

# UC Riverside

## UC Riverside Previously Published Works

### Title

MAB-Malware: A Reinforcement Learning Framework for Attacking Static Malware Classifiers

### Permalink

<https://escholarship.org/uc/item/38m6n4vm>

### Authors

Song, Wei  
Li, Xuezixiang  
Afroz, Sadia  
[et al.](#)

### Publication Date

2020-03-06

Peer reviewed

# Automatic Generation of Adversarial Examples for Interpreting Malware Classifiers

Wei Song<sup>1</sup>, Xuezixiang Li<sup>1</sup>, Sadia Afroz<sup>2,3</sup>, Deepali Garg<sup>2</sup>, Dmitry Kuznetsov<sup>2</sup>, and Heng Yin<sup>1</sup>

<sup>1</sup>UC Riverside , {wsong008, xli287}@ucr.edu, heng@cs.ucr.edu

<sup>2</sup>Avast , {sadia.afroz, deepali.garg, kuznetsov}@avast.com

<sup>3</sup>ICSI , sadia@icsi.berkeley.edu

## Abstract

Recent advances in adversarial attacks have shown that machine learning classifiers based on static analysis are vulnerable to adversarial attacks. However, real-world antivirus systems do not rely only on static classifiers, thus many of these static evasions get detected by dynamic analysis whenever the malware runs. The real question is to what extent these adversarial attacks are actually harmful to the real users? In this paper, we propose a systematic framework to create and evaluate realistic adversarial malware to evade real-world systems. We propose new adversarial attacks against real-world antivirus systems based on code randomization and binary manipulation, and use our framework to perform the attacks on 1000 malware samples and test four commercial antivirus software and two open-source classifiers. We demonstrate that the static detectors of real-world antivirus can be evaded 24.3%–41.9% of the cases and often by changing only one byte. We also find that the adversarial attacks are transferable between different antivirus up to 16% of the cases. We also tested the efficacy of the complete (i.e. static + dynamic) classifiers in protecting users. While most of the commercial antivirus use their dynamic engines to protect the users' device when the static classifiers are evaded, we are the first to demonstrate that for one commercial antivirus, static evasions can also evade the offline dynamic detectors and infect users' machines. Our framework can also help explain which features are responsible for evasion and thus can help improve the robustness of malware detectors.

## 1 Introduction

Malware attacks continue to be one of the most pressing security issues users face today. Recent research showed that during the first nine months of 2019, at least 7.2 billion malware attacks and 151.9 million ransomware attacks have been reported.<sup>1</sup> The traditional signature-based method cannot keep

<sup>1</sup><https://www.msspalert.com/cybersecurity-research/sonic-wall-research-malware-attacks-2019/>

up with this rampant inflation of malware. Hence commercial antivirus companies started using machine learning [7, 40]. Machine-learning-based detectors are scalable and efficient at protecting against the huge influx of malware, which is why since the first paper in 2001 on detecting malware using machine learning [55], there has been an explosion of academic research papers on predicting malicious content using machine learning, many of them flaunting high accuracy and being able to detect new malware unseen during the training [6, 12, 48, 51, 54].

On the flip side, research has also demonstrated that machine-learning-based detectors can be easily evaded by making trivial changes to a malware [2, 4, 9, 10, 15–17, 20, 22, 23, 25–27, 29, 31–35, 38, 39, 45, 47, 52, 53, 57, 58, 60, 61]. The results in adversarial attacks are especially significant in the malware domain where there are already existing adversaries trying to evade detection using commercial packers and crypters [4, 18].

We argue that research on the adversarial attack against malware detection has not yet convincingly demonstrated that adversarial attacks are harmful to users. Specifically, we notice the following shortcomings:

1. **The attacks do not generate actual malware:** The majority of the prior attacks are performed only on the feature vectors, without actually modifying the malware binaries [2, 10, 21]. Attacks performed in the feature space might not be able to translate to actual binaries. Some of these adversarial modifications can corrupt malware format or affect functionality.
2. **The attacks do not verify the modified malware:** Some previous research [4] uses transformations that, in theory, should preserve functionality but in practice (e.g. obfuscation tricks) even supposedly safe transformations break functionality. Therefore, it is crucial to verify the malicious behaviors of adversarial samples. However, no paper on Windows malware systematically checked if the functionalities of the modified malware were still preserved. Moreover, all the real world antivirus products

use both static and dynamic classifiers. Demonstrating only static evasion is not enough to show real-world ramifications of adversarial attacks because static adversarial samples can still be blocked by dynamic classifiers based on their runtime behaviors.

3. **The attacks consider unrealistic adversaries:** Among 14 papers on adversarial attacks on windows malware, 9 papers [3, 10, 14, 26, 28, 30, 35, 56, 58] performed white-box attacks and 5 papers [4, 18, 20, 52, 53] performed blackbox attacks. The whitebox attack scenario on commercial AVs is unrealistic because attackers cannot obtain insider knowledge. Even for the blackbox attack scenario, some papers assume that the adversaries know the features used by the target model [20, 21], and some assume that they have access to the confidence score of decisions from a model under attack [34, 61]. However, commercial antivirus products do not expose features or return the confidence score.
4. **The attacks do not explain why evasion happened:** We found only one paper tried interpreting the decisions of a blackbox deep neural network model, MalConv [59], using integrated gradients [15]. However, the paper manipulated the byte-level features without manipulating the actual file. Their explanation also cannot be used for other malware detectors that do not use only byte-level features.

Moreover, adversarial examples provide a unique opportunity to understand how malware detectors work. None of the current work paper investigated which feature changes lead to the adversarial evasions on commercial antivirus. Identifying the fragile features responsible for adversarial evasions can help us improve the robustness of malware detectors.

In this paper, we seek to answer the following research questions:

- **RQ1:** How robust are the commercial antivirus engines against adversarial examples of windows malware?
- **RQ2:** How can we explain the commercial antivirus engines using adversarial examples?
- **RQ3:** Given that antivirus often employs a combination of static and dynamic classifiers, how harmful are adversarial attacks to the end users?

To answer these questions, we propose an open-source automatic framework that provides generic actions to modify windows malware binaries to evade a variety of static malware detectors (Section 3). This framework integrates a sandbox to verify the functionalities of generated evasive examples. We use our framework to test four top antivirus engines according to PCMag [46], one signature-based open-source AV (ClamAV), and one machine-learning-based classifier (EMBER). We find that even with a completely blackbox attack, where an adversary only has external access to a system, the static

classifiers of commercial antivirus engines can be evaded 24.3%–41.9% of the cases (Section 4).

Generating these adversarial examples are surprisingly easy. We find that more than 80% of the adversarial samples are generated only using a single action, 2% - 19% of which only need to change 1 byte (e.g., append 1 byte in a particular section, or change one byte of section name) (Section 4.3). On average, a blackbox adversary has to spend about 13–61 seconds to generate these adversarial malware samples. These adversarial attacks are transferable among different malware detectors. We find that the adversarial samples generated using one malware detector can evade other malware detectors up to 16% of the cases (Section 4.3.3).

Although our adversarial examples are not designed to evade dynamic detection employed by antivirus products, we indeed observed that one out of four commercial AV products failed to detect our adversarial examples during execution, indicating flaws in its dynamic detection component (Section 4.3.4).

In summary, our contributions are as follows:

- **An open-source<sup>2</sup> framework for generating adversarial examples for commercial antivirus:** Our framework can automatically generate and verify functional adversarial malware.
- **A novel algorithm for explaining why evasion happens:** We propose a novel explanation algorithm to reveal the root cause for evasive samples for many different kinds of malware detectors. It identifies the minimum number of features needed to be changed to alter the decision of a malware detector. Our technique is generic for any malware detectors.
- **Measuring harm to end-users:** We are the first to measure how harmful adversarial examples can be to end-users. Our results show that for one out of four antivirus engines we tested, we found that adversarial ransomware samples that evade static classifier can also evade its offline dynamic classifier and encrypt user files. This result shows that to demonstrate real harm, research in adversarial attacks should focus on evading the entire antivirus pipelines, instead of only the static classifiers. On the other hand, antivirus systems should provide offline dynamic detection to protect users against static adversarial attacks.

## 2 Problem

### 2.1 Threat Model

We follow the study by Carlini et al. [8] to describe our threat model, from three aspects: adversarial goal, adversarial capabilities, and adversarial knowledge.

<sup>2</sup>[https://github.com/bitsecurerlab/adversarial\\_malware.git](https://github.com/bitsecurerlab/adversarial_malware.git)

**Adversarial Goal.** The adversary’s goal is to manipulate malware samples to evade the detection of static PE malware classifiers. Other types of malware like PDF malware or Android malware are not within the scope of this study. This is an untargeted attack because we only consider a binary classification (benign or malicious) not specific malware families in this classification task and we are only interested in causing the malicious samples to be classified as benign.

**Adversarial Capabilities.** In this work, we assume that the adversary does not have access to the training phase of the malware classifiers. For instance, the adversary cannot inject poisonous data in the training dataset.

In addition, the adversary cannot arbitrarily change the input data. In most scenarios of adversarial attacks, such as image recognition, the adversary is required to make only “small” changes to the original sample to keep the manipulation visually imperceptible. However, when attacking malware classification, the restriction is not on the number or size of changes, but on the preservation of malicious functionality. If “small” changes on a malware sample indeed confuse a malware classifier but prevent the malware from acting maliciously, this manipulation is not considered successful.

