗	–	0	11, 713	11, 713
Dahl <i>et al.</i> [46]	Logistic Regression, Neural Networks	Byte sequences and APIs/system calls	The authors obtain a high two-class error rate.	✗	Microsoft	✗	Mostly manual	817, 485	1, 843, 359	3, 760, 844
Hu <i>et al.</i> [47]	Prototype-based clustering	Opcodes	Obfuscation techniques reduce the effectiveness of their prototype for malware family selection.	✗	–	✗	Manual and automated	0	137, 055	137, 055

Continue on the next page

Table 3: Characterization of surveyed papers having malware families selection as objective. (Continued)

Paper	Algorithms	Features	Limitations	Dataset samples						
				Public	Source	Available	Labeling	Benign	Malicious	Total
Islam <i>et al.</i> [48]	Decision Tree, k -NN, Random Forest, SVM	Strings, byte sequences and APIs/system calls	The proposed approach is less effective novel samples. The dataset is small.	✗	CA Labs	✗	—	51	2,398	2,939
Kong and Yan [49]	SVM, k -NN	Opcodes, memory, file system, and CPU registers	Significant differences in samples belonging to the same family reduce the proposed approach accuracy.	✗	Offensive Computing	✗	Automated	0	526,179	526,179
Nari and Ghorbani [50]	Decision Tree	Network	Network features are extracted by a commercial traffic analyzer. The dataset used in the experimental evaluations is small.	✗	Communication Research Center Canada	✗	Automated	0	3,768	3,768
Ahmadi <i>et al.</i> [31]	SVM, Random Forest, Gradient Boosting Decision Tree	Byte sequences, opcodes, APIs/system calls, Windows Registry, CPU registers, and PE file characteristics	Selected features can be further reduced to have a clearer view of the reasons behind sample classification.	✓	Microsoft's malware classification challenge	✗	—	0	21,741	21,741
Asquith [70]	_2	APIs/System calls, memory, file system, PE file characteristics, and raised exceptions	_5	—	—	—	—	—	—	—
Lin <i>et al.</i> [52]	SVM	Byte sequences, APIs/system calls, file system, and CPU registers	Selected API/system call set could be not optimal. Evaluation techniques and samples requiring user interactions reduce the accuracy of the proposed approach. The dataset is small.	✗	Own sandbox	✗	—	389	3,899	4,288

Continue on the next page

Table 3: Characterization of surveyed papers having malware families selection as objective. (Continued)

Paper	Algorithms	Features	Limitations	Dataset samples						
				Public	Source	Available	Labeling	Benign	Malicious	Total
Kawaguchi and Omote [51]	Decision Tree, Random Forest, k -NN, Naïve Bayes	APIs/System calls	This classification approach can be easily evaded by real-world malware. The dataset used in the experimental evaluations is very small.	✗	FFRI Inc.	✗	–	236	408	644
Mohaisen <i>et al.</i> [53]	Decision Tree, k -NN, SVM, Clustering with different similarity measures, Hierarchical clustering	File system, Windows Registry, CPU registers, and network	Evasion techniques reduce the accuracy of the proposed solution.	✗	AMAL system	✗	Manual and automated	0	115,157	115,157
Pai <i>et al.</i> [54]	k -Means, Expectation Maximization	Opcodes	Obfuscation techniques reduce the effectiveness of the employed approach. The dataset is small.	✗	Cygwin utility files and Malicia	✗	–	213	8,052	8,265
Raff and Nicholas [36]	k -NN with Lempel-Ziv Jaccard distance	Byte sequences	Obfuscation techniques reduce detection accuracy.	✗	Industry partner	✗	–	240,000	237,349	477,349

Table 4: Characterization of surveyed papers having malware similarities detection as objective. ³SVM is used only for computing the optimal values of weight factors associated to each feature chosen to detect similarities among malicious samples.

Paper	Algorithms	Features	Limitations	Dataset samples						
				Public	Source	Available	Labeling	Benign	Malicious	Total
Bailey <i>et al.</i> [55]	Hierarchical clustering with normalized compression distance	APIs/System calls, file system, Windows Registry, and network	Evasion techniques and samples requiring user interactions reduce the accuracy of the proposed classification method. The dataset is small.	✗	Albor Malware Library and public repository	✗	Automated	0	8, 228	8, 228
Bayer <i>et al.</i> [56]	Clustering with locality sensitive hashing	APIs/System calls	Evasion techniques and samples requiring user interactions reduce the approach accuracy.	✗	Anubis website	✗	Manual and automated	0	75, 692	75, 692
Rieck <i>et al.</i> [57]	Prototype-based classification and clustering with Euclidean distance	Byte sequences and APIs/system calls	Evasion techniques and samples requiring user interactions reduce the accuracy of the proposed framework.	✗	CWSandbox and Sumbelt Software	✗	Automated	0	36, 831	36, 831

Continue on the next page

Table 4: Characterization of surveyed papers having malware similarities detection as objective. (Continued)

Paper	Algorithms	Features	Limitations	Dataset samples						
				Public	Source	Available	Labeling	Benign	Malicious	Total
Palahan <i>et al.</i> [58]	Logistic Regression	APIs/System calls	Evasion techniques and samples requiring user interactions reduce the accuracy of the proposed framework, while unknown observed behaviors are classified as malicious. The dataset used in the experimental evaluations is very small.	X	Own honeypot	X	-	49	912	961
Egele <i>et al.</i> [59]	SVM ³	APIs/System calls, memory, and CPU registers	The accuracy of computed PE function similarities drops when different compiler toolchains or aggressive optimization levels are used. The dataset is small.	✓	coreutils-8.13 program suite	X	-	1,140	0	1,140

Table 5: Characterization of surveyed papers having malware differences detection as objective.

Paper	Algorithms	Features	Limitations	Dataset samples						
				Public	Source	Available	Labeling	Benign	Malicious	Total
Bayer <i>et al.</i> [56]	Clustering with locality sensitive hashing	APIs/System calls	Evasion techniques and samples requiring user interactions reduce the approach accuracy.	x	Anubis website	x	Manual and automated	0	75,692	75,692
Lindorfer <i>et al.</i> [60]	Rule-based classifier	APIs/System calls and network	Sophisticated evasion techniques and samples requiring user interactions can still bypass detection processes. The dataset used in the experimental evaluations is small.	x	Anubis Sandbox	x	Automated	0	1,871	1,871
Rieck <i>et al.</i> [57]	Prototype-based classification and clustering with Euclidean distance	Byte sequences and APIs/system calls	Evasion techniques and samples requiring user interactions reduce the accuracy of the proposed framework.	x	CWSandbox and Sunbelt Software	x	Automated	0	36,831	36,831

Continue on the next page

Table 5: Characterization of surveyed papers having malware differences detection as objective. (Continued)

Paper	Algorithms	Features	Limitations	Dataset samples						
				Public	Source	Available	Labeling	Benign	Malicious	Total
Palahan <i>et al.</i> [58]	Logistic Regression	APIs/System calls	Evasion techniques and samples requiring user interactions reduce the accuracy of the proposed framework, while unknown observed behaviors are classified as malicious. The dataset used in the experimental evaluations is small.	✗	Own honeypot	✗	–	49	912	961
Santos <i>et al.</i> [61]	Decision Tree, k -NN, Bayesian Network, Random Forest	Opcodes	Opcode sequence is not optimal and the dataset size is small. The proposed method is not effective against packed malware. The dataset is small.	✓	Own machines and VXHeavens	✗	Automated	1,000	1,000	2,000

Continue on the next page

Table 5: Characterization of surveyed papers having malware differences detection as objective. (Continued)

Paper	Algorithms	Features	Limitations	Dataset samples							
				Public	Source	Available	Labeling	Benign	Malicious	Total	
Polino <i>et al.</i> [62]	Clustering with Jac-card similarity	Jac-APIs/System calls	Evasion techniques, packed malware, and samples requiring user interactions reduce the accuracy of the proposed framework. API calls sequence used to identify sample behaviors is not optimal. The dataset size is small.	-	-	-	-	-	?	?	2, 136

Table 6: Characterization of surveyed papers having malware category detection as objective. ⁴Instead of using machine learning techniques, these articles rely on Hidden Markov Models to detect metamorphic viruses [63, 64].

