

# Intriguing Properties of Adversarial ML Attacks in the Problem Space [Extended Version]

JACOPO CORTELLAZZI, King's College London, UK  
FEARGUS PENDLEBURY, University College London, UK  
DANIEL ARP, Technische Universität Wien, Austria  
ERWIN QUIRING, Ruhr University Bochum and ICSI, Germany  
FABIO PIERAZZI, University College London, UK  
LORENZO CAVALLARO, University College London, UK

Recent research efforts on adversarial machine learning (ML) have investigated problem-space attacks, focusing on the generation of real evasive objects in domains where, unlike images, there is no clear inverse mapping to the feature space (e.g., software). However, the design, comparison, and real-world implications of problem-space attacks remain underexplored.

This article makes three major contributions. Firstly, we propose a general formalization for adversarial ML evasion attacks in the problem-space, which includes the definition of a comprehensive set of constraints on available transformations, preserved semantics, absent artifacts, and plausibility. We shed light on the relationship between feature space and problem space, and we introduce the concept of *side-effect features* as the by-product of the inverse feature-mapping problem. This enables us to define and prove necessary and sufficient conditions for the existence of problem-space attacks. Secondly, building on our general formalization, we propose a novel problem-space attack on Android malware that overcomes past limitations in terms of semantics and artifacts. We have tested our approach on a dataset with 150K Android apps from 2016 and 2018 which show the practical feasibility of evading a state-of-the-art malware classifier along with its hardened version. Thirdly, we explore the effectiveness of adversarial training as a possible approach to enforce robustness against adversarial samples, evaluating its effectiveness on the considered machine learning models under different scenarios.

Our results demonstrate that “adversarial-malware as a service” is a realistic threat, as we automatically generate thousands of realistic and inconspicuous adversarial applications at scale, where on average it takes only a few minutes to generate an adversarial instance.

CCS Concepts: • **Security and privacy** → **Malware and its mitigation**; **Software security engineering**; • **Computing methodologies** → **Artificial intelligence**.

Additional Key Words and Phrases: machine learning, malware, adversarial training, AI security

## ACM Reference Format:

Jacopo Cortellazzi, Feargus Pendlebury, Daniel Arp, Erwin Quiring, Fabio Pierazzi, and Lorenzo Cavallaro. 2025. Intriguing Properties of Adversarial ML Attacks in the Problem Space [Extended Version]. *ACM Trans. Priv. Sec.* 1, 1 (June 2025), 36 pages. <https://doi.org/TOPS-2024-04-0123.R2>

---

Authors' addresses: Jacopo Cortellazzi, [jacopo.cortellazzi@kcl.ac.uk](mailto:jacopo.cortellazzi@kcl.ac.uk), King's College London, UK; Feargus Pendlebury, [f.pendlebury@ucl.ac.uk](mailto:f.pendlebury@ucl.ac.uk), University College London, UK; Daniel Arp, [daniel.arp@tuwien.ac.at](mailto:daniel.arp@tuwien.ac.at), Technische Universität Wien, Austria; Erwin Quiring, [erwin.quiring@rub.de](mailto:erwin.quiring@rub.de), Ruhr University Bochum and ICSI, Germany; Fabio Pierazzi, [f.pierazzi@ucl.ac.uk](mailto:f.pierazzi@ucl.ac.uk), University College London, UK; Lorenzo Cavallaro, [l.cavallaro@ucl.ac.uk](mailto:l.cavallaro@ucl.ac.uk), University College London, UK.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2025 Association for Computing Machinery.

2471-2566/2025/6-ART \$15.00

<https://doi.org/TOPS-2024-04-0123.R2>

## 1 INTRODUCTION

Adversarial machine learning (ML) attacks have been extensively studied across various domains [49], and they present a significant threat to the widespread adoption of machine learning solutions in security-critical scenarios. This paper focuses on test-time evasion attacks in the so-called *problem space*, wherein the challenge lies in modifying real input-space objects to correspond to an adversarial feature vector.