**Adversarial Knowledge.** Based on what knowledge the adversary can obtain, the adversarial attack can be divided into two types: 1) whitebox attacks in which the adversary has unlimited access to the model; and 2) blackbox attacks in which the adversary has no knowledge about the model and can obtain the classification results only through a limited number of attempts. A classification result can be a score or simply a label.

In this work, we consider an adversary with only blackbox access. The adversary does not know anything about the internals of the deployed classifiers, can perform a limited number of attempts to the classifiers and can observe the classifiers’ actions when the samples are considered malicious. Commercial antivirus systems delete malicious samples. If a file gets deleted, the adversary can be certain that the AV considered the file malicious. However, if the file does not get deleted the adversary cannot be certain that the file is considered benign because it might get uploaded to the cloud for in-depth processing. For the purpose of our experimental setup where we only use offline antivirus systems, we assume that if a file is not deleted by an AV, the AV considers it as benign.

## 2.2 Problem Definition

We aim to automatically generate adversarial examples for malware classifiers and explain the root cause of the evasions. The problem can be split into two sub-problems: adversarial example generation and feature interpretation.

**Adversarial Example Generation.** We aim to manipulate a malware sample without breaking its malicious function-

alities, such that the malware classifier misclassify it as a benignware.

Let malware classifier be a mapping  $f$  that maps from a binary sample  $s$  to a classification label. It outputs 1 if  $s$  is a malware, and outputs 0 if  $s$  is a benignware. We then design the action set  $A = \{a_1, a_2, \dots a_n\}$  that can be used to safely manipulate the malware samples. We also design a verifier function  $v(s, s')$  to check if the behaviors of a malware sample  $s$  are preserved for modified sample  $s'$ . It outputs 1 if the modified sample  $s'$  preserves the functionalities of the original sample  $s$ . Otherwise, it outputs 0 if the functionalities are altered.

Our goal is to select an action sequence  $seq = \{a_1, a_2, \dots a_m\}$  from Action Set  $A$  such that when applied to sample  $s$ , the probability of  $f(s') = 0$  and  $v(s, s') = 1$  is maximized.

**Feature Interpretation.** We want to pinpoint why a classifier misclassified an adversarial sample by identifying the group of features that are responsible for changing the decision. However, we do not know the features used by black box malware classifiers. All we can control is the manipulation actions used to generate adversarial examples. Let  $s$  be a malware sample,  $seq = \{a_1, a_2, \dots a_m\}$  be the action sequence used to generate an adversarial example  $s'$ , and  $F = \{f_1, f_2, \dots f_k\}$  be the changed features between  $s$  and  $s'$ . The problem is, given  $s$  and  $seq$ , how to find out a minimal subset of  $F$  that is essential to the evasion?

## 3 Methodology

### 3.1 Action Set and Features

At first, we define the actions that can be applied to malware to create adversarial samples, as shown in Table 1. Every action manipulates a set of features that a classifier uses to detect malware. For a blackbox model, we do not know the exact features the classifier is using. However, we can make an educated guess, by examining open-source malware detectors such as EMBER and ClamAV, about the categories of features necessary for detecting malware. We assume three large categories of features: hash-based signatures (file hash and section hash), rule-based signatures (section count, section name, section padding, debug info, checksum, certificate, and code sequences) and data distribution based features (byte histogram, and byte entropy histogram).

**Macro-actions.** We use the actions proposed by Anderson et al. [4] that do not break the functionality of malware samples (Table 1). We also adopt a code randomization action (CR) from Pappas et al. [44]. It is a defense method originally proposed to prevent Return Oriented Programming (ROP) attack. This approach breaks the chains of code gadgets that ROP attacks use by making narrow-scope code transformations statically. Since all transformations are equivalent and do not alter the locations of the basic blocks, it is considered

Table 1: Action Set is the collection of macro and micro actions that are used to generate adversarial examples. Micro actions are modified macro actions such that the number of features impacted is reduced.

Type	Abbr	Name	Description
Macro	OA	Overlay Append	Appends benign contents at the end of a binary
	SP	Section Append	Appends random bytes to the unused space at the end of a section.
	SA	Section Add	Adds a new section with benign contents.
	SR	Section Rename	Change the a section name to a name in benign binaries.
	RC	Remove Certificate	Zero out the signed certificate of a binary.
	RD	Remove Debug	Zero out the debug information in a binary.
	BC	Break Checksum	Zero out the checksum value in the optional header.
	CR	Code Randomization	Replace instruction sequence with semantically equivalent one
Micro	OA1	Overlay Append	Appends 1 byte at the end of a binary
	SP1	Section Append 1 Byte	Appends 1 byte to the unused space at the end of a section.
	SA1	Section Add 1 Byte	Adds a new section with 1 byte content.
	SR1	Section Rename 1 Byte	Change 1 byte of a section name.
	CPI	Code Section Append 1 Byte	Appends 1 byte to the unused space at the end of the code section.

Table 2: Affected Features by Actions. Each action that is applied to a binary can impact multiple features.

		CR	OA	SP	SA	SR	RC	RD	BC	OA1	SP1	SA1	SR1	CPI
Hash-Based Signatures	$F_1$ : File Hash	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	$F_2$ : Section Hash	✓		✓				✓			✓			✓
Rule-based Signatures	$F_3$ : Section Count				✓							✓		
	$F_4$ : Section Name					✓							✓	
	$F_5$ : Section Padding			✓										
	$F_6$ : Debug Info							✓						
	$F_7$ : Checksum								✓					
	$F_8$ : Certificate							✓						
	$F_9$ : Code Sequence	✓												
Data Distribution	$F_{10}$ : Data Distribution		✓		✓									

a safe randomization method for binaries. In our case, we use it as an adversarial attack technique to break the code patterns that may serve as rule-based features for detecting malware.

**Micro-actions.** Micro is a relative concept. If an action  $a$  changes feature set  $F = \{f_1, f_2, \dots, f_k\}$  of a malware sample, then another action that only changes a subset of  $F$  is a micro-action of  $a$ . We implement 5 micro-actions for macro-action OA, SP, SA, SR and CR. Table 2 shows all the actions used in our framework and the corresponding affected features for each action. Take the action Section Append (SP) as an example, by looking up Table 2, we get that SP action affects features File Hash ( $F_1$ ), Section Hash ( $F_2$ ) and Rule-Based Signatures in the section padding data ( $F_5$ ).  $F_1, F_2$  are affected if any modification is made to the file and section content.  $F_5$  may be affected if SP action modifies the padding content that may contain body-based signatures.

### 3.2 Workflow

The workflow includes 3 steps: adversarial example generation, action sequence minimization, and feature interpretation (Figure 1).

Firstly, to generate adversarial samples we randomly select and apply macro-actions to change the original samples. Macro actions affect a large number of features, thus ensures the probability of the original sample crossing the decision boundary is higher.

Secondly, we remove unnecessary macro-actions from the action sequence to generate a minimized evasive sample. This step is necessary to identify the cause of evasion. After action sequence minimization, we know the relative effectiveness of the macro-actions for causing evasions. Due to some obfuscation tricks in malware samples, supposedly safe actions may still break functionality. Removing redundant actions also helps to reduce the possibility of generating broken samples.

Thirdly, to gain fine-grained insights into the cause of the evasion, we need to disentangle the macro-actions into micro-actions and find the precise reason for the evasion. We traverse

every macro-action in the minimized action sequence, and try to replace it with different micro-actions. In this way, more new samples are generated around the adversarial sample. Using the feature changes of each new sample and its classification label, we can explain how the classification decision is made, and which features are essential to the evasion.




Figure 1: **Workflow.** The seed malware sample  $S$  resides in the malicious region of the feature space. We perform a sequence of single actions  $a_1, a_2$  and  $a_3$  until the generated sample  $S_{123}$  successfully reaches the benign region.  $S_{123}$  is an adversarial example. In the feature interpretation phase, first, we remove useless actions. The action  $a_2$  is essential, because by removing action  $a_2$ , the generated sample  $S_{13}$  is no longer evasive. The action  $a_1$  is useless because by removing action  $a_1$ , the generated sample  $S_{23}$  has no effect in the classifier’s decision. Then we disentangle macro actions into micro ones.  $a_2$  can be replaced with micro-actions  $a'_2, a''_2$ .  $a_3$  can be replaced with micro-actions  $a'_3$  or  $a''_3$ . We generated three samples  $S_{2'3'}$ ,  $S_{2'3''}$  and  $S_{2''3''}$ . Finally, we have an adversarial sample  $S_{2'3''}$  with a minimized action sequence ( $a'_2, a'_3$ ).

### 3.3 Framework

Our framework consists of five modules: Binary Rewriter, Static Classifier, Action Sequence Minimizer, Verifier and Feature Interpreter (Figure 2). As inputs, the framework uses original malware samples, a set of actions that can be applied to a malware sample, and a set of content that can be added when applying actions that require the addition of new content. For each input malware sample, the Binary Rewriter chooses one action from the action set and applies this action to the malware sample and rewrites the binary. The Static Classifier tests this rewritten binary to determine if it is classified as malicious. The process is repeated until the malicious sample is classified as benign. The Action Sequence Minimizer reduces the redundant actions that have no contribution to evasions. The Verifier validates the minimized evasive samples. If the functionalities of the evaded sample are similar to the original sample, we call it the adversarial sample. Lastly, the Feature Interpreter disentangles the large actions into smaller micro-actions to ensure a small number of features are affected. The next couple of subsections elaborate on each module in detail.