Paper	Algorithms	Features	Limitations	Dataset samples						
				Public	Source	Available	Labeling	Benign	Malicious	Total
Wong and Stamp [63]	⁴	Opcodes	Detection fails if metamorphic malware are similar to benign files. The dataset is extremely small.	X	Cyghwin and VX Heavens generators	X	-	40	25	65

Table 6: Characterization of surveyed papers having malware category detection as objective. (Continued)

Paper	Algorithms	Features	Limitations	Dataset samples						
				Public	Source	Available	Labeling	Benign	Malicious	Total
Attaluri <i>et al.</i> [64]	⁶	Opcodes	The proposed approach is not effective against all types of metamorphic viruses. The dataset size is very small.	✗	Cygwin, legitimate DLLs and VX Heavens generators	✗	-	240	70	310
Tian <i>et al.</i> [78]	Rule-based classifier	Function length	Function lengths alone are not sufficient to detect Trojans and the dataset used in the experimental evaluations is very small.	-	-	-	-	0	721	721
Siddiqui <i>et al.</i> [72]	Decision Tree, Random Forest	Opcodes	Advanced packing techniques could reduce detection accuracy. The dataset used in the experimental evaluations is small.	✓	Windows XP and VX Heavens	✗	-	1,444	1,330	2,774
Chen <i>et al.</i> [65]	Random Forest, SVM	Byte sequences	The proposed framework heavily relies on security companies' encyclopedias.	✗	Trend Micro	✗	-	0	330,248	330,248

Continue on the next page

Table 6: Characterization of surveyed papers having malware category detection as objective. (Continued)

Paper	Algorithms	Features	Limitations	Dataset samples					Total	
				Public	Source	Available	Labeling	Benign		Malicious
Comar <i>et al.</i> [66]	Random Forest, SVM	Network	Network features are extracted by a commercial traffic analyzer.	✗	Internet Service Provider	✗	Manual and automated	212,505	4,394	216,899
Kwon <i>et al.</i> [32]	Random Forest	Network	Not able to detect bots with rootkit capabilities.	✓	Symantec's Worldwide Intelligence Network Environment	✗	-	?	?	$24 * 10^6$
Sexton <i>et al.</i> [67]	Rule-based classifier, Logistic Regression, Naïve Bayes, SVM	Byte sequences and opcodes	Obfuscation techniques reduce detection accuracy. The dataset used in the experimental evaluations is small.	-	-	-	-	4,622	197	4,819

5. Issues and Challenges

Based on the characterization detailed in section 4, this section identifies the main issues and challenges of surveyed papers. In the specific, the main problems regard the usage of anti-analysis techniques by malware (§ 5.1), what operation set to consider 5.2 and used dataset 5.3.

5.1. Anti-analysis Techniques

Malware developers want to avoid their samples to be analysed, so they devise and refine several *anti-analysis techniques* that are effective in hindering the reverse engineering of executables. Indeed many surveyed works claim that the solution they propose does not work or loses in accuracy when samples using such techniques are considered (§ 4).

Static analysis (§ 3.2.1) is commonly prevented by rendering sample binary and resources unreadable through *obfuscation*, *packing* or *encryption*. Anyway, at runtime, code and any other concealed data has to be either deobfuscated, unpacked or decrypted to enable the correct execution of the payload. This implies that such a kind of anti-analysis techniques can be overcome by using dynamic analysis (§ 3.2.1) to make the sample unveil hidden information and load them in memory, where they can then be extracted by creating a dump. Refer to Ye *et al.* [1] for a detailed disquisition on how obfuscation, packing and encryption are used by malware developers.

More advanced anti-analysis techniques exist to keep malware internals secret even when dynamic analysis is used. One approach, commonly referred to as *environmental awareness*, consists in the malware trying to detect whether it is being executed in a controlled setting where an analyst is trying to dissect it, for example by using a virtual machine or by running the sample in debug mode. If any cue is found of possibly being under analysis, then the malware does not execute its malicious payload. Miramirkhani *et al.* [79] show that a malware can easily understand if it is running into an artificial environment. Other approaches rely on *timing-based evasion*, i.e. they only show their malicious behaviour at predetermined dates and times. Other malware instead

require or wait for some *user interaction* to start their intended activity, in order to make any kind of automatic analysis infeasible.

Identifying and overcoming these anti-analysis techniques is an important direction to investigate to improve the effectiveness of malware analysis. Recent academic and not-academic literature are aligned on this aspect. Karpin and Dorfman [80] highlight the need to address very current problems such as discovering where malware configuration files are stored and whether standard or custom obfuscation/packing/encryption algorithms are employed. De-obfuscation [81, 82] and other operations aimed at supporting binary reverse engineering, such as function similarity identification [83], are still very active research directions. Symbolic execution techniques [84] are promising means to understand what execution paths trigger the launch of the malicious payload.

5.2. Operation Set

Opcodes, instructions, APIs and system calls (hereinafter, we refer to them in general as *operations*) are the most used and powerful features employed for malware analysis (§ 4), as they allow to directly and accurately model sample behaviour. Normally, to reduce complexity and required computational power, only a subset of all the available operations is considered. This choice of ignoring part of the operations at disposal can reduce the accuracy of malware behaviour model, which in turn reflects on the reliability of analysis outcomes. This issue has been raised explicitly in some surveyed papers, including [18, 37, 38], while others are anyway affected although they do not mention it, such as [30, 40, 43].

On one hand, this challenge can be addressed by either improving or using different machine learning techniques to achieve a more effective feature selection. On the other hand, program analysis advances can be leveraged to enhance the accuracy of disassemblers and decompilers, indeed these tools are known to be error-prone [85, 86] and are thus likely to affect negatively the whole analyses. Approaches that improve the quality of generated disassembly and decompiled code, as in [87], can reduce the impact due to these errors.

5.3. Datasets

More than 72% of surveyed works use datasets with both malicious and benign samples, while about 28% rely on datasets with malware only. Just two works rely on benign datasets only [59, 73], because their objectives are identifying sample similarities and attributing the ownership of some source codes under analysis, respectively.

Figure 2 shows the dataset sources for malicious and benign samples. It is worth noting that most of benign datasets consists of legitimate applications (e.g. software contained in “Program Files” or “system” folders), while most of malware have been obtained from public repositories, security vendors and popular sandboxed analysis services. The most popular public repository in the examined works is VX Heavens [88], followed by Offensive Computing [89] and Malicia Project [90]. The first two repositories are still actively maintained at the time of writing, while Malicia Project has been permanently shut down due to dataset ageing and lack of maintainers.