The main challenge lies in the *inverse feature-mapping problem* [66, 20, 70, 13, 18, 17], as in many settings, converting a feature vector into a problem-space object is not feasible due to the non-invertible and non-differentiable nature of the feature-mapping function. Additionally, the modified problem-space object must be a valid and inconspicuous member of the considered domain while remaining robust against non-ML detection attempts, such as program analysis. Prior research has explored problem-space attacks in diverse areas, including text [47, 65], malicious PDFs [20, 16, 18, 35, 15, 24], Android malware [39, 57, 46], Windows malware [53], source code attribution [70], malicious Javascript [63], and eyeglass frames [34]. However, while there is a good understanding of how to perform feature-space attacks [36], the requirements for attacks in the problem space—along with principled comparisons of existing solutions’ strengths and weaknesses—remain less clear.

We introduce the concept of *side-effect features* as the by-product of attempting to generate a problem-space transformation that perturbs the feature space in a specific direction. Defining side-effect features allows us to gain deeper insights into the relationships between the feature space and problem space. We establish necessary and sufficient conditions for the existence of problem-space attacks and identify two main types of search strategies (feature-driven and problem-driven) for generating adversarial objects in the problem space. Additionally, we analyze how these samples can enhance the robustness of a target model through adversarial training, considering scenarios where the attacker and defender capabilities depend on gadget characteristics. We explore how representing the same code through multiple representations may impact the potential robustness built during adversarial training. Furthermore, we employ our formalization to describe various interesting attacks proposed in both problem space and feature space. This analysis reveals that some promising prior problem-space attacks in the Android malware domain [57, 46, 41] have limitations, particularly concerning semantics and artifacts. The limitations of the mentioned works inspire us to propose a novel problem-space attack in the Android malware domain, addressing these issues and overcoming existing solutions’ shortcomings.

At last, we evaluate adversarial training as a potential defense mechanism against problem space attacks and consider the impact and the differences on using multiple generative strategies for hardening. This approach is widely accepted as the most effective method in practice to improve the adversarial robustness of traditional machine learning and deep learning models. The fundamental premise of adversarial training, introduced by [25] [23], is to train a model to withstand adversarial attacks by exposing it to carefully crafted adversarial examples, which are designed to perturb the model’s predictions. Investigating adversarial training methods is crucial in the current landscape of AI and machine learning, where models are increasingly deployed in safety-critical applications such as autonomous driving [76] and medical diagnostics [64]. By understanding and improving upon adversarial training techniques, researchers can enhance the security and reliability of AI systems, ultimately contributing to the responsible deployment of artificial intelligence.

In summary, this paper presents three major contributions:

- We propose a novel general formalization of problem-space attacks (§2), which defines key requirements and commonalities across different domains, establishes necessary and sufficient conditions for problem-space attacks, and facilitates the comparison of strengths and weaknesses of prior approaches. Notably, existing strategies for adversarial malware

generation are found to be among the weakest in terms of attack robustness. We introduce the concept of *side-effect features*, revealing connections between feature space and problem space, enabling principled reasoning about search strategies for problem-space attacks.

- Building on our formalization, we propose a novel problem-space attack in the Android malware domain, utilizing automated software transplantations [26], which addresses limitations of prior work regarding semantics and artifacts. Through experiments on a dataset of 150k apps from 2016-2018 (§4), we demonstrate the feasibility of evading a state-of-the-art malware classifier, DREBIN [21], and its hardened version, Sec-SVM [39]. Compared to the conference version, we have enhanced the gadgets dataset, being able to extract the whole set of DREBIN features. The time required to generate an adversarial example is in the order of minutes, highlighting the realistic threat of the "adversarial-malware as a service" scenario and the insufficiency of existing defenses.
- We further expand our reasoning from the conference version and propose a comprehensive study of whether different types of adversarial hardening techniques can serve as a valuable defense against this type of adversarial attacks. We consider both adversarial retraining [79] and adversarial training [43]. We show that adversarial training can improve model robustness but the impact of it heavily depends on the considered model.

## 2 PROBLEM-SPACE ADVERSARIAL ML ATTACKS

Our research is centered on *evasion attacks* [18, 36, 13], where adversaries manipulate objects during the testing phase to produce specific misclassifications. We draw upon previous literature to give insights into *feature-space* attacks (see §2.1), and subsequently present a fresh perspective on *problem-space* attacks (§2.2). Subsequently, we shed light on the key aspects of our framework, illustrating its application to both the well-established feature-space and the newer problem-space attacks, as informed by related studies across multiple domains (§2.3). Our approach to threat modeling, rooted in understanding the attacker's knowledge and abilities, aligns with existing studies [49, 58, 62]. A comprehensive overview can be found in Appendix B. For reader convenience, we have cataloged symbols in Appendix A.