#### 3.3.1 Binary Rewriter

The Binary Rewriter generates multiple variants of an original malware sample and interacts with static classifiers to test if the generated samples evade detection. We design two types of manipulations to rewrite a malware file: in-place code randomization and code-agnostic manipulation that changes the malware file headers, sections and overlay data.

**Adversarial Example Generation.** With the two types of manipulation actions, we now describe how to generate the evasive examples (see Algorithm 1).

---

#### Algorithm 1: Adversarial Example Generation

---

**Input** : malware sample  $s$ , malware classifier  $f$   
**Output** : evasive examples  $s_{eva}$

- 1 Initialize action set  $A$ , max actions on single sample  $max\_attempt$ , max action sequence length  $max\_length$ , wait time on VM  $T$ , probability  $p$  to select saved action content.
- 2  $s' \leftarrow s$
- 3  $W \leftarrow \text{getActionWeights}(f)$
- 4 **for**  $\_ \in \{1, \dots, max\_attempt\}$  **do**
- 5      $action \leftarrow \text{randomSelect}(A, W)$
- 6      $content \leftarrow \text{selectContent}(p)$
- 7      $s' \leftarrow \text{applyAction}(s', action, content)$
- 8      $\text{copyToVM}(s')$
- 9     **while**  $True$  **do**
- 10          $isExt \leftarrow \text{checkExistedOnVM}(s')$
- 11          $isMod \leftarrow \text{checkModifiedByVM}(s')$
- 12         **if**  $!isExt$  **or**  $isMod$  **then**
- 13             **break**
- 14         **else**
- 15             **if**  $curTime - s'.creatTime \geq T$  **then**
- 16                  $s_{eva} \leftarrow s'$
- 17                 **return**  $s_{eva}$
- 18             **end**
- 19         **end**
- 20     **end**
- 21     **if**  $len(\text{getActionSeq}(s')) \geq max\_length$  **then**
- 22          $s' \leftarrow s$
- 23     **end**
- 24 **end**

---

For a seed malware sample  $s$ , we aim to generate a variant of the sample as  $s_{eva}$ , such that it evades the target classifier ( $f(s_{eva}) = 0$ ) while preserving the original malicious behaviors ( $v(s, s_{eva}) = 1$ ). We apply various actions iteratively until we get an evasive sample or the total number of attempts exceeds  $max\_attempt$ . For each attempt, an action is randomly selected according to weights  $W$  from the action set. If the selected action requires content (for example, a new content is required for the Section Add action), we select content from the content set (generated in Algorithm 3) for that action with probability  $p$  or generate random content with probability  $1 - p$ . The Binary Rewriter generates a new binary with the chosen actions applied to the original samples. This rewritten

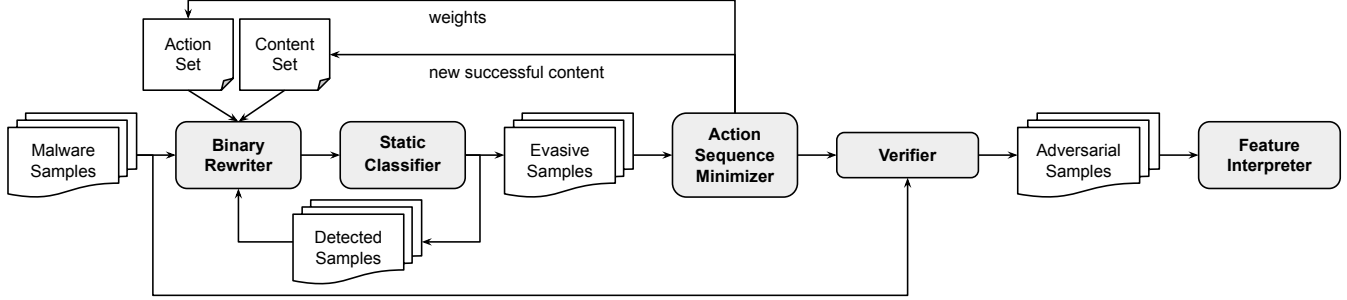


Figure 2: **Framework.** This framework takes as input original malware samples, a set of actions that can be applied to a malware sample, and a set of content that can be added when applying actions that require adding new content. As output, it produces evasive samples that are marked benign by the Static Classifier and that preserve the functionality of the original sample. Additionally, these evasive samples are further analyzed to learn about features of the Static Classifier.

sample is copied to the shared folder of a virtual machine (VM) with Static Classifier installed. If the Static Classifier detects the sample, it either deletes the sample or disinfects it. A sample is considered “evasive” if it is neither detected nor disinfects (MD5 checksum does not change) by the Static Classifier after a fixed time  $T$ .

We restrict the number of actions to a certain length ( $max\_length$ ) to ensure that the size of the rewritten binary does not grow too large for our sandbox. The sandbox we are using, the Cuckoo sandbox, does not support the analysis of extremely large binaries. Also, it reduces the possibility that the generated sample is broken by one of these actions.

### 3.3.2 Action Sequence Minimizer

The Action Sequence Minimizer aims to remove unnecessary actions from the action set to produce a “minimized” sample using only actions required to make the original malware sample evasive. Details about this module can be found in Algorithm 3 in Appendix A.

With minimized samples, we design two strategies to reuse effective actions: 1) increasing the weights of effective actions to increase their probability of getting selected during the random selection process of the Binary Rewriter, 2) reusing both the action sequence and associated content to new samples.

For a successful action sequence  $q = (a_1, a_2, \dots)$  of an evasive malware sample  $s_{eva}$ , we traverse every action  $a_i$  in  $q$  trying to remove the action and get a shorter sequence. Then, we apply the new sequence on the original sample  $s$  to generate a new sample  $s'$ . To replay the same actions as in generating  $s_{eva}$ , we first find the action log using the action index, then parse the log to get the action content, such as the new section content in action Section Add. We apply all actions in the shorter sequence with corresponding action content. The function  $isEvasive(f, s')$  (the same procedure between line 9 and 20 in Algorithm 1) is used to determine whether  $s'$  can still evade the target classifier  $f$ . If so, we consider the action  $a_i$  useless in generating evasive sample  $s_{eva}$ , and remove  $a_i$

from the action sequence. If not, we consider the action  $a_i$  as an essential action. We need to increase the weight of  $a_i$  in  $W$  to select  $a_i$  with a higher possibility in the future, and also save the corresponding content to reuse later. When all the useless actions are removed from the sequence  $q$ , the remaining action sequence is the minimized sequence, and the corresponding sample  $s_{min}$  is the minimized sample.

To calculate the weights of the actions, for each classifier we collect a set of minimized sequences. For each action  $a_i \in A$ , we count the number of occurrences in sequences as  $N_{a_i}$ . Then, the weight of an action  $a_i$  is calculated as  $\frac{N_{a_i}}{\sum_{j=1}^n N_{a_j}}$ ,  $a_i, a_j \in A$ . The weights are used in Binary Rewriter when randomly selecting actions.

### 3.3.3 Verifier

The Verifier is used in two phases: 1) to select the original malware sample to ensure that the samples demonstrate malicious behavior and 2) to check whether the functionalities of the evasive sample are still intact.

The Verifier uses the Cuckoo sandbox to verify the functionalities of the malware. The Cuckoo sandbox collects the behaviors of samples and converts them into readable descriptive signatures. Each signature is a text string summarizing one specific behavior of a sample. Some examples of signature are “Drops 2302 unknown file mime types indicative of ransomware writing encrypted files back to disk (50 out of 2280 events)”, or “Resumed a suspended thread in a remote process potentially indicative of process injection (4 events).” It also provides a maliciousness score based on the behavior. The Verifier considers a sample as malicious if its maliciousness score is higher than a threshold, which is 7 for our experiments. To compare the malicious behaviors of our minimized samples with the original samples, we compare the behaviors of the minimized sample with the original one. If a minimized sample (classified as benign) has the same behaviors as the original one, we consider it as an evasive

example.

Specifically, suppose the original sample  $s$  has a signature set  $G$ , and the evasive sample  $s_{eva}$  has a signature set  $G_{eva}$ . We check each signature in  $G$  and search in  $G_{eva}$  to see if the same signature exists in  $G_{eva}$ . We count the total number of the similar signatures in  $G$  and  $G_{eva}$  as  $hits(G, G_{eva})$ . The behavior similarity between two samples is defined as  $sim(G, G_{eva}) = \frac{hits(G, G_{eva})}{\#G}$ . If the  $sim(G, G_{eva}) \geq 0.8$ , the Verifier considers the two samples  $s$  and  $s_{eva}$  share the same behaviors. Otherwise, the Verifier considers the manipulation has changed the behaviors of the original sample. We do not set the threshold of similarity to 1 because the behaviors of binaries may vary slightly during different executions.