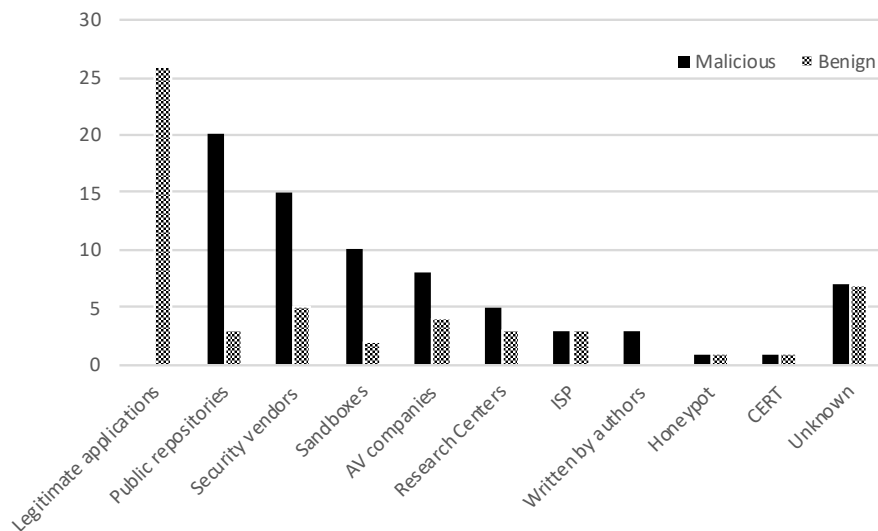


Figure 2: Frequency histogram showing how many reviewed papers use each type of source (e.g. public repositories, honeypot) to collect their datasets, and whether it is used to gather malware or benign samples.

Security vendors, popular sandboxed analysis services, and AV companies have access to a huge number of samples. Surveyed works rely on CWSandbox, developed by ThreatTrack Security [91], and Anubis [92]. As can be observed from Figure 2, these sandboxes are mainly used for obtaining malicious samples. Internet Service Providers (ISPs), honeypots and Computer Emergency Response Teams (CERTs) share with researchers both benign and malicious datasets. A few works use malware developed by the authors [37, 38], created using malware toolkits [63] such as Next Generation Virus Construction Kit, Virus Creation Lab, Mass Code Generator and Second Generation Virus Generator, all available on VX Heavens [88]. A minority of analysed papers do not mention the source of their datasets.

Among surveyed papers, a recurring issue is the size of used dataset. Many works, including [12, 13, 15], carry out evaluations on less than 1,000 samples. Just 39% of reviewed studies test their approaches on a population greater than 10,000 samples.

When both malicious and benign samples are used for the evaluation, it is crucial to reflect their real distribution [11, 12, 13, 72, 15, 44, 16, 17, 18, 93, 19, 46, 21, 48, 77, 58, 61, 20, 23, 26, 28, 29, 30, 51, 52, 33, 54, 34, 74, 35, 36]. Indeed, there needs to be a huge imbalance because non-malware executables are the overwhelming majority. 48% of surveyed works do not take care of this aspect and use datasets that either are balanced between malware and non-malicious software or, even, have more of the former than the latter. In [19], Yonts supports his choice of using a smaller benign dataset by pointing out that changes in standard system files and legitimate applications are little. 38% of examined papers employ instead datasets having a proper distribution of malware and non-malware, indeed they are either unbalanced towards benign samples or use exclusively benign or malicious software. As an example, the majority of surveyed papers having malware similarities detection as objective (see Table 4) contains exclusively either malware or legitimate applications [55, 56, 57, 59]. The remaining 14% does not describe how datasets are composed.

Differently from other research fields, no reference benchmark is available for

malware analysis to compare accuracy and performance with other works. Furthermore, published results are known to be biased towards good results [94]. In addition, since the datasets used for evaluations are rarely shared, it is nearly impossible to compare works. Only two surveyed works have shared their dataset [11, 39], while a third one plans to share it in the future [53]. It is worth mentioning that one of the shared dataset is from 2001, hence almost useless today. Indeed, temporal information is crucial to evaluate malware analysis results [95] and determine whether machine learning models have become obsolete [96, 97].

Given such lack of reference datasets, we propose three desiderata for malware analysis benchmarks.

1. Benchmarks should be labeled accordingly to the specific objectives to achieve. As an example, benchmarks for families selection should be labeled with samples' families.
2. Benchmarks should model realistically the sample distributions of real-world scenarios, considering the objectives to attain. For example, benchmarks for malware detection should contain a set of legitimate applications orders of magnitude greater than the number of malware samples.
3. Benchmarks should be actively maintained and updated over time with new samples, trying to keep pace with the malware industry. Samples should also be provided with temporal information, e.g., when they have been spotted first.

Datasets used in [11] and [39] are correctly labeled according to malware detection and malware variants selection objectives, respectively. Both datasets are not balanced. In Shultz *et al.* [11], the described dataset is biased towards malicious programs, while in [39] diverse groups of variants contain a different number of samples, ranging from 3 to 20. Finally, analysed datasets are not actively maintained and do not contain any temporal information (in [11], the authors do not mention if such information has been included into the dataset).

6. Topical Trends

This section outlines a list of topical trends in malware analysis, i.e. topics that are currently being investigated but have not reached the same level of maturity of the other areas described in previous sections.

6.1. *Malware Development Detection*

Malware developers can use online public services like VirusTotal [98] and Malwr [99] to test the effectiveness of their samples in evading most common antiviruses. Malware analysts can leverage such behaviour by querying these online services to obtain additional information useful for the analysis, such as submission time and how many online antiviruses classify a sample as malicious. Graziano *et al.* [69] leverage submissions to an online sandbox for identifying cases where new samples are being tested, with the final aim to detect novel malware during their development process. Surprisingly, it turned out that samples used in infamous targeted campaigns had been submitted to public sandboxes months or years before.

With reference to the proposed taxonomy, advances in the state of the art in malware analysis could be obtained by analysing submissions to online malware analysis services, to extract additional machine learning features and gather intelligence on what next malware are likely to be.

6.2. *Malware Attribution*

Another aspect of interest for malware analysts is the identification of who developed a given sample, i.e. the attribution of a malware to a specific malicious actor. There are a number of features in a binary to support this process: used programming language, included IP addresses and URLs, and the language of comments and resources. Additional, recently proposed features which can be used for attribution are the time slot when the malware communicates with a command and control centre and what digital certificates are used [100]. Features related to the coding style can also reveal important details on developer's identity, at least for arguing whether different malware have been developed

by the same person or group. In [73], the author’s coding style of a generic software (i.e. not necessarily malicious) is accurately profiled through syntactic, lexical, and layout features. Unfortunately, this approach requires the availability of source code, which happens only occasionally, e.g. in case of leaks and/or public disclosures.

Malware attribution can be seen as an additional analysis objective, according to the proposed taxonomy. Progresses in this direction through machine learning techniques are currently hindered by the lack of ground truth on malware authors, which proves to be really hard to provision. Recent approaches leverage on public reports referring to APT groups and detailing what malware they are supposed to have developed: those reports are parsed to mine the relationships between malicious samples and corresponding APT group authors [101]. The state of the art in malware attribution through machine learning can be advanced by researching alternative methods to generate reliable ground truth on malware developers, or on what malware have been developed by the same actor.