### 2.1 Feature-Space Attacks

It's worth noting that the various definitions of feature-space attacks (refer to §2.1) have been previously summarized in existing literature [49, 36, 39, 41, 25, 14, 5, 6]. We present them here for clarity and to set the stage for examining the connections between feature-space and problem-space attacks in subsequent sections.

We introduce a *problem space*  $\mathcal{Z}$ , also known as the *input space*, encompassing entities of a specific domain, such as images [36], sound recordings [50], software [70], or PDF documents [16]. Every entity  $z \in \mathcal{Z}$  is paired with a corresponding authentic label  $y \in \mathcal{Y}$ , with  $\mathcal{Y}$  symbolizing the realm of potential labels. As most machine learning approaches predominantly operate on vectorized numerical data [7], entities in  $\mathcal{Z}$  must undergo a transformation to be suitable for machine learning procedures.

**DEFINITION 1 (FEATURE MAPPING).** A *feature mapping* denotes a function  $\varphi : \mathcal{Z} \rightarrow \mathcal{X} \subseteq \mathbb{R}^n$  that transforms a problem-space entity  $z \in \mathcal{Z}$  into an  $n$ -dimensional feature vector  $\mathbf{x} \in \mathcal{X}$ , resulting in  $\varphi(z) = \mathbf{x}$ . This encapsulates implicit/latent mappings as well, where features are not immediately discernible upon input but are internally computed by the model, for instance, in deep learning frameworks [31].

**DEFINITION 2 (DISCRIMINANT FUNCTION).** For an  $m$ -class machine learning classifier represented as  $g : \mathcal{X} \rightarrow \mathcal{Y}$ , a discriminant function  $h : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$  yields a real value  $h(\mathbf{x}, i)$ , often expressed

succinctly as  $h_i(\mathbf{x})$ . This value characterizes the suitability of object  $\mathbf{x}$  for class  $i \in \mathcal{Y}$ . Elevated returns from the discriminant function  $h_i$  indicate a stronger association with class  $i$ . Specifically, the anticipated label for an entity  $\mathbf{x}$  is determined as  $g(\mathbf{x}) = \hat{y} = \arg \max_{i \in \mathcal{Y}} h_i(\mathbf{x})$ .

The goal of a *targeted* feature-space attack is to alter an entity  $\mathbf{x} \in \mathcal{X}$  assigned with the label  $y \in \mathcal{Y}$  into another entity  $\mathbf{x}'$  that correlates with a designated class  $t \in \mathcal{Y}$ , where  $t \neq y$  (meaning, to adapt  $\mathbf{x}$  to be wrongly classified under target class  $t$ ). Attackers can discern a disturbance  $\delta$  to tweak  $\mathbf{x}$  such that  $g(\mathbf{x} + \delta) = t$  by refining a meticulously designed *attack objective function*. For a comprehensive definition of this function, readers are directed to [36] and [49], where *high-confidence* attacks and multi-class scenarios are taken into consideration.

**DEFINITION 3 (ATTACK OBJECTIVE FUNCTION).** *Given an object  $\mathbf{x} \in \mathcal{X}$  and a target label  $t \in \mathcal{Y}$ , an attack objective function  $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$  is defined as follows:*

$$f(\mathbf{x}, t) = \max_{i \neq t} \{h_i(\mathbf{x})\} - h_t(\mathbf{x}), \quad (1)$$

for which we use the shorthand  $f_t(\mathbf{x})$ . Generally,  $\mathbf{x}$  is classified as a member of  $t$  if and only if  $f_t(\mathbf{x}) < 0$ . An adversary can also enforce a desired attack confidence  $\kappa \in \mathbb{R}$  such that the attack is considered successful if and only if  $f_t(\mathbf{x}) < -\kappa$ .

The parameter  $\kappa$  allows to specify the desired confidence of an adversarial attack. If  $\kappa = 0$ , a *low-confidence* attack is performed where the object is merely misclassified; if  $\kappa > 0$  then the object is misclassified with higher confidence. It is worth observing that it is a common misconceptions of adversarial attacks that the feature-space perturbation must be minimal [49]. This was derived from the popularity of adversarial attacks in the image domain, in which you want to minimize the perceptual perturbation. It has been already shown in works [36, 50] that *high-confidence* attacks are much more effective in evading a classifier. The intuition is to minimize  $f_t$  by modifying  $\mathbf{x}$  in directions that follow the negative gradient of  $f_t$ , i.e., to get  $\mathbf{x}$  closer to the target class  $t$ .