---

### Algorithm 2: Feature Interpretation

---

```

Input :malware sample  $s$ , adversarial sample  $s_{adv}$ 
Output:essential feature set for the evasion
          $essential\_features$ 
1  $q \leftarrow getActionSeq(s_{min})$ 
2  $action\_idxs \leftarrow [0, \dots, len(q) - 1]$ 
3  $kept\_idx\_micro \leftarrow []$ 
4  $essential\_features \leftarrow []$ 
5 foreach  $idx \in action\_idxs$  do
6    $action \leftarrow q[idx]$ 
7    $alter\_micros \leftarrow getAlterMicroActions(action)$ 
8    $q' \leftarrow q$ 
9   foreach  $micro \in alter\_micros$  do
10     $q'[idx] \leftarrow micro$ 
11    foreach  $kept\_idx, kept\_micro \in kept\_idx\_micros$  do
12       $q'[kept\_idx] \leftarrow kept\_micro$ 
13    end
14     $s' \leftarrow s$ 
15    foreach  $action \in q'$  do
16       $content \leftarrow parseLog(action.idx)$ 
17       $s' \leftarrow applyAction(s', action, content)$ 
18    end
19    if  $isEvasive(f, s')$  then
20       $kept\_idx\_micros.append(idx, micro)$ 
21       $feature \leftarrow actionFeatureMapping(micro)$ 
22       $essential\_features.append(feature)$ 
23    end
24  end
25 end
26 return  $essential\_features$ 

```

---

### 3.3.4 Feature Interpreter

A single action may change multiple features and can have a varying impact on a classifier. From an attacker's perspective, we only know if a particular action causes an evasion, but we do not know which feature change is responsible for the evasion. To gain insights into the cause of the evasion, we need to break one large action into several micro-actions.

Take action Section Append (SP) as an example, Figure 3

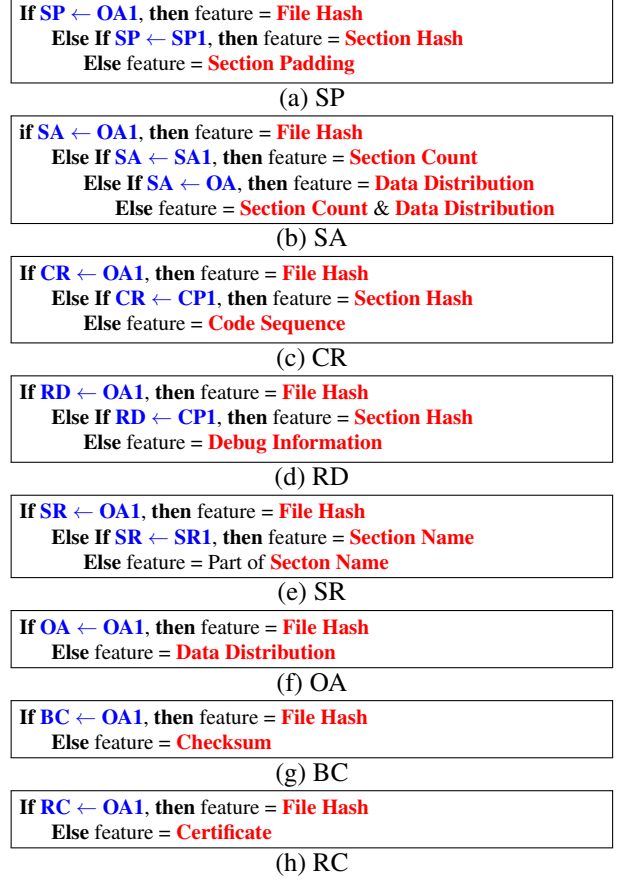


Figure 3: Decision rules are used to map actions to feature space, which enables us to understand the cause of evasion.

(a) shows how we break it into micro-actions. First, by looking up Table 2 to see that SP changes feature set  $F$  consisting of features related to File Hash ( $F_1$ ), Section Hash ( $F_2$ ) and Section padding ( $F_5$ ). The actions that only change a subset of  $F$  are one-byte overlay append (OA1) that changes  $F_1$  and one-byte section append (SP1) that changes  $F_1, F_2$ . They are alternative actions of SP. Starting from the minimum change, we try to replace SP with OA1 and check if the file is still evasive (denoted as  $SP \leftarrow OA1$ ). If so, we can conclude the essential effect of SP is breaking file hash ( $F_1$ ). If not, we then replace SP with SP1. If still evasive, the essential effect of SP is breaking section hash ( $F_2$ ). Otherwise, the evasion is breaking signatures in Section Padding ( $F_5$ ).

The Feature Interpreter accepts a sample  $s$  and corresponding adversarial sample  $s_{adv}$  as input, and try to replace each action used in  $s_{adv}$  one after another with alternative actions if possible (Algorithm 2). Specifically, for each action, it first finds all the alternative actions from Table 2, and then replaces the action with one micro action at a time. If after the replacements, the new sample is still evasive, we then keep the new minimized sample. Otherwise, we try the rest of the actions and repeat the replacement and validation process



until all macro-actions are processed.

## 4 Evaluation

### 4.1 Experiment Setup

**Dataset:** In this paper, we generate adversarial examples for Windows PE binaries. To ensure the executability and functionality of the generated samples, the format and constraints of PE files must remain intact.

To guarantee the quality of malware samples, we randomly select 1000 samples from VirusTotal that meet the following requirements: 1) More than 80% antivirus engines of VirusTotal label them malicious. 2) Top 5 antivirus engines label those samples as malicious. 3) The execution of those samples in a Cuckoo sandbox shows malicious behavior. Presence of malicious behaviour is ensured using the malicious score, which is computed by Cuckoo Sandbox as a weighted sum of identified signatures.

Some malware samples may detect a VM and not execute, while others may connect to a C&C servers that is not longer available. Such samples, while labeled malicious, do not display malicious behavior.

Visual Basic (VB) programs are excluded from the dataset because the IDA Pro in the implementation of code randomization [44] cannot generate CFG for them. These implementation issues are left as future work.

**Setup:** The experiments are performed on a desktop computer with Intel® Core™ i5-6400 CPU @ 2.70 GHz x 4, 16 GiB RAM, 1TB SSD, Ubuntu 18.04.3 LTS. All the scripts such as Binary Rewriter, Action Sequence Minimizer, Feature Interpreter, and result analysis are implemented in Python 3.6.9. The Binary Rewriter also requires pefile and IDA Pro 6.8. For all the antivirus software under testing, free versions and default settings are used. We choose four real-world commercial antivirus software for blackbox testing, which are anonymized as AV1, AV2, AV3, AV4. Each antivirus is installed on a guest machine (Windows 7 Version 6.1 Build 7601: Service Pack 1) in Virtualbox 5.2.32.

The network of the virtual machine is disabled for two reasons: 1) to ensure the malware will not infect other machines in the network and 2) to ensure the stability and reproducibility of our experiments. Network connected commercial AVs are continually updated based on the samples they scan, which can impact the evasion results over time.

We also choose two open-source models as baselines:

- **EMBER** [6] is an open-source machine-learning-based classifier that uses a tree-based classifier model LightGBM to detect malware. It generates a 2350-dimensional feature vector for each sample consisting of two main types of features: raw features (e.g. Byte-Histogram, ByteEntropyHistogram, Strings) and parsed

features (e.g. GeneralFileInfo, HeaderFileInfo, Section-Info, ImportsInfo, ExportsInfo). We use the implementation provided by Endgame in Machine Learning Static Evasion Competition (MLSEC) [41].

- **ClamAV** [11] (0.102.1) is an open-source signature-based antivirus engine that utilizes hash-based and rule-based signatures to detect malware. All the generated signatures for a file are available.

### 4.2 Adversarial Example Generation

#### 4.2.1 Evasion Rate

To test the effectiveness of adversarial examples of both a state-of-the-art classifier and antivirus software, we measure detection rate and evasion rate as follows.

**Detection Rate  $R_d$ :** measures the fraction of the original samples that a static malware scanner can detect. That is, detection rate  $R_d = N_d/N$ , where  $N_d$  is the number of the detected original samples and  $N$  is the total number of samples under testing.

**Evasion Rate  $R_e$ :** measures the fraction of verified adversarial malware that a static malware scanner failed to detect. That is, evasion Rate  $R_e = N_e/N_d$ , where  $N_e$  is the total number of verified adversarial malware that successfully evades a static malware scanner and  $N_d$  is the number of detected samples.

Table 3 shows the evasion ability of adversarial examples of different malware classifiers. From these results we can see that:

- For classifier EMBER, the evasion rate is the highest at 56.0%, which is much higher than the result of Anderson et al. [4] at only 10%-24%. When generating adversarial examples, their strategy aims at selecting correct actions at each step, while our strategy selects actions and content. Section 4.2.3 demonstrates the impact of selective content on the evasion rate for EMBER.
- ClamAV is much harder to evade, although its detection rate is also the lowest amongst the classifiers in this study. Our framework can only obtain a 17.2% evasion rate for the detected malware samples. However, this does not mean ClamAV is the most robust AV. On the contrary, with open source signatures, it is feasible to scan the target malware using a sliding window to find the byte sequences that match these signatures, and then to design customized actions to break these signatures. This would be an ad-hoc approach and not fair to other machine-learning-based models or closed-source AVs. We aim to provide a generic and automatic approach for all types of AV. Moreover, we find that for ClamAV, when the size of a rewritten binary exceeds 25 MB, the sample is always classified as benign. This case is not considered a successful evasion as it reveals not unreliable features but a design limitation.

Table 3: Evasion results show for each AV, the detection rate, the evasion rate, and the number of evasive and functional samples. The average time required to generate samples and the average number of actions required for each sample are shown.

	Detection Rate	# Evasive	# Functional	Evasion Rate	AVG Time	AVG Action Count (Before/After Mini)
<b>EMBER</b>	83.0%	486	467 (96.1%)	56.0%	21.9 s	5.33 / 2.06
<b>ClamAV</b>	78.4%	142	135 (95.1%)	17.2%	13.1 s	4.1 / 1.07
<b>AV1</b>	95.2%	279	274 (98.2%)	28.8%	22 s	5.14 / 1.92
<b>AV2</b>	100%	369	345 (93.5%)	34.5%	35.4 s	5.51 / 1.57
<b>AV3</b>	99.6%	436	417 (95.6%)	41.9%	50.7 s	6.45 / 2.44
<b>AV4</b>	93.6%	250	227 (90.8%)	24.3%	61 s	5.05 / 1.34

- For commercial antivirus systems, the evasion rates are in the range from 24.3% to 41.9%, which shows that the static scanning of current antivirus products is easy to evade using our framework.