6.3. Malware Triage

Given the huge amount of new malware that need to be analysed, a fast and accurate prioritisation is required to identify what samples deserve more in depth analyses. This can be decided on the basis of the level of similarity with already known samples. If a new malware resembles very closely other binaries that have been analysed before, then its examination is not a priority. Otherwise, further analyses can be advised if a new malware really looks differently from everything else observed so far. This process is referred to as *malware triage* and shares some aspects with malware similarity analysis, as they both provide key information to support malware analysis prioritisation. Anyway, they are different because triage requires faster results at the cost of worse accuracy, hence different techniques are usually employed [75, 77, 101, 86].

Likewise attribution, triage can be considered as another malware analysis objective. One important challenge of malware triage is finding the proper

trade-off between accuracy and performance, which fits with the problems we address in the context of malware analysis economics (see Section 7).

6.4. Prediction of Future Variants

Compared to malware analysts, malware developers have the advantage of knowing current anti-malware measures and thus novel variants can be designed accordingly. A novel trend in malware analysis is investigating the feasibility to fill that gap by predicting how future malware will look like, so as to allow analysts to update anti-malware measures ahead. Howard *et al.* [102] use machine learning techniques to model patterns in malware family evolutions and predict future variants.

This problem can be seen as yet another objective in the malware analysis taxonomy. It has not been investigated much yet, only a couple of works seem to address that topic [103, 102]. Given its great potential to proactively identify novel malware, and considering the opportunity to exploit existing malware families datasets, we claim the worthiness to boost the research on malware evolution prediction through machine learning techniques.

6.5. Other Features

This section describes features different from those analysed in Section 3.2.2 and that have been used by just a few papers so far. In view of advancing the state of the art in malware analysis, additional research is required on the effectiveness of using such features to improve the accuracy of machine learning techniques.

6.5.1. Memory Accesses

Any data of interest, such as user generated content, is temporary stored in main memory, hence analysing how memory is accessed can reveal important information about the behaviour of an executable [104]. Kong *et al.* rely on statically trace reads and writes in main memory [49], while Egele *et al.* dynamically trace values read from and written to stack and heap [59].

6.5.2. *Function Length*

Another characterising feature is the function length, measured as the number of bytes contained in a function. This input alone is not sufficient to discriminate malicious executables from benign software, indeed it is usually combined with other features. This idea, formulated in [78], is adopted in [48], where function length frequencies, extracted through static analysis, are used together with other static and dynamic features.

6.5.3. *Raised Exceptions*

The analysis of the exceptions raised during the execution can help understanding what strategies a malware adopts to evade analysis systems [20, 70]. A common trick to deceive analysts is throwing an exception to run a malicious handler, registered at the beginning of malware execution. In this way, examining the control flow becomes much more complex.

7. **Malware Analysis Economics**

Analysing samples through machine learning techniques requires complex computations for extracting desired features and running chosen algorithms. The time complexity of these computations has to be carefully taken into account to ensure they complete fast enough to keep pace with the speed new malware are developed. Space complexity has to be considered as well, indeed feature space can easily become excessively large (e.g., using n-grams), and also the memory required by machine learning algorithms can grow to the point of saturating available resources.

Time and space complexities can be either reduced to adapt to processing and storage capacity at disposal, or they can be accommodated by supplying more resources. In the former case, the analysis accuracy is likely to worsen, while, in the latter, accuracy levels can be preserved at the cost of providing more computing machines, storage and network. There exist therefore trade-offs between maintaining high accuracy and performance of malware analysis

on one hand, and supplying the required equipment on the other. We refer to the study of these trade-offs as *malware analysis economics*, and in this section we provide some initial qualitative discussions on this novel topic.

The time needed to analyse a sample through machine learning is mainly spent in feature extraction and algorithm execution. While time complexity of machine learning algorithms is widely discussed in literature, the same does not apply for the study of feature extraction execution time. The main aspect to take into account is whether desired features come from static or dynamic analysis, which considerably affects execution time because the former does not require to run the samples, while the latter does.

To deepen even further this point, Table 7 reports for each feature type whether it can be extracted through static or dynamic analysis. It is interesting to note that certain types of features can be extracted both statically and dynamically, with significant differences on execution time as well as on malware analysis accuracy. Indeed, while certainly more time-consuming, dynamic analysis enables to gather features that contribute relevantly to the overall analysis effectiveness [105]. As an example, we can consider the features derived from API calls (see Table 7), which can be obtained both statically and dynamically. Tools like IDA provide the list of imports used by a sample and can statically trace what API calls are present in the sample code. Malware authors usually hide their suspicious API calls by inserting in the source code a huge number of legitimate APIs. By means of dynamic analysis, it is possible to obtain the list

Table 7: Type of analysis required for extracting the inputs presented in Sections 3.2.2 and 6.5: strings, byte sequences, opcodes, APIs/system calls, file system, CPU registers, PE file characteristics, network, AV/Sandbox submissions, code stylometry, memory accesses, function length, and raised exceptions.

Analysis	Str.	Byte seq.	Ops	APIs Sys. calls	File sys.	CPU reg.	PE file char.	Net.	Sub-mis.	Code stylo.	Mem.	Func. len.	Exc.
Static	✓	✓	✓	✓		✓	✓		✓	✓		✓	
Dynamic				✓	✓	✓		✓	✓		✓		✓

of the APIs that the sample has actually invoked, thus simplifying the identification of those suspicious APIs. By consequences, in this case dynamic analysis is likely to generate more valuable features compared to static analysis. Maze-Walker [106] is a typical example of how dynamic information can integrate static analysis.

Although choosing dynamic analysis over, or in addition to, static seems obvious, its inherently higher time complexity constitutes a potential performance bottleneck for the whole malware analysis process, which can undermine the possibility to keep pace with malware evolution speed. The natural solution is to provision more computational resources to parallelise analysis tasks and thus remove bottlenecks. In turn, such solution has a cost to be taken into account when designing a malware analysis environment, such as the one presented by Laurenza *et al.* [107].

The qualitative trade-offs we have identified are between accuracy and time complexity (i.e., higher accuracy requires larger times), between time complexity and analysis pace (i.e., larger times implies slower pace), between analysis pace and computational resources (faster analysis demands using more resources), and between computational resources and economic cost (obviously, additional equipment has a cost). Similar trade-offs also hold for space complexity. As an example, when using n -grams as features, it has been shown that larger values of n lead to more accurate analysis, at cost of having the feature space grow exponentially with n [26, 52]. As another example, using larger datasets in general enables more accurate machine learning models and thus better accuracy, provided the availability of enough space to store all the samples of the dataset and the related analysis reports.