In addition to the attack objective function, a considered problem-space domain may also come with constraints on the modification of the feature vectors. For example, in the image domain the value of pixels must be bounded between 0 and 255 [36]; in software, some features in  $\mathbf{x}$  may only be added but not removed (e.g., API calls [39]).

**DEFINITION 4 (FEATURE-SPACE CONSTRAINTS).** *We define  $\Omega$  as the set of feature-space constraints, i.e., a set of constraints on the possible feature-space modifications. The set  $\Omega$  reflects the requirements of realistic problem-space objects. Given an object  $\mathbf{x} \in \mathcal{X}$ , any modification of its feature values can be represented as a perturbation vector  $\delta \in \mathbb{R}^n$ ; if  $\delta$  satisfies  $\Omega$ , we borrow notation from model theory [29] and write  $\delta \models \Omega$ .*

These are in part due to domain constraints (e.g., the features are bounded between 0 and 255 in images, and some features cannot be changed) and in part due to allowed modifications constraints (e.g., only addition of software permissions may be allowed, or a maximum perturbation  $\delta_{max}$ ).

As examples of feature-space constraints, in the image domain [e.g., 36, 49] the perturbation  $\delta$  is subject to an upper bound based on  $l_p$  norms ( $\|\delta\|_p \leq \delta_{max}$ ), to preserve similarity to the original object; in the software domain [e.g., 39, 41], only some features of  $\mathbf{x}$  may be modified, such that  $\delta_{lb} \leq \delta \leq \delta_{ub}$  (where  $\delta_1 \leq \delta_2$  implies that each element of  $\delta_1$  is  $\leq$  the corresponding  $i$ -th element in  $\delta_2$ ).

Common instances of  $\Omega$  in the image domain are obtained by constraining the maximum number of perturbations with  $l_p$  norms ( $\|\delta\|_p \leq \delta_{max}$ ); we are more interested in constraints in the feature space, that determine which amount of perturbations the attacker do on each features ( $\delta_{lb} \leq \delta \leq \delta_{ub}$ , where  $v_1 \leq v_2$  implies that each element of  $v_1$  is  $\leq$  of the corresponding  $i$ -th element in  $v_2$ ).

We can now formalize the traditional feature-space attack as in related work [36, 18, 39, 49, 33].

**DEFINITION 5 (FEATURE-SPACE ATTACK).** *Given a machine learning classifier  $g$ , an object  $x \in \mathcal{X}$  with label  $y \in \mathcal{Y}$ , and a target label  $t \in \mathcal{Y}$ ,  $t \neq y$ , the adversary aims to identify a perturbation vector  $\delta \in \mathbb{R}^n$  such that  $g(x + \delta) = t$ . The desired perturbation can be achieved by solving the following optimization problem:*

$$\delta^* = \arg \min_{\delta \in \mathbb{R}^n} f_t(x + \delta) \quad (2)$$

$$\text{subject to: } \delta \models \Omega. \quad (3)$$

A feature-space attack is successful if  $f_t(x + \delta^*) < 0$  (or less than  $-\kappa$ , if the attacker enforces a desired attack confidence).

Without loss of generality, we observe that the feature-space attacks definition can be extended to ensure that the adversarial example is closer to the training data points (e.g., through the tuning of a parameter  $\lambda$  that penalizes adversarial examples generated in low density regions, as in the mimicry attacks of [18]).

## 2.2 Problem-Space Attacks

This section presents a novel formalization of problem-space attacks and introduces insights into the relationship between feature space and problem space.