#### 4.2.2 Functionality Verification

Some rewritten samples lose their functionality after manipulations. Examples of broken samples can be found in Appendix C.

Table 3 shows the number of evasive samples and functional samples. A sample may evade static scanning either because it has fooled the classifier or because it is no longer malicious due to the manipulations. As shown, even with carefully designed actions and minimizers that reduce redundant actions, it is still possible that evasive samples become nonfunctional. The functional rate ranges from 90.8% to 98.2%.

#### 4.2.3 Estimated Cost to the Adversary

We compute the number of queries and the time required for an adversary to generate a successful adversarial sample.

The evasion rate increases as the number of attempts increase (Figure 4). For ClamAV and commercial antivirus products, our framework reaches a relatively stable evasion rate after 30 attempts. But for EMBER, which is a machine-learning-based model, the evasion rate keeps increasing even after 50 attempts. We also noticed that the signature-based ClamAV has the lowest evasion rate, the ML-based EMBER has the highest evasion rate. The commercial AVs usually use both signature and machine learning models, and their evasion rates are between ClamAV and EMBER.

For commercial AVs, each attempt requires copying verified adversarial sample to a VM to check whether or not the sample is deleted by the AV within  $T$  seconds. For ClamAV and Ember, each attempt requires directly querying an ML model (EMBER) or a native Linux program (ClamAV), which is much quicker (Table 3). Sample generation is accelerated by placing multiple samples in the VM for evaluation in parallel.

As shown in Figure 5, the cost to the adversary is substantially reduced when action weights and successful content are




Figure 4: Evasion rate increases with the number of attempts but for most AVs, the evasion rate plateaus at about 30 attempts, indicating that additional manipulations are not fruitful. The evasion rate for EMBER, however, keeps increasing beyond 50% as the number of attempts is increased.

used to manipulate original samples as they result in more successful samples in fewer attempts. This figure only shows the result of EMBER because the improvement of EMBER with weight and content is the most outstanding. The improvements of others are under 3%.

#### 4.2.4 Action Sequence Length

According to our experiments, minimization halves the action sequence lengths for most samples (Table 3), resulting in majority of evasive samples requiring at most 3 actions.

This result demonstrates that using random search is sufficient for generating effective adversarial examples. Techniques such as reinforcement learning [4] do not provide additional improvement over random search.

#### 4.2.5 Number of bytes changed

To measure the difference between the adversarial example and the original malware, we compute the total number of bytes appended or modified to generate the adversarial example.




Figure 5: With weighted actions and selective content, more evasive samples can be generated with fewer attempts.

By positioning the samples in a line sorted by byte changes (Figure 6), we noticed that:

- By only changing one byte of the original malware, we can generate 90 evasive samples for ClamAV, 34 for AV1, 20 for AV2, 9 for AV3 and 11 for AV4.
- EMBER is not vulnerable to small changes. As a machine-learning-based model, it represents the target malware sample as a high-dimension feature vector (i.e. 2350 features), where tiny modifications of the vector are less likely to cause any significant impact in the feature space. Big changes that make the features appear similar to benign features (e.g. append a large amount of benign content at the end of the sample), can move the vector from malicious region to benign region in the feature space and cause a misclassification.

In Figure 7, we show 4 special cases comparing the binary content differences between the original malware samples and the generated evasive examples. In all cases, the difference is within 4 bytes. We can rename the section from “idata” to “rdata”, to cause an evasion, as shown in case (a). Or, we can remove the 4-byte debug information to cause another evasion, as shown in case (d). It indicates that the features used by static scanners are very fragile and superficial.

## 4.3 Explanation

### 4.3.1 Verification of Action-Feature Space Mapping using ClamAV

Using the decision rules of Figure 3, the Feature Interpreter extracts the essential feature changes for each macro-actions. Although the classification process of the black box AVs is unknown, the open-source ClamAV can be used for this validation.

For example, our framework reports the evasion reason for 13 samples is feature  $F_5$  (Section Padding) (see Table 2). For these 13 samples, the ClamAV all report them are caused by feature  $F_5$  too. This indicates that our explanations of all these 13 samples are correct. We then validated the remaining features reported by our framework, 21 samples for  $F_1$ , 62

samples for  $F_2$ , 1 sample for  $F_3$ , 2 samples for  $F_4$ , 17 samples for  $F_6$ , 11 samples for  $F_9$ , all of them are reported by ClamAV due to the breaking of the same reasons as the framework reports.

### 4.3.2 Explaining Why Evasion Happens

Understanding why evasion happens can help improve the robustness of a classifier against adversarial examples. In generating an adversarial sample for a target blackbox classifier, some of the actions can be removed without affecting the evasiveness of the sample, or can be replaced by micro-actions. To reveal the essential reason that causes the evasion, we use Action Sequence Minimizer and Feature Interpreter to find the minimum changes to evade the classifier.

We summarize the minimized actions used for all evasive samples among different classifiers, as shown in Figure 10 in Appendix B, and cross-reference it with the Action-Feature mapping in Table 2 to count the features changed in all samples, as shown in Figure 8.

1) For EMBER, more than 80% of evasive samples are generated using Overlay Append (OA). The OA action can affect multiple features, such as Byte Histogram (256 dimensions), Byte Entropy Histogram (256 dimensions), and Strings (104 dimensions). These features all belong to the data distribution feature category in our Action-Feature.

2) For ClamAV, the most effective actions are SP1, OA1, RD, SP, and CR. All of these actions impact hash based and rule-based signatures Table 2, which may exist in debug directory, special padding content or code sections.

3) For commercial AVs, the most effective action for AV1, AV3 and AV4, and also the second most effective action for AV2 is SA1. While this action corresponds to either breaking file hash or breaking the section count, the major cause for evasion for all AVs is section count, as seen in Figure 8. It is reassuring that commercial AVs do not use only file hash to differentiate malware and benign programs, but they still use the fragile feature section count.

As in Figure 8, AV1 and AV4 are not sensitive to data distribution change, but for AV2 and AV3, the data distribution change is the most important reason for the evasions of adversarial samples. This maybe because for AV2 and AV3, machine learning-based models are employed more heavily while AV1 and AV4 rely more on signature-based models. OA1 action can be used to generation 6 and 7 evasive samples for AV3 and AV4, respectively. This implies that these AVs have no mechanisms other than a file hash feature to identify these particular malware samples.

### 4.3.3 Transferability

Transferability refers to the property that allows the same adversarial examples to fool multiple models. If the adversarial examples are transferable, then evading one classifier also




Figure 6: Number of Changed Bytes of Adversarial Examples. By changing only one byte, adversarial examples are created for all tested AVs, with the exception of EMBER, which is not vulnerable to a few bytes of change in the binary.

```

root@ubuntu:~/hd# diff 52f46d37fff48f30bfa3b69ea070c46aebce4f0c1e14c21fc81658bd360fc678.hd*
32c32
< 00000210 00 00 00 20 00 00 60 2e 72 64 61 74 61 00 00 |...|.data..|
...
> 00000210 00 00 00 20 00 00 60 2e 69 64 61 74 61 00 00 |...|.data..|
(a) Section Rename

root@ubuntu:~/hd# diff 51568f63346fd261144ca3a34154852eab1bdbcba58570e11f47e812c189d17f.hd*
20c20
< 00000130 00 90 0c 00 00 10 00 00 da 28 00 ed 02 00 00 00 |.....|
...
> 00000130 00 90 0c 00 00 10 00 00 00 00 00 02 00 00 00 |.....|
(b) Break Checksum

root@ubuntu:~/hd# diff c1fcab4d6d19cc97f30e8cd0d8c7fa8ed6fe206f8160082a1b22594a86362067.hd*
39c39
< 00000410 00 00 00 33 c0 89 08 50 45 43 6f 6d 70 61 63 |...|.PECompac|
...
> 00000410 00 00 00 31 c0 89 08 50 45 43 6f 6d 70 61 63 |...|.PECompac|
(c) Code Randomization

root@ubuntu:~/hd# diff 62da3ab0ab3cba40826e0d6bccbd494681b184017302350a8e6e116b1d005d70.hd*
24,25d23
< 00000170 00 00 00 00 00 00 00 70 10 00 00 1c 00 00 00 |.....p.....|
< 00000180 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
(d) Remove Debug

```

Figure 7: Difference in binary content when various actions are applied.




Figure 8: Feature changes that cause evasion. EMBER is most impacted by data distribution changes, which also impact AV2 and AV3 heavily. Half of all ClamAV evasions changed the section hash, and section count changes impacted all commercial AVs.

imperils the other one.

Figure 9 shows the percentage of evasive samples generated for one classifier that can also evade other classifiers. From this table, we can see that 1) The evasive samples generated for AV3 have the highest transferability to other classifiers, especially to AV1 (16.55%). 2) The evasive samples generated

for AV4 have the lowest transferability to other AVs while the evasive samples generated for other AVs have a much higher chance to transfer to AV4 (i.e. 1.32% vs 8.76%, 7.05% vs 14.49%, 8.81% vs 13.19%). 3) None of the adversarial samples generated for ClamAV transfer to EMBER. Similarly, AV1 and AV4’s samples do not transfer to EMBER. It indicates that ClamAV, AV1 and AV4 do not rely (much) on machine learning models for static malware detection, or use machine learning models in a very different way. This is further confirmed by Figure 8, which shows that data distribution changes, which are associated with evading machine learning models, are not very effective at generating evasive samples for AV1 and AV4.