Table 8: Relationship between n and number of features.

n	feature count
1	187
2	6740
3	46216
4	130671
5	342663

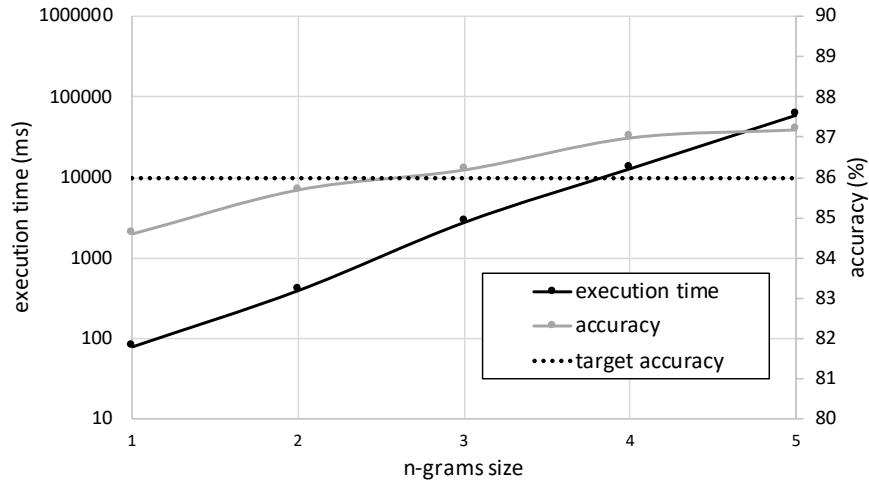


Figure 3: Relationship between execution time (in logarithmic scale) and detection accuracy as n varies. The target accuracy of 86% is also reported.

We present a qualitative, simplified example of analysis that leverages on the trade-offs just introduced. The scenario we target regards detecting malware families of new malicious samples (§ 3.1.2) using as features n -grams computed over invoked APIs (§ 3.2.2), recorded through dynamic analysis (§ 3.2.1). We want here to explore the trade-offs between family detection accuracy, execution time, analysis pace and cost, in terms of required computational resources. For what concerns the scenario and qualitative numbers on the relationships between n , the number of features, accuracy and execution time, we take inspiration from the experimental evaluation presented by Lin *et al.* [52]. Table 8 shows the relationship between n and feature count. We introduce a few simplifying assumptions and constraints to make this qualitative example as consistent as possible. We assume that the algorithm used to detect families is parallelisable and ideally scalable, meaning that by doubling available machines we also double the throughput, i.e. the number of malware analysed per second. We want to process one million malware per day with an accuracy of at least 86%.

Figure 3 highlights the trade-off between execution time (in logarithmic scale) and detection accuracy as n is varied. As n grows, the accuracy in-

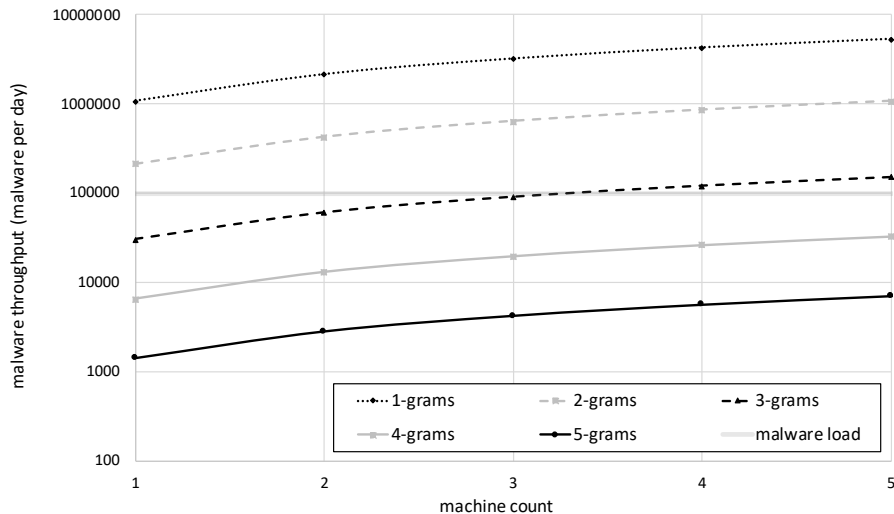


Figure 4: Relationship between machine count and malware throughput (in logarithmic scale) for different n-grams sizes. The one million malware per day to sustain is also reported.

creases almost linearly while the execution time has an exponential rise, which translates to an exponential decrease of how many malware per second can be processed. It can be noted that the minimum n-grams size to meet the accuracy requirement of 86% is 3. The trade-off between analysis pace and cost can be observed in Figure 4 where, by leveraging on the assumption of ideal scalability of the detection algorithm, it is shown that sustainable malware throughput (in logarithmic scale) increases linearly as the algorithm is parallelised on more machines. 4-grams and 5-grams cannot be used to cope with the expected malware load of one million per day, at least when considering up to five machines. On the other hand, by using four machines and 3-grams, we can sustain the target load and at the same time meet the constraint on detection accuracy.

The presented toy example is just meant to better explain how malware analysis economics can be used in practical scenarios. We claim the significance of investigating these trade-offs more in detail, with the aim of outlining proper guidelines and strategies to design a malware analysis environment in compliance with requirements on analysis accuracy and pace, while respecting budget

constraints.

8. Conclusion

We presented a survey on existing literature on malware analysis through machine learning techniques. There are five main contributions of our work. First, we proposed an organization of reviewed works according to three orthogonal dimensions: *the objective of the analysis*, *the type of features extracted from samples*, *the machine learning algorithms used to process these features*. Such characterization provides an overview on how machine learning algorithms can be employed in malware analysis, emphasising which specific feature classes allow to achieve the objective(s) of interest. Second, we have arranged existing literature on PE malware analysis through machine learning according the proposed taxonomy, providing a detailed comparative analysis of surveyed works. Third, we highlighted the current issues of machine learning for malware analysis: anti-analysis techniques used by malware, what operation set to consider for the features and used datasets. Fourth, we identified topical trends on interesting objectives and features, such as malware attribution and triage. Fifth, we introduced the novel concept of *malware analysis economics*, concerning the investigation and exploitation of existing trade-offs between performance metrics of malware analysis (e.g., analysis accuracy and execution time) and economical costs.

Noteworthy research directions to investigate can be linked to those contributions. Novel combinations of objectives, features and algorithms can be investigated to achieve better accuracy compared to the state of the art. Moreover, observing that some classes of algorithms have never been used for a certain objective may suggest novel directions to examine further. The discussion on malware analysis issues can provide further ideas worth to be explored. In particular, defining appropriate benchmarks for malware analysis is a priority of the whole research area. The novel concept of malware analysis economics can encourage further research directions, where appropriate tuning strategies

can be provided to balance competing metrics (e.g. accuracy and cost) when designing a malware analysis environment.

Acknowledgment

This work has been partially supported by a grant of the Italian Presidency of Ministry Council and by the Laboratorio Nazionale of Cyber Security of the CINI (Consorzio Interuniversitario Nazionale Informatica).

References

- [1] Y. Ye, T. Li, D. Adjeroh, S. S. Iyengar, A survey on malware detection using data mining techniques, *ACM Computing Surveys (CSUR)* 50 (3) (2017) 41.
- [2] AV-TEST, Security Report 2016/17, https://www.av-test.org/fileadmin/pdf/security_report/AV-TEST_Security_Report_2016-2017.pdf (2017).
- [3] A. Shabtai, R. Moskovitch, Y. Elovici, C. Glezer, Detection of malicious code by applying machine learning classifiers on static features: A state-of-the-art survey, *Inf. Secur. Tech. Rep.* 14 (1) (2009) 16–29.
- [4] M. K. Sahu, M. Ahirwar, A. Hemlata, A review of malware detection based on pattern matching technique, *Int. J. of Computer Science and Information Technologies (IJCSIT)* 5 (1) (2014) 944–947.
- [5] A. Souri, R. Hosseini, A state-of-the-art survey of malware detection approaches using data mining techniques, *Human-centric Computing and Information Sciences* 8 (1) (2018) 3.
- [6] C. LeDoux, A. Lakhotia, Malware and machine learning, in: *Intelligent Methods for Cyber Warfare*, Springer, 2015, pp. 1–42.