**Inverse Feature-Mapping Problem.** The major challenge that complicates (and, in most cases, prevents) the direct applicability of gradient-driven feature-space attacks to find problem-space adversarial examples is the so-called *inverse feature-mapping problem* [66, 20, 70, 13, 18, 17]. As an extension, [70] discuss the *feature-problem space dilemma*, which highlights the difficulty of moving in both directions: from feature space to problem space, and from problem space to feature space. In most cases, the feature mapping function  $\varphi$  is not bijective, i.e., *not injective* and *not surjective*. This means that given  $z \in \mathcal{Z}$  with features  $x$ , and a feature-space perturbation  $\delta^*$ , there is no one-to-one mapping that allows going from  $x + \delta^*$  to an adversarial problem-space object  $z'$ . Nevertheless, there are two additional scenarios. If  $\varphi$  is not invertible but is *differentiable*, then it is possible to backpropagate the gradient of  $f_t(x)$  from  $\mathcal{X}$  to  $\mathcal{Z}$  to derive how the input can be changed in order to follow the negative gradient (e.g., to know which input pixels to perturbate to follow the gradient in the deep-learning latent feature space). If  $\varphi$  is not invertible and not differentiable, then the challenge is to find a way to map the adversarial feature vector  $x' \in \mathcal{X}$  to an adversarial object  $z' \in \mathcal{Z}$ , by applying a transformation to  $z$  in order to produce  $z'$  such that  $\varphi(z')$  is “as close as possible” to  $x'$ ; i.e., to follow the gradient towards the transformation that most likely leads to a successful evasion [53]. In problem-space settings such as software, the function  $\varphi$  is typically not invertible and not differentiable, so the search for transforming  $z$  to perform the attack cannot be purely gradient-based.

In this section, we consider the general case in which the feature mapping  $\varphi$  is not differentiable and not invertible (i.e., the most challenging setting), and we refer to this context to formalize problem-space evasion attacks.

First, we define a *problem-space transformation* operator through which we can alter problem-space objects. Due to their generality, we adapt the code transformation definitions from the *compiler engineering* literature [9, 70] to formalize general problem-space transformations.

**DEFINITION 6 (PROBLEM-SPACE TRANSFORMATION).** *A problem-space transformation  $T : \mathcal{Z} \rightarrow \mathcal{Z}$  takes a problem-space object  $z \in \mathcal{Z}$  as input and modifies it to  $z' \in \mathcal{Z}$ . We refer to the following notation:  $T(z) = z'$ .*

The possible problem-space transformations are either *addition*, *removal*, or *modification* (i.e., combination of addition and removal). In the case of computer programs, *obfuscation* is a special case of modification.

**DEFINITION 7 (TRANSFORMATION SEQUENCE).** *A transformation sequence  $\mathbf{T} = T_n \circ T_{n-1} \circ \dots \circ T_1$  is the subsequent application of problem-space transformations to an object  $z \in \mathcal{Z}$ .*

Intuitively, given a problem-space object  $z \in \mathcal{Z}$  with label  $y \in \mathcal{Y}$ , the purpose of the adversary is to find a transformation sequence  $\mathbf{T}$  such that the transformed object  $\mathbf{T}(z)$  is classified into any target class  $t$  chosen by the adversary ( $t \in \mathcal{Y}, t \neq y$ ). One way to achieve such a transformation is to first compute a feature-space perturbation  $\delta^*$ , and then modify the problem-space object  $z$  so that features corresponding to  $\delta^*$  are carefully altered. However, in the general case where the feature mapping  $\varphi$  is neither invertible nor differentiable, the adversary must perform a search in the problem-space that approximately follows the negative gradient in the feature space. However, this search is not unconstrained, because the adversarial problem-space object  $\mathbf{T}(z)$  must be realistic.

**Problem-Space Constraints.** Given a problem-space object  $z \in \mathcal{Z}$ , a transformation sequence  $\mathbf{T}$  must lead to an object  $z' = \mathbf{T}(z)$  that is valid and realistic. To express this formally, we identify four main types of constraints common to any problem-space attack:

- (1) *Available transformations*, which describe which modifications can be performed in the problem-space by the attacker (e.g., only addition and not removal).
- (2) *Preserved semantics*, the semantics to be preserved while mutating  $z$  to  $z'$ , with respect to specific feature abstractions which the attacker aims to be resilient against (e.g., in programs, the transformed object may need to produce the same dynamic call traces). Semantics may also be preserved by construction [e.g., 70].
- (3) *Plausibility* (or *Inconspicuousness*), which describes which (qualitative) properties must be preserved in mutating  $z$  to  $z'$ , so that  $z$  appears realistic upon manual inspection. For example, often an adversarial image must look like a valid image from the training distribution [36]; a program's source code must look manually written and not artificially or inconsistently altered [70]. In the general case, verification of plausibility may be hard to automate and may require human analysis.
- (4) *Robustness to preprocessing*, which determines which non-ML techniques could disrupt the attack (e.g., filtering in images, dead code removal in programs).