### 4.3.4 Testing Harm to Users

We want to test to what extent adversarial examples evade the full AV pipeline and infect users. We hypothesize that the dynamic and behavioral classifiers of the AVs will detect and stop the adversarial examples evading the the static-only classifiers when they are executed, thus, posing no real harm to the users. To answer this question, we create adversarial samples by modifying 30 ransomware samples and test whether the samples that evade static classifiers can infect users’ machines. We select the ransomware samples from our dataset by choosing samples whose Cuckoo sandbox signa-




Figure 9: Transferability of adversarial samples. 16% of samples that were generated using AV3 also evade AV1. AV3’s samples show the highest transferability to other AV, while AV4’s samples show the lowest. AV4 is also the most receptive to adversarial samples generated using other AVs.

ture reports contain “ransomware” or “writing encrypted files back to disk.” To test the adversarial ransomware, we place several decoy files (selfie.png and taxes.docs in the Desktop folder) in the VM and executed the adversarial ransomware. Explicit observable malicious behaviour, such as change in the content of the decoy files, or notifications from the ransom indicates the successful execution of ransomware.

For AV1, AV3, AV4, blocked the execution of adversarial ransomware samples – none of the samples can encrypt the user’s personal files. However, AV2 is the exception. All of the 30 adversarial ransomware samples evade the behavior detector of AV2; files are encrypted and blackmail messages are shown on the screen. AV2 shows alert for 4 samples but does not stop the encryption for any of the 30 samples.

We then connect to the Internet and redo the experiment for all the antivirus. We can still evade the static scanner of AV1. AV2 can detect all the samples this time. We have contacted the R&D team of AV2 regarding our findings. They stated their products rely heavily on cloud techniques since they think malware authors can easily reverse engineer the AV in offline mode. We think devices should be protected even when offline. An attacker might apply other techniques to disconnect the Internet or use a USB to load the malware to the victim machine and execute the sample. Without offline protection, the malware sample might infect a machine and cause damage to it.

This represents a new attack surface for adversarial examples and a recommendation for future antivirus systems. Anti-virus systems rely on both static and dynamic detection to detect malware. However, since static detectors are easy to evade, AVs need a robust dynamic detection to protect users from malicious programs. Three of the four antivirus systems we tested already provide this protection, however,

one system did not have a robust offline dynamic detection to detect the adversarial malware when they are executed.

## 5 Discussions

**Triviality of Defense.** The triviality of the defense depends on the type of attack. To defend the overlay append attack, the defender can ignore the overlay data when training models. To defend the SA attack, the defender can lower the importance of benign features in models, and only consider malware features. To defend the RD, RS, BC attack, the defender should avoid using such fragile patterns as malware features. The code randomization (CR) attack is hard to defend because the defender cannot locate the small striped of binary that is randomized.

**Recommendation for Antivirus Systems.** Our attacks demonstrate that some antivirus systems still rely on fragile signature-based detection for some type of malware. This practice can make them vulnerable to trivial adversarial attacks. We recommend redesigning the malware features, reducing fragile features, incorporating dynamic features of malware. It will provide better protection to the compute when static scanners fail. Another important point to note is the reliance on static-only classifiers. Several prior work show high accuracy for a static-only malware detector that might inspire commercial systems to heavily rely on static classifiers. However, we demonstrate that static classifiers are easy to evade, thus dynamic classification is crucial to protect users from harm.

**Recommendation for Researchers.** We demonstrate how adversarial examples can be used to explain how a complex blackbox system makes decisions. When training malware classifiers, the researchers should use explanation techniques to understand the behaviors of the classifiers and check if the learned features are fragile features that can be easily evaded or conflict with expert knowledge. We also argue that for security applications, demonstrating harm to real users is crucial to understand the real ramification of an attack.

**Fully automatic generation of adversarial malware.** In some cases, automatically generating an adversarial malware against a real-world system might be difficult. When an antivirus system detects a malware, it usually actively delete it or disinfect it in static scanning. However, sometimes it may take a long time to react to certain malware, or some antivirus products do not delete samples without user interactions. This will block the generation of adversarial examples as well as reduce the throughout of the framework. To mitigate this, we can either wait a long time for the antivirus system to take action or write Marco scripts to simulate user delete actions.

## 6 Related Work

Adversarial attacks on machine learning is a rapidly growing field. Since 2014, there has been more than 1200 papers on adversarial attacks and defense<sup>3</sup>. However, only 35 papers focused on the malware domain, the rest focuses on the image domain. These works performed attacks on Android malware [16, 25, 29, 32, 34, 39, 43, 47, 60], PDF malware [10, 17, 38, 61], Windows malware (PE files) [2, 4, 9, 15, 20, 22–24, 26, 27, 31, 33, 35, 45, 52, 53, 57, 58], IoT malware [1] and Flash-based malware [37]. The recent papers performed both whitebox attacks where the adversary has complete control of the machine learning model and black box attacks where the adversary has limited control. The black box attacks usually use a surrogate model to construct the attacks, which is then transferred to another model. Because of the transferability of machine learning models [36, 42], an adversarial sample generated from one machine learning model can often fool another model with different structure and parameters. Attackers can train a local surrogate model to substitute for the target model, allowing them to bypass the black-box constraint and generate adversarial samples in a white-box setting. However, though this approach is successful in the image domain, we demonstrate that in the malware domain, adversarial examples that evade an open-source surrogate model are unlikely to evade commercial antivirus systems.

Prior work proposed various ways of constructing adversarial attacks against a variety of machine-learning-based systems. Xu et al. [61] used genetic algorithms to generate evasive PDF malware samples. Dang et al. [13] proposed a hill-climbing-based evasion method, EvadeHC, to generate PDF files that can evade PDF malware classifiers. The goal of their approach is to find a sample that has crossed the boundary of the detector but not the boundary of the tester. Hu et al. [21] used Generative Adversarial Networks (GAN) to generate adversarial examples against a black-box model. Their approach, MalGAN, consists of a generator that transforms a malware feature vector into its adversarial version by only adding a random number of API calls. Anderson et al. [5] used deep reinforcement learning to generate adversarial examples against a gradient boosted decision tree model. The deep reinforcement learning model consists of an agent and an environment. However, their preliminary RL model provided only 1% improvement of the evasion rate over random selection. Kolosnjaji et al. [28] proposed to conduct white-box gradient-based adversarial attacks on deep learning model MalConv [49], which directly takes the raw bytes of the binaries as input. The paper observes that to achieve a high evasion rate an attacker needs to append a significant amount of garbage data at the end of a malware sample, which makes these attacks easy to detect as they cannot manipulate bytes at arbitrary locations. By generate adversarial samples

by changing few tens of bytes in the file header, Luca et al. [15] demonstrated that MalConv differentiates benign and malicious samples mostly based on characteristics of the file header and ignores the data and code section.

A few papers discussed attacks against a dynamic classifier [56], although all the attacks are performed in the feature space. We are the first to demonstrate functionally active malware can evade real-world dynamic classifiers and harm real users.

Demetrio [15] leverages integrated gradients to find the influential features for black-box decisions of an ML-based malware classifier MalConv. Their results show that MalConv only learns patterns from the file header part of binaries, not important data and code sections. However, the paper manipulated the byte-level features without manipulating the actual file. LIME [50] is an explanatory model that distorts the inputs by blacking out random regions and then checking its impact on outputs to identify the influential regions to classification. LEMNA [19] is a modified version of LIME customized for a security domain, which learns the interpretation with the sequential RNN (Recurrent Neural Network) model. The manipulations on samples in both LIME and LEMNA do not guarantee the sample is still executable and functionality-preserving when applied in malware adversarial domain. However, all these three approaches are conducted on the byte level, not the feature level. A one-byte change that alters the classification result may be because it breaks the format of PE and thus its functionality, not because it is important to the identification of malware. So their explanation method cannot be used for commercial AVs. Our approach generates functional variants around the adversarial samples using our safe action set, and provides a mapping from action space to feature space. In this way, we can explain which features are essential for the evasion of adversarial examples.

## 7 Conclusions

We design a generalized, systematic, automatic framework to perform realistic adversarial attacks on real-world malware detectors. We use in-place code randomization and code-agnostic manipulations to generate adversarial examples. To explain why evasion occurs, we design an algorithm to filter out the actions that are ineffective for adversarial sample generation. For each commercial and open-source antivirus system, we compute the cost to the attacker and the effectiveness of the attack on static and dynamic classifiers, the effectiveness of each action and the key features that cause evasions. Our results show that it is easy to evade the static detectors of commercial AV, and that many of the adversarial attacks are transferable between different antivirus systems. Additionally, when the static classifier of antivirus systems are evaded, these adversarial samples can be used to test the dynamic classifier, which is a new attack surface for adversarial examples to do real harm to users.