- [7] Z. Bazrafshan, H. Hashemi, S. M. H. Fard, A. Hamzeh, A survey on heuristic malware detection techniques, in: *Information and Knowledge Technology (IKT), 2013 5th Conference on*, IEEE, 2013, pp. 113–120.
- [8] I. Basu, Malware detection based on source data using data mining: A survey, *American Journal Of Advanced Computing* 3 (1).
- [9] J. J. Barriga, S. G. Yoo, Malware detection and evasion with machine learning techniques: A survey, *International Journal of Applied Engineering Research* 12 (318).
- [10] J. Gardiner, S. Nagaraja, On the security of machine learning in malware c&c detection: A survey, *ACM Comput. Surv.* 49 (3) (2016) 59:1–59:39.
- [11] M. G. Schultz, E. Eskin, F. Zadok, S. J. Stolfo, Data mining methods for detection of new malicious executables, in: *Security and Privacy, 2001. S P 2001. Proceedings. 2001 IEEE Symposium on*, 2001, pp. 38–49.
- [12] J. Z. Kolter, M. A. Maloof, Learning to detect and classify malicious executables in the wild, *J. Mach. Learn. Res.* 7 (2006) 2721–2744.
- [13] F. Ahmed, H. Hameed, M. Z. Shafiq, M. Farooq, Using spatio-temporal information in api calls with machine learning algorithms for malware detection, in: *Proceedings of the 2nd ACM workshop on Security and artificial intelligence*, ACM, 2009, pp. 55–62.
- [14] D. H. Chau, C. Nachenberg, J. Wilhelm, A. Wright, C. Faloutsos, Polonium: Tera-scale graph mining for malware detection, in: *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2010, pp. 131–142.
- [15] I. Firdausi, C. Lim, A. Erwin, A. S. Nugroho, Analysis of machine learning techniques used in behavior-based malware detection, in: *ACT '10*, IEEE, 2010, pp. 201–203.
- [16] B. Anderson, D. Quist, J. Neil, C. Storlie, T. Lane, Graph-based malware detection using dynamic analysis, *Journal in Computer Virology* 7 (4) (2011) 247–258.

- [17] I. Santos, J. Nieves, P. G. Bringas, International Symposium on Distributed Computing and Artificial Intelligence, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011, Ch. Semi-supervised Learning for Unknown Malware Detection, pp. 415–422.
- [18] B. Anderson, C. Storlie, T. Lane, Improving malware classification: bridging the static/dynamic gap, in: Proceedings of the 5th ACM workshop on Security and artificial intelligence, ACM, 2012, pp. 3–14.
- [19] J. Yonts, Attributes of malicious files, Tech. rep., The SANS Institute (2012).
- [20] I. Santos, J. Devesa, F. Brezo, J. Nieves, P. G. Bringas, Opem: A static-dynamic approach for machine-learning-based malware detection, in: CI-SIS '12-ICEUTE' 12-SOCO', Springer, 2013, pp. 271–280.
- [21] M. Eskandari, Z. Khorshidpour, S. Hashemi, Hdm-analyser: a hybrid analysis approach based on data mining techniques for malware detection, Journal of Computer Virology and Hacking Techniques 9 (2) (2013) 77–93.
- [22] P. Vadrevu, B. Rahbarinia, R. Perdisci, K. Li, M. Antonakakis, Measuring and detecting malware downloads in live network traffic, in: Computer Security – ESORICS 2013: 18th European Symposium on Research in Computer Security, Egham, UK, September 9-13, 2013. Proceedings, Springer Berlin Heidelberg, Berlin, Heidelberg, 2013, pp. 556–573.
- [23] J. Bai, J. Wang, G. Zou, A malware detection scheme based on mining format information, The Scientific World Journal.
- [24] M. Kruczkowski, E. N. Szyrkiewicz, Support vector machine for malware analysis and classification, in: Web Intelligence (WI) and Intelligent Agent Technologies (IAT), IEEE Computer Society, 2014, pp. 415–420.
- [25] A. Tamersoy, K. Roundy, D. H. Chau, Guilt by association: large scale malware detection by mining file-relation graphs, in: Proceedings of the 20th ACM SIGKDD, ACM, 2014, pp. 1524–1533.

- [26] D. Uppal, R. Sinha, V. Mehra, V. Jain, Malware detection and classification based on extraction of api sequences, in: ICACCI, IEEE, 2014, pp. 2337–2342.
- [27] L. Chen, T. Li, M. Abdulhayoglu, Y. Ye, Intelligent malware detection based on file relation graphs, in: Semantic Computing (ICSC), 2015 IEEE International Conference on, 2015, pp. 85–92.
- [28] E. Elhadi, M. A. Maarof, B. Barry, Improving the detection of malware behaviour using simplified data dependent api call graph, *Journal of Security and Its Applications*.
- [29] Z. Feng, S. Xiong, D. Cao, X. Deng, X. Wang, Y. Yang, X. Zhou, Y. Huang, G. Wu, Hrs: A hybrid framework for malware detection, in: Proceedings of the 2015 ACM International Workshop on Security and Privacy Analytics, ACM, 2015, pp. 19–26.
- [30] M. Ghiasi, A. Sami, Z. Salehi, Dynamic VSA: a framework for malware detection based on register contents , *Engineering Applications of Artificial Intelligence* 44 (2015) 111 – 122.
- [31] M. Ahmadi, G. Giacinto, D. Ulyanov, S. Semenov, M. Trofimov, Novel feature extraction, selection and fusion for effective malware family classification, *CoRR* abs/1511.04317.
- [32] B. J. Kwon, J. Mondal, J. Jang, L. Bilge, T. Dumitras, The dropper effect: Insights into malware distribution with downloader graph analytics, in: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, ACM, 2015, pp. 1118–1129.
- [33] W. Mao, Z. Cai, D. Towsley, X. Guan, Probabilistic inference on integrity for access behavior based malware detection, in: International Workshop on Recent Advances in Intrusion Detection, Springer, 2015, pp. 155–176.
- [34] J. Saxe, K. Berlin, Deep neural network based malware detection using two dimensional binary program features, in: Malicious and Unwanted

- Software (MALWARE), 2015 10th International Conference on, IEEE, 2015, pp. 11–20.
- [35] T. Wüchner, M. Ochoa, A. Pretschner, Robust and effective malware detection through quantitative data flow graph metrics, in: *Detection of Intrusions and Malware, and Vulnerability Assessment*, Springer, 2015, pp. 98–118.
- [36] E. Raff, C. Nicholas, An alternative to ncd for large sequences, lempel-ziv jaccard distance, in: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, 2017, pp. 1007–1015.
- [37] M. Gharacheh, V. Derhami, S. Hashemi, S. M. H. Fard, Proposing an hmm-based approach to detect metamorphic malware, in: *Fuzzy and Intelligent Systems (CFIS)*, 2015, pp. 1–5.
- [38] P. Khodamoradi, M. Fazlali, F. Mardukhi, M. Nosrati, Heuristic metamorphic malware detection based on statistics of assembly instructions using classification algorithms, in: *Computer Architecture and Digital Systems (CADS)*, 2015 18th CSI International Symposium on, IEEE, 2015, pp. 1–6.
- [39] J. Upchurch, X. Zhou, Variant: a malware similarity testing framework, in: *2015 10th International Conference on Malicious and Unwanted Software (MALWARE)*, IEEE, 2015, pp. 31–39.
- [40] G. Liang, J. Pang, C. Dai, A behavior-based malware variant classification technique, *International Journal of Information and Education Technology* 6 (4) (2016) 291.
- [41] P. Vadrevu, R. Perdisci, MAXS: Scaling Malware Execution with Sequential Multi-Hypothesis Testing, in: *ASIA CCS '16, ASIA CCS '16*, ACM, New York, NY, USA, 2016, pp. 771–782.