These constraints have been sparsely mentioned in prior literature [18, 70, 35, 49], but have never been identified together as a set for problem-space attacks. When designing a novel problem-space attack, it is fundamental to explicitly define these four types of constraints, to clarify strengths and weaknesses. While we believe that this framework captures all nuances of the current state-of-the-art for a thorough evaluation and comparison, we welcome future research that uses this as a foundation to identify new constraints.

We now introduce formal definitions for the constraints. First, similarly to [39, 49], we define the space of available transformations.

**DEFINITION 8 (AVAILABLE TRANSFORMATIONS).** *We define  $\mathcal{T}$  as the space of available transformations, which determines which types of automated problem-space transformations  $T$  the attacker can perform. In general, it determines if and how the attacker can add, remove, or edit parts of the original object  $z \in \mathcal{Z}$  to obtain a new object  $z' \in \mathcal{Z}$ . We write  $\mathbf{T} \in \mathcal{T}$  if a transformation sequence consists of available transformations.*

For example, the pixels of an image may be modified only if they remain within the range of integers 0 to 255 [e.g., 36]; in programs, an adversary may only add valid no-op API calls to ensure that modifications preserve functionality [e.g., 57].



Moreover, the attacker needs to ensure that some semantics are preserved during the transformation of  $z$ , according to some feature abstractions. Semantic equivalence is known to be generally undecidable [26, 70]; hence, as in [26], we formalize semantic equivalence through *testing*, by borrowing notation from *denotational semantics* [4].

**DEFINITION 9 (PRESERVED SEMANTICS).** *Let us consider two problem-space objects  $z$  and  $z' = \mathbf{T}(z)$ , and a suite of automated tests  $\Upsilon$  to verify preserved semantics. We define  $z$  and  $z'$  to be semantically equivalent with respect to  $\Upsilon$  if they satisfy all its tests  $\tau \in \Upsilon$ , where  $\tau : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{B}$ . In particular, we denote semantics equivalence with respect to a test suite  $\Upsilon$  as follows:*

$$\llbracket z \rrbracket^\tau = \llbracket z' \rrbracket^\tau, \forall \tau \in \Upsilon, \quad (4)$$

where  $\llbracket z \rrbracket^\tau$  denotes the semantics of  $z$  induced during test  $\tau$ .

Informally,  $\Upsilon$  consists of tests that are aimed at evaluating whether  $z$  and  $z'$  (or parts of them) lead to the same abstract representations in a certain feature space. In other words, the tests in  $\Upsilon$  model preserved semantics. For example, in programs a typical test aims to verify that malicious functionality is preserved; this is done through tests where, given a certain test input, the program produces exactly the same output [26]. Additionally, the attacker may want to ensure that an adversarial program ( $z'$ ) leads to the same instruction trace as its benign version ( $z$ )—so as not to raise suspicion in feature abstractions derived from dynamic analysis.

Plausibility is more subjective than semantic equivalence, but in many scenarios it is critical that an adversarial object appear inconspicuous when manually audited by a human. To be plausible, an analyst must believe the adversarial object is a valid member of the problem-space distribution.

**DEFINITION 10 (PLAUSIBILITY).** *We define  $\Pi$  as the set of (typically) manual tests to verify plausibility. We say  $z$  looks like a valid member of the data distribution to a human being if it satisfies all tests  $\pi \in \Pi$ , where  $\pi : \mathcal{Z} \rightarrow \mathbb{B}$ .*

Plausibility is often hard to verify automatically; previous work has often relied on user studies with domain experts to judge the plausibility of the generated objects (e.g., program plausibility in [70], realistic eyeglass frames in [34]). Plausibility in software-related domains may also be enforced by construction during the transformation process, e.g., by relying on automated software transplantation [26, 46].

In addition to semantic equivalence and plausibility, the adversarial problem-space objects need to ensure they are robust to non-ML automated *preprocessing* techniques that could alter properties on which the adversarial attack depends, thus compromising the attack.