<sup>3</sup><https://nicholas.carlini.com/writing/2019/all-adversarial-example-papers.html>

## A Algorithm of Action Sequence Minimizer

See Algorithm 3

---

**Algorithm 3:** Action Sequence Minimization

---

**Input** : malware sample  $s$ , evasive sample  $s_{eva}$ , malware classifier  $f$   
**Output** : minimized sample  $s_{min}$

```
1  $q \leftarrow \text{getActionSeq}(s_{eva})$ 
2  $action\_idxs \leftarrow [0, \dots, \text{len}(q) - 1]$ 
3  $deleted\_idxs \leftarrow []$ 
4 foreach  $idx \in action\_idxs$  do
5    $q' \leftarrow q$ 
6    $q'[idx] \leftarrow \text{None}$ 
7   foreach  $idx' \in deleted\_idxs$  do
8      $q'[idx'] \leftarrow \text{None}$ 
9   end
10   $s' \leftarrow s$ 
11  foreach  $action \in q'$  do
12    if  $action \neq \text{None}$  then
13       $content \leftarrow \text{parseLog}(action.idx)$ 
14       $s' \leftarrow \text{applyAction}(s', action, content)$ 
15    end
16  end
17  if  $isEvasive(f, s')$  then
18     $deleted\_idxs.append(idx)$ 
19     $s_{min} \leftarrow s'$ 
20  else
21     $action \leftarrow q[idx]$ 
22     $\text{updateActionWeight}(f, action)$ 
23     $content \leftarrow \text{parseLog}(action.idx)$ 
24     $\text{saveSuccContent}(content)$ 
25  end
26 end
27 return  $s_{min}$ 
```

---

## B Effective Action Combinations

See Figure 10.

## C Broken Malware Examples

Case 1: Implementation errors in instruction replacement. As shown in Figure 11, the original implementation of code randomization only supports 8-bit and 32-bit, not 16-bit instruction. It tries to replace a 16-bit add/sub instruction in a wrong way (assuming 32-bit format). It treats the last four bytes as the second operand, but actually, only the last two bytes are the second operand. This would break the CFG of the program.

Case 2: Overwriting overlay. As shown in Figure 12, when adding a new section at the end of the last section, if the

sample has overlay data, the added new data may affect the overlay data extraction of the malware.

## References

- [1] Ahmed Abusnaina, Aminollah Khormali, Hisham Alasmary, Jeman Park, Afsah Anwar, Ulku Meteriz, and Aziz Mohaisen. Examining adversarial learning against graph-based iot malware detection systems. *arXiv preprint arXiv:1902.04416*, 2019.
- [2] Abdullah Al-Dujaili, Alex Huang, Erik Hemberg, and Una-May O'Reilly. Adversarial deep learning for robust detection of binary encoded malware. In *2018 IEEE Security and Privacy Workshops (SPW)*, pages 76–82. IEEE, 2018.
- [3] Abdullah Al-Dujaili, Alex Huang, Erik Hemberg, and Una-May O'Reilly. Adversarial deep learning for robust detection of binary encoded malware. In *2018 IEEE Security and Privacy Workshops (SPW)*, pages 76–82. IEEE, 2018.
- [4] Hyrum S Anderson, Anant Kharkar, Bobby Filar, David Evans, and Phil Roth. Learning to evade static pe machine learning malware models via reinforcement learning. *arXiv preprint arXiv:1801.08917*, 2018.
- [5] Hyrum S Anderson, Anant Kharkar, Bobby Filar, and Phil Roth. Evading machine learning malware detection. *Black Hat*, 2017.
- [6] Hyrum S Anderson and Phil Roth. Ember: an open dataset for training static pe malware machine learning models. *arXiv preprint arXiv:1804.04637*, 2018.
- [7] Ai & machine learning. <https://www.avast.com/en-us/technology/ai-and-machine-learning>.
- [8] Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian Goodfellow, and Aleksander Madry. On evaluating adversarial robustness. *arXiv preprint arXiv:1902.06705*, 2019.
- [9] Li Chen. Understanding the efficacy, reliability and resiliency of computer vision techniques for malware detection and future research directions. *arXiv preprint arXiv:1904.10504*, 2019.
- [10] Lingwei Chen, Yanfang Ye, and Thirimachos Bourlai. Adversarial machine learning in malware detection: Arms race between evasion attack and defense. In *2017 European Intelligence and Security Informatics Conference (EISIC)*, pages 99–106. IEEE, 2017.

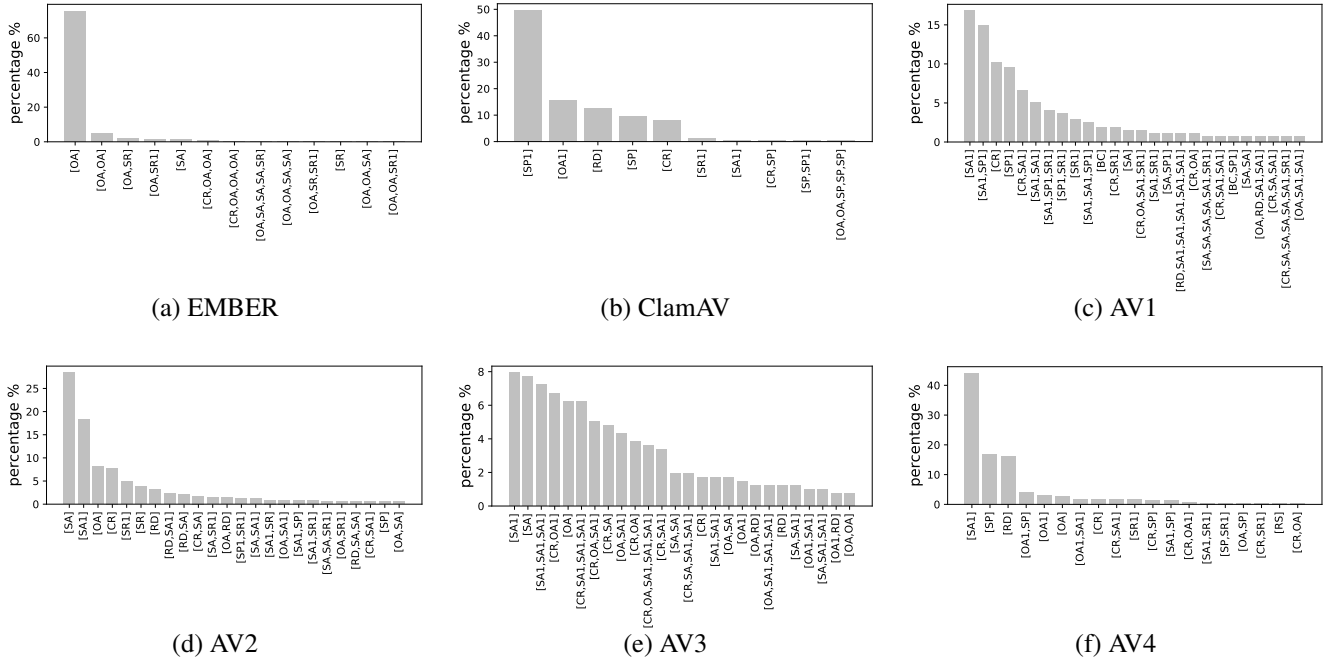


Figure 10: Effective Action Combinations for Adversarial Sample Generation: For EMBER, addition of random bytes to the end of the binary is overwhelmingly the most effect action for fooling the classifier. Addition of just one byte to the end of a section or to the end of the file can create adversarial examples for ClamAV. And, addition of a new section is the most effective action for creating adversarial samples for all the four commercial AVs. For AV1, AV3 and AV4, addition of a section with just one byte can create adversarial samples.

```

mov     ecx, dword [ebp+0x30]
mov     ecx, word [ebp+0x32]
xor     ecx, 0x2a06
mov     dword [esp+0x42] (var_3c), ecx
add     ecx, 0xffffffff (0x0)
mov     dword [esp+0x18] (var_58), ecx
mov     dword [esp+0x14] (var_4c), ecx (eax)
mov     dword [esp+0x10] (var_40), ecx (eax)
xor     ecx, 0x2a06
add     ecx, 0x2a06 => sub ecx, -0x2a06

```

Figure 11: Implementation Errors in CR

- [11] Clamav. <https://www.clamav.net/>.
- [12] George E Dahl, Jack W Stokes, Li Deng, and Dong Yu. Large-scale malware classification using random projections and neural networks. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 3422–3426. IEEE, 2013.
- [13] Hung Dang, Yue Huang, and Ee-Chien Chang. Evading classifiers by morphing in the dark. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 119–133. ACM, 2017.
- [14] Park Daniel, Khan Haidar, and Yener Bülent. Generation & evaluation of adversarial examples for malware obfuscation. *arXiv preprint arXiv:1904.04802*, 2019.