- [42] T. Lee, J. J. Mody, Behavioral classification, in: EICAR Conference, 2006, pp. 1–17.
- [43] K. Huang, Y. Ye, Q. Jiang, Ismcs: an intelligent instruction sequence based malware categorization system, in: Anti-counterfeiting, Security, and Identification in Communication, 2009, IEEE, 2009, pp. 509–512.
- [44] Y. Park, D. Reeves, V. Mulukutla, B. Sundaravel, Fast malware classification by automated behavioral graph matching, in: Workshop on Cyber Security and Information Intelligence Research, ACM, 2010, p. 45.
- [45] Y. Ye, T. Li, Y. Chen, Q. Jiang, Automatic malware categorization using cluster ensemble, in: Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining, ACM, 2010, pp. 95–104.
- [46] G. E. Dahl, J. W. Stokes, L. Deng, D. Yu, Large-scale malware classification using random projections and neural networks, in: Acoustics, Speech and Signal Processing (ICASSP), IEEE, 2013, pp. 3422–3426.
- [47] X. Hu, K. G. Shin, S. Bhatkar, K. Griffin, Mutantx-s: Scalable malware clustering based on static features, in: USENIX Annual Technical Conference, 2013, pp. 187–198.
- [48] R. Islam, R. Tian, L. M. Batten, S. Versteeg, Classification of malware based on integrated static and dynamic features, *Journal of Network and Computer Applications* 36 (2) (2013) 646–656.
- [49] D. Kong, G. Yan, Discriminant malware distance learning on structural information for automated malware classification, in: ACM SIGKDD '13, KDD '13, ACM, New York, NY, USA, 2013, pp. 1357–1365.
- [50] S. Nari, A. A. Ghorbani, Automated malware classification based on network behavior, in: Computing, Networking and Communications (ICNC), 2013 International Conference on, IEEE, 2013, pp. 642–647.

- [51] N. Kawaguchi, K. Omote, Malware function classification using apis in initial behavior, in: Information Security (AsiaJCIS), 2015 10th Asia Joint Conference on, IEEE, 2015, pp. 138–144.
- [52] C.-T. Lin, N.-J. Wang, H. Xiao, C. Eckert, Feature selection and extraction for malware classification, *Journal of Information Science and Engineering* 31 (3) (2015) 965–992.
- [53] A. Mohaisen, O. Alrawi, M. Mohaisen, Amal: High-fidelity, behavior-based automated malware analysis and classification, *computers & security* 52 (2015) 251–266.
- [54] S. Pai, F. Di Troia, C. A. Visaggio, T. H. Austin, M. Stamp, Clustering for malware classification, *J Comput Virol Hack Tech*.
- [55] M. Bailey, J. Oberheide, J. Andersen, Z. M. Mao, F. Jahanian, J. Nazario, Automated classification and analysis of internet malware, in: *Recent advances in intrusion detection*, Springer, 2007, pp. 178–197.
- [56] U. Bayer, P. M. Comparetti, C. Hlauschek, C. Kruegel, E. Kirda, Scalable, behavior-based malware clustering, in: *NDSS*, Vol. 9, 2009, pp. 8–11.
- [57] K. Rieck, P. Trinius, C. Willems, T. Holz, Automatic analysis of malware behavior using machine learning, *Journal of Computer Security* 19 (4) (2011) 639–668.
- [58] S. Palahan, D. Babić, S. Chaudhuri, D. Kifer, Extraction of statistically significant malware behaviors, in: *Computer Security Applications Conference*, ACM, 2013, pp. 69–78.
- [59] M. Egele, M. Woo, P. Chapman, D. Brumley, Blanket execution: Dynamic similarity testing for program binaries and components, in: *USENIX Security '14*, USENIX Association, San Diego, CA, 2014, pp. 303–317.
- [60] M. Lindorfer, C. Kolbitsch, P. M. Comparetti, Detecting environment-sensitive malware, in: *Recent Advances in Intrusion Detection*, Springer, 2011, pp. 338–357.

- [61] I. Santos, F. Brezo, X. Ugarte-Pedrero, P. G. Bringas, Opcode sequences as representation of executables for data-mining-based unknown malware detection, *Information Sciences* 231 (2013) 64–82.
- [62] M. Polino, A. Scorti, F. Maggi, S. Zanero, Jackdaw: Towards Automatic Reverse Engineering of Large Datasets of Binaries, in: *Detection of Intrusions and Malware, and Vulnerability Assessment, Lecture Notes in Computer Science*, Springer International Publishing, 2015, pp. 121–143.
- [63] W. Wong, M. Stamp, Hunting for metamorphic engines, *Journal in Computer Virology* 2 (3) (2006) 211–229.
- [64] S. Attaluri, S. McGhee, M. Stamp, Profile hidden markov models and metamorphic virus detection, *Journal in Computer Virology* 5 (2) (2009) 151–169.
- [65] Z. Chen, M. Roussopoulos, Z. Liang, Y. Zhang, Z. Chen, A. Delis, Malware characteristics and threats on the internet ecosystem, *Journal of Systems and Software* 85 (7) (2012) 1650–1672.
- [66] P. M. Comar, L. Liu, S. Saha, P. N. Tan, A. Nucci, Combining supervised and unsupervised learning for zero-day malware detection, in: *INFOCOM, 2013 Proceedings IEEE*, 2013, pp. 2022–2030.
- [67] J. Sexton, C. Storlie, B. Anderson, Subroutine based detection of APT malware, *Journal of Computer Virology and Hacking Techniques* (2015) 1–9.
- [68] H.-S. Park, C.-H. Jun, A simple and fast algorithm for k-medoids clustering, *Expert Systems with Applications* 36 (2009) 3336 – 3341.
- [69] M. Graziano, D. Canali, L. Bilge, A. Lanzi, D. Balzarotti, Needles in a haystack: Mining information from public dynamic analysis sandboxes for malware intelligence, in: *USENIX Security '15*, 2015, pp. 1057–1072.