**DEFINITION 11 (ROBUSTNESS TO PREPROCESSING).** *We define  $\Lambda$  as the set of preprocessing operators an object  $z' = \mathbf{T}(z)$  should be resilient to. We say  $z'$  is robust to preprocessing if  $\mathbf{A}(\mathbf{T}(z)) = \mathbf{T}(z)$  for all  $\mathbf{A} \in \Lambda$ , where  $\mathbf{A} : \mathcal{Z} \rightarrow \mathcal{Z}$  simulates an expected preprocessing.*

Examples of preprocessing operators in  $\Lambda$  include compression to remove pixel artifacts (in images), filters to remove noise (in audio), and program analysis to remove dead or redundant code (in programs).

Properties affected by preprocessing are often related to *fragile and spurious features* learned by the target classifier. While taking advantage of such features may be necessary to demonstrate the weaknesses of the target model, an attacker should be aware that these brittle features are usually the first to change when a model is improved. Given this, a stronger attack is one that does not rely on them.

As a concrete example, in an attack on authorship attribution, [70] purposefully omit layout features (such as the use of spaces vs. tabs) which are trivial to change. Additionally, [35] discovered

the presence of font objects was a critical (but erroneously discriminative) feature following their problem-space attack on PDF malware. These are features that are cheap for an attacker to abuse but can be easily removed by the application of some preprocessing. As a defender, investigation of this constraint will help identify features that are weak to adversarial attacks. Note that knowledge of preprocessing can also be exploited by the attacker (e.g., in *scaling attacks* [75]).

We can now define a fundamental set of constraint elements in the problem space based on the previous definitions.

**DEFINITION 12 (PROBLEM-SPACE CONSTRAINTS).** *We define the problem-space constraints  $\Gamma = \{\mathcal{T}, \Upsilon, \Pi, \Lambda\}$  as the set of all constraints satisfying  $\mathcal{T}, \Upsilon, \Pi, \Lambda$ . We write  $\mathbf{T}(z) \models \Gamma$  if a transformation sequence applied to object  $z \in \mathcal{Z}$  satisfies all the problem-space constraints, and we refer to this as a valid transformation sequence. The problem-space constraints  $\Gamma$  determine the feature-space constraints  $\Omega$ , and we denote this relationship as  $\Gamma \vdash \Omega$  (i.e.,  $\Gamma$  determines  $\Omega$ ); with a slight abuse of notation, we can also write that  $\Omega \subseteq \Gamma$ , because some constraints may be specific to the problem space (e.g., program size similar to that of benign applications) and may not be possible to enforce in the feature space  $\mathcal{X}$ .*

**Side-Effect Features.** Satisfying the problem-space constraints  $\Gamma$  further complicates the inverse feature mapping, as  $\Gamma$  is a superset of  $\Omega$ . Moreover, enforcing  $\Gamma$  may require substantially altering an object  $z$  to ensure satisfaction of all constraints during mutations. Let us focus on an example in the software domain, so that  $z$  is a program with features  $\mathbf{x}$ ; if we want to transform  $z$  to  $z'$  such that  $\varphi(z') = \mathbf{x} + \delta$ , we may want to add to  $z$  a program  $o$  where  $\varphi(o) = \delta$ . However, the union of  $z$  and  $o$  may have features different from  $\mathbf{x} + \delta$ , because other consolidation operations are required (e.g., name deduplication, class declarations, resource name normalization)—which cannot be feasibly computed in advance for each possible object in  $\mathcal{Z}$ . Hence, after modifying  $z$  in an attempt to obtain a problem-space object  $z'$  with certain features (e.g., close to  $\mathbf{x} + \delta$ ), the attacker-modified object may have some additional features that are not related to the intended transformation (e.g., adding an API which maps to a feature in  $\delta$ ), but are required to satisfy all the problem-space constraints in  $\Gamma$  (e.g., inserting valid parameters for the API call, and importing dependencies for its invocation). We call *side-effect features*  $\boldsymbol{\eta}$  the features that are altered in  $z' = \mathbf{T}(z)$  specifically for the satisfaction of problem-space constraints. We observe that these features do not follow any particular direction of the gradient, and hence they could have both a positive or negative impact on the classification score.