Figure 12: Errors after Overlay Data Append

- [15] Luca Demetrio, Battista Biggio, Giovanni Lagorio, Fabio Roli, and Alessandro Armando. Explaining vulnerabilities of deep learning to adversarial malware binaries. *arXiv preprint arXiv:1901.03583*, 2019.
- [16] Ambra Demontis, Marco Melis, Battista Biggio, Davide Maiorca, Daniel Arp, Konrad Rieck, Iginio Corona, Giorgio Giacinto, and Fabio Roli. Yes, machine learning can be more secure! a case study on android malware de-



- tection. *IEEE Transactions on Dependable and Secure Computing*, 2017.
- [17] Saeed Ehteshamifar, Antonio Barresi, Thomas R Gross, and Michael Pradel. Easy to fool? testing the anti-evasion capabilities of pdf malware scanners. *arXiv preprint arXiv:1901.05674*, 2019.
- [18] William Fleshman, Edward Raff, Richard Zak, Mark McLean, and Charles Nicholas. Static malware detection & subterfuge: Quantifying the robustness of machine learning and current anti-virus. In *2018 13th International Conference on Malicious and Unwanted Software (MALWARE)*, pages 1–10. IEEE, 2018.
- [19] Wenbo Guo, Dongliang Mu, Jun Xu, Purui Su, Gang Wang, and Xinyu Xing. Lemna: Explaining deep learning based security applications. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 364–379. ACM, 2018.
- [20] Weiwei Hu and Ying Tan. Generating adversarial malware examples for black-box attacks based on gan. *arXiv preprint arXiv:1702.05983*, 2017.
- [21] Weiwei Hu and Ying Tan. Generating adversarial malware examples for black-box attacks based on gan. *arXiv preprint arXiv:1702.05983*, 2017.
- [22] Weiwei Hu and Ying Tan. Black-box attacks against rnn based malware detection algorithms. In *Workshops at the Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [23] Alex Huang, Abdullah Al-Dujaili, Erik Hemberg, and Una-May O’Reilly. On visual hallmarks of robustness to adversarial malware. *arXiv preprint arXiv:1805.03553*, 2018.
- [24] Yonghong Huang, Utkarsh Verma, Celeste Fralick, Gabriel Infante-Lopez, Brajesh Kumar, and Carl Woodward. Malware evasion attack and defense. *2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)*, Jun 2019.
- [25] ElMouatez Billah Karbab, Mourad Debbabi, Abdelouahid Derhab, and Djedjiga Mouheb. Android malware detection using deep learning on api method sequences. *arXiv preprint arXiv:1712.08996*, 2017.
- [26] Aminollah Khormali, Ahmed Abusnaina, Songqing Chen, DaeHun Nyang, and Aziz Mohaisen. Copycat: Practical adversarial attacks on visualization-based malware detection. *arXiv preprint arXiv:1909.09735*, 2019.
- [27] Bojan Kolosnjaji, Ambra Demontis, Battista Biggio, Davide Maiorca, Giorgio Giacinto, Claudia Eckert, and Fabio Roli. Adversarial malware binaries: Evading deep learning for malware detection in executables. In *2018 26th European Signal Processing Conference (EU-SIPCO)*, pages 533–537. IEEE, 2018.
- [28] Bojan Kolosnjaji, Ambra Demontis, Battista Biggio, Davide Maiorca, Giorgio Giacinto, Claudia Eckert, and Fabio Roli. Adversarial malware binaries: Evading deep learning for malware detection in executables. In *2018 26th European Signal Processing Conference (EU-SIPCO)*, pages 533–537. IEEE, 2018.
- [29] Alex Kouzemtchenko. Defending malware classification networks against adversarial perturbations with non-negative weight restrictions. *arXiv preprint arXiv:1806.09035*, 2018.
- [30] Felix Kreuk, Assi Barak, Shir Aviv-Reuven, Moran Baruch, Benny Pinkas, and Joseph Keshet. Adversarial examples on discrete sequences for beating whole-binary malware detection. *arXiv preprint arXiv:1802.04528*, 2018.
- [31] Felix Kreuk, Assi Barak, Shir Aviv-Reuven, Moran Baruch, Benny Pinkas, and Joseph Keshet. Deceiving end-to-end deep learning malware detectors using adversarial examples. *arXiv preprint arXiv:1802.04528*, 2018.
- [32] Deqiang Li, Ramesh Baral, Tao Li, Han Wang, Qianmu Li, and Shouhuai Xu. Hashtran-dnn: A framework for enhancing robustness of deep neural networks against adversarial malware samples. *arXiv preprint arXiv:1809.06498*, 2018.
- [33] Deqiang Li, Qianmu Li, Yanfang Ye, and Shouhuai Xu. Enhancing robustness of deep neural networks against adversarial malware samples: Principles, framework, and aics’2019 challenge. *arXiv preprint arXiv:1812.08108*, 2018.
- [34] Xiaolei Liu, Xiaojiang Du, Xiaosong Zhang, Qingxin Zhu, Hao Wang, and Mohsen Guizani. Adversarial samples on android malware detection systems for iot systems. *Sensors*, 19(4):974, 2019.
- [35] Xinbo Liu, Jiliang Zhang, Yaping Lin, and He Li. Atmpa: Attacking machine learning-based malware visualization detection methods via adversarial examples. 2019.
- [36] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into transferable adversarial examples and black-box attacks. *arXiv preprint arXiv:1611.02770*, 2016.

- [37] Davide Maiorca, Battista Biggio, Maria Elena Chiappe, and Giorgio Giacinto. Adversarial detection of flash malware: Limitations and open issues. *CoRR*, abs/1710.10225, 2017.
- [38] Davide Maiorca, Battista Biggio, and Giorgio Giacinto. Towards robust detection of adversarial infection vectors: Lessons learned in pdf malware. *arXiv preprint arXiv:1811.00830*, 2018.
- [39] Marco Melis, Davide Maiorca, Battista Biggio, Giorgio Giacinto, and Fabio Roli. Explaining black-box android malware detection. In *2018 26th European Signal Processing Conference (EUSIPCO)*, pages 524–528. IEEE, 2018.
- [40] New machine learning model sifts through the good to unearth the bad in evasive malware. <https://www.microsoft.com/security/blog/2019/07/25/new-machine-learning-model-sifts-through-the-good-to-unearth-the-bad-in-evasive-malware/>.
- [41] MLSEC. [https://github.com/endgameinc/malware\\_evasion\\_competition](https://github.com/endgameinc/malware_evasion_competition).
- [42] Nicolas Papernot, Patrick McDaniel, and Ian Goodfellow. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *arXiv preprint arXiv:1605.07277*, 2016.
- [43] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia conference on computer and communications security*, pages 506–519. ACM, 2017.
- [44] Vasilis Pappas, Michalis Polychronakis, and Angelos D Keromytis. Smashing the gadgets: Hindering return-oriented programming using in-place code randomization. In *2012 IEEE Symposium on Security and Privacy*, pages 601–615. IEEE, 2012.
- [45] Daniel Park, Haidar Khan, and Bülent Yener. Short paper: Creating adversarial malware examples using code insertion. *CoRR*, abs/1904.04802, 2019.
- [46] pcmag. <https://www.pcmag.com/picks/the-best-antivirus-protection>.
- [47] Robert Podschwadt and Hassan Takabi. Effectiveness of adversarial examples and defenses for malware classification. *arXiv preprint arXiv:1909.04778*, 2019.
- [48] Edward Raff, Jon Barker, Jared Sylvester, Robert Brandon, Bryan Catanzaro, and Charles K Nicholas. Malware detection by eating a whole exe. In *Workshops at the Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [49] Edward Raff, Jon Barker, Jared Sylvester, Robert Brandon, Bryan Catanzaro, and Charles K Nicholas. Malware detection by eating a whole exe. In *Workshops at the Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [50] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Why should i trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144. ACM, 2016.
- [51] Konrad Rieck, Philipp Trinius, Carsten Willems, and Thorsten Holz. Automatic analysis of malware behavior using machine learning. *Journal of Computer Security*, 19(4):639–668, 2011.
- [52] Ishai Rosenberg, Asaf Shabtai, Yuval Elovici, and Lior Rokach. Query-efficient gan based black-box attack against sequence based machine and deep learning classifiers. *arXiv preprint arXiv:1804.08778*, 2018.
- [53] Ishai Rosenberg, Asaf Shabtai, Lior Rokach, and Yuval Elovici. Generic black-box end-to-end attack against state of the art api call based malware classifiers. In *International Symposium on Research in Attacks, Intrusions, and Defenses*, pages 490–510. Springer, 2018.
- [54] Joshua Saxe and Konstantin Berlin. Deep neural network based malware detection using two dimensional binary program features. In *2015 10th International Conference on Malicious and Unwanted Software (MALWARE)*, pages 11–20. IEEE, 2015.
- [55] Matthew G Schultz, Eleazar Eskin, F Zadok, and Salvatore J Stolfo. Data mining methods for detection of new malicious executables. In *Proceedings 2001 IEEE Symposium on Security and Privacy. S&P 2001*, pages 38–49. IEEE, 2000.
- [56] Jack W Stokes, De Wang, Mady Marinescu, Marc Marino, and Brian Bussone. Attack and defense of dynamic analysis-based, adversarial neural malware classification models. *arXiv preprint arXiv:1712.05919*, 2017.
- [57] Jack W Stokes, De Wang, Mady Marinescu, Marc Marino, and Brian Bussone. Attack and defense of dynamic analysis-based, adversarial neural malware detection models. In *MILCOM 2018-2018 IEEE Military Communications Conference (MILCOM)*, pages 1–8. IEEE, 2018.
- [58] Octavian Suciuc, Scott E Coull, and Jeffrey Johns. Exploring adversarial examples in malware detection. In *2019 IEEE Security and Privacy Workshops (SPW)*, pages 8–14. IEEE, 2019.

- [59] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3319–3328. JMLR. org, 2017.
- [60] Rahim Taheri, Reza Javidan, Mohammad Shojafar, Zahra Pooranian, Ali Miri, and Mauro Conti. On defending against label flipping attacks on malware detection systems. *arXiv preprint arXiv:1908.04473*, 2019.
- [61] Weilin Xu, Yanjun Qi, and David Evans. Automatically evading classifiers. In *Proceedings of the 2016 network and distributed systems symposium*, pages 21–24, 2016.