- [70] M. Asquith, Extremely scalable storage and clustering of malware meta-data, *Journal of Computer Virology and Hacking Techniques* (2015) 1–10.
- [71] M. Egele, T. Scholte, E. Kirda, C. Kruegel, A survey on automated dynamic malware-analysis techniques and tools, *ACM computing surveys (CSUR)* 44 (2) (2012) 6.
- [72] M. Siddiqui, M. C. Wang, J. Lee, Detecting internet worms using data mining techniques, *Journal of Systemics, Cybernetics and Informatics* (2009) 48–53.
- [73] A. Caliskan-Islam, R. Harang, A. Liu, A. Narayanan, C. Voss, F. Yamaguchi, R. Greenstadt, De-anonymizing programmers via code stylometry, in: *USENIX Security '15*, USENIX Association, 2015, pp. 255–270.
- [74] S. Srakaew, W. Piyanuntcharatsr, S. Adulkasem, On the comparison of malware detection methods using data mining with two feature sets, *Journal of Security and Its Applications* 9 (2015) 293–318.
- [75] J. Jang, D. Brumley, S. Venkataraman, Bitshred: feature hashing malware for scalable triage and semantic analysis, in: *Computer and communications security*, ACM, 2011, pp. 309–320.
- [76] F. E. Allen, Control flow analysis, in: *Proceedings of a Symposium on Compiler Optimization*, ACM, New York, NY, USA, 1970, pp. 1–19.
- [77] D. Kirat, L. Nataraj, G. Vigna, B. Manjunath, Sigmal: A static signal processing based malware triage, in: *Proceedings of the 29th Annual Computer Security Applications Conference*, ACM, 2013, pp. 89–98.
- [78] R. Tian, L. M. Batten, S. C. Versteeg, Function length as a tool for malware classification, in: *Malicious and Unwanted Software, 2008. MALWARE 2008. 3rd International Conference on*, 2008, pp. 69–76.
- [79] N. Miramirkhani, M. P. Appini, N. Nikiforakis, M. Polychronakis, Spotless sandboxes: Evading malware analysis systems using wear-and-tear

- artifacts, in: 2017 IEEE Symposium on Security and Privacy (SP), 2017, pp. 1009–1024.
- [80] J. Karpin, A. Dorfman, Crypton - Exposing Malware’s Deepest Secrets, <https://recon.cx/2017/montreal/resources/slides/RECON-MTL-2017-crypton.pdf>, last accessed: 2018-10-18 (2017).
- [81] T. Blazytko, M. Contag, C. Aschermann, T. Holz, Syntia: Synthesizing the semantics of obfuscated code, in: 26th USENIX Security Symposium (USENIX Security 17), USENIX Association, 2017, pp. 643–659.
- [82] V. Kotov, M. Wojnowicz, Towards generic deobfuscation of windows API calls, CoRR abs/1802.04466.
- [83] Y. Liao, R. Cai, G. Zhu, Y. Yin, K. Li, Mobilefindr: Function similarity identification for reversing mobile binaries, in: ESORICS (1), Vol. 11098 of Lecture Notes in Computer Science, Springer, 2018, pp. 66–83.
- [84] R. Baldoni, E. Coppa, D. C. D’Elia, C. Demetrescu, I. Finocchi, A survey of symbolic execution techniques, ACM Comput. Surv. 51 (3).
- [85] I. Guilfanov, Decompilers and beyond, Black Hat USA.
- [86] A. Rosseau, R. Seymour, Finding Xori: Malware Analysis Triage with Automated Disassembly, <https://i.blackhat.com/us-18/Wed-August-8/us-18-Rousseau-Finding-Xori-Malware-Analysis-Triage-With-Automated-Disassembly.pdf>, last accessed: 2018-10-14 (2018).
- [87] E. Schulte, J. Ruchti, M. Noonan, D. Ciarletta, A. Loginov, Evolving exact decompilation, in: Binary Analysis Research (BAR), 2018, 2018.
- [88] Vxheaven, <https://github.com/opsxcq/mirror-vxheaven.org>, accessed: 2018-06-03.
- [89] Offensive Computing, <http://www.offensivecomputing.net>, accessed: 2018-06-03.

- [90] Malicia Project, <http://malicia-project.com>, accessed: 2018-06-03.
- [91] ThreatTrack, <https://www.threattrack.com/malware-analysis.aspx>, accessed: 2018-06-03.
- [92] Anubis, <https://seclab.cs.ucsb.edu/academic/projects/projects/anubis/>, accessed: 2018-06-03.
- [93] L. Bilge, D. Balzarotti, W. Robertson, E. Kirda, C. Kruegel, Disclosure: detecting botnet command and control servers through large-scale netflow analysis, in: ACSAC '12, ACM, 2012, pp. 129–138.
- [94] H. Sanders, Garbage In, Garbage Out - How purportedly great ML models can be screwed up by bad data , <https://www.blackhat.com/docs/us-17/wednesday/us-17-Sanders-Garbage-In-Garbage-Out-How-Purportedly-Great-ML-Models-Can-Be-Screwed-Up.pdf>, last accessed: 2018-10-18 (2017).
- [95] B. Miller, A. Kantchelian, S. Afroz, R. Bachwani, R. Faizullahoy, L. Huang, V. Shankar, M. Tschantz, T. Wu, G. Yiu, et al., Back to the future: Malware detection with temporally consistent labels, CORR.
- [96] R. Jordaney, K. Sharad, S. K. Dash, Z. Wang, D. Papini, I. Nouretdinov, L. Cavallaro, Transcend: Detecting concept drift in malware classification models, in: 26th USENIX Security Symposium (USENIX Security 17), USENIX Association, Vancouver, BC, 2017, pp. 625–642.
URL <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/jordaney>
- [97] R. Harang, F. Ducau, Measuring the Speed of the Red Queens Race, <https://i.blackhat.com/us-18/Wed-August-8/us-18-Harang-Measuring-the-Speed-of-the-Red-Queens-Race.pdf>, last accessed: 2018-10-18 (2018).
- [98] VirusTotal, <https://www.virustotal.com>, accessed: 2018-06-03.

- [99] Malwr, <https://malwr.com>, accessed: 2018-06-03.
- [100] M. Ruthven, A. Blaich, Fighting targeted malware in the mobile ecosystem, <https://www.blackhat.com/docs/us-17/wednesday/us-17-Ruthven-Fighting-Targeted-Malware-In-The-Mobile-Ecosystem.pdf>, last accessed: 2018-10-16 (2017).
- [101] G. Laurenza, L. Aniello, R. Lazzeretti, R. Baldoni, Malware triage based on static features and public apt reports, in: S. Dolev, S. Lodha (Eds.), Cyber Security Cryptography and Machine Learning, Springer International Publishing, Cham, 2017, pp. 288–305.
- [102] M. Howard, A. Pfeffer, M. Dalai, M. Reposa, Predicting signatures of future malware variants, in: 2017 12th International Conference on Malicious and Unwanted Software (MALWARE), 2017, pp. 126–132.
- [103] V. Juzonis, N. Goranin, A. Cenys, D. Oliner, Specialized genetic algorithm based simulation tool designed for malware evolution forecasting, *Annales Universitatis Mariae Curie-Sklodowska, sectio AI-Informatica* 12 (4).
- [104] H. Pomeranz, Detecting malware with memory forensics, http://www.deer-run.com/~hal/Detect_Malware_w_Memory_Forensics.pdf, last accessed: 2018-05-14 (2012).
- [105] A. Damodaran, F. Di Troia, C. A. Visaggio, T. H. Austin, M. Stamp, A comparison of static, dynamic, and hybrid analysis for malware detection, *Journal of Computer Virology and Hacking Techniques* (2015) 1–12.
- [106] Y. Kulakov, Mazewalker, <https://recon.cx/2017/montreal/resources/slides/RECON-MTL-2017-MazeWalker.pdf> (2017).
- [107] G. Laurenza, D. Ucci, L. Aniello, R. Baldoni, An architecture for semi-automatic collaborative malware analysis for cis, in: Dependable Systems and Networks Workshop, 2016 46th Annual IEEE/IFIP International Conference on, IEEE, 2016, pp. 137–142.