**Analogy with Projection.** Figure 1 presents an analogy between side-effect features  $\boldsymbol{\eta}$  and the notion of *projection* in numerical optimization [7], which helps explain the nature and impact of  $\boldsymbol{\eta}$  in problem-space attacks. The right half corresponds to higher values of a discriminant function  $h(\mathbf{x})$  and the left half to lower values. The vertical central curve (where the heatmap value is equal to zero) represents the decision boundary: objects on the left-half are classified as negative (e.g., benign), and objects on the right-half as positive (e.g., malicious). The goal of the adversary is to conduct a *maximum confidence attack* that has an object misclassified as the negative class. The thick solid line represents the *feasible feature space* determined by constraints  $\Omega$ , and the thin solid line the *feasible problem space* determined by  $\Gamma$  (which corresponds to two unconnected areas). We assume that the initial object  $\mathbf{x} \in \mathcal{X}$  is always within the feasible problem space. In this example, the attacker first conducts a gradient-based attack in the feature space on object  $\mathbf{x}$ , which results in a feature vector  $\mathbf{x} + \delta^*$ , which is classified as negative with high-confidence. However, this point is not in the feasibility space of constraints  $\Gamma$ , which is more restrictive than that of  $\Omega$ . Hence, the attacker needs to find a *projection* that maps  $\mathbf{x} + \delta^*$  back to the feasible problem-space regions, which leads to the addition of a side-effect feature vector  $\boldsymbol{\eta}$ .

**DEFINITION 13 (SIDE-EFFECT FEATURE VECTOR).** *We define  $\boldsymbol{\eta}$  as the side-effect feature vector that results from enforcing  $\Gamma$  while choosing a sequence of transformations  $\mathbf{T}$  such that  $\mathbf{T}(z) \models \Gamma$ . In other*






Fig. 1. Example of projection of the feature-space attack vector  $\mathbf{x} + \delta^*$  in the *feasible* problem space, resulting in side-effect features  $\eta$ . The background displays the value of the discriminant function  $h(\mathbf{x})$ , where negative values indicate the target class of the evasion attack. Small arrows represent directions of the negative gradient. The thick solid line represents the *feasible* feature space determined by  $\Omega$ , and the thin solid line that determined by  $\Gamma$  (which is more restrictive). The dotted arrow represents the gradient-based attack  $\mathbf{x} + \delta^*$  derived from  $\mathbf{x}$ , which is then projected into  $\mathbf{x} + \delta^* + \eta$  to fit into the feasible problem space.

words,  $\eta$  are the features derived from the projection of a feature-space attack onto a feasibility region that satisfies problem-space constraints  $\Gamma$ .

We observe that in settings where the feature mapping  $\varphi$  is neither differentiable nor invertible, and where the problem-space representation is very different from the feature-space representation (e.g., unlike in images or audio), it is generally infeasible or impossible to compute the exact impact of side-effect features on the objective function in advance—because the set of problem-space constraints  $\Gamma$  cannot be expressed analytically in closed-form. Hence the attacker needs to find a transformation sequence  $\mathbf{T}$  such that  $\varphi(\mathbf{T}(z)) = \varphi(z')$  is within the feasibility region of problem-space constraints  $\Gamma$ .

It is relevant to observe that, in the general case, if an object  $z_o$  is added to (or removed from) two different objects  $z_1$  and  $z_2$ , it is possible that the resulting side-effect feature vectors  $\eta_1$  and  $\eta_2$  are different (e.g., in the software domain [70]).

Semantic preservation constraints are essential, as they ensure that any modifications or perturbations maintain its original functionality. In certain situations, it may be necessary to minimize overall perturbations while adhering to these constraints, and their impact varies depending on the considered domain. For instance, in domains where the feature mapping function is invertible and differentiable, and there is no inverse feature mapping issue, such as in computer vision tasks, the side effect features may also be influenced. Conversely, in areas like source code tasks, the side effect features might remain unaffected. This highlights that the decision to minimize perturbations depends on the specific requirements for preserving semantics.

**Considerations on Attack Confidence.** There are some important characteristics of the impact of the side-effect features  $\eta$  on the attack objective function. If the attacker performs a *maximum-confidence attack* in the feature space under constraints  $\Omega$ , then the confidence of the problem-space attack will always be *lower or equal* than the one in the feature-space attack. This is intuitively represented in Figure 1, where the point is moved to the maximum-confidence attack area within  $\Omega$ , and the attack confidence is reduced after projection to the feasibility space of the problem space, induced by  $\Gamma$ . In general, the confidence of the feature- and problem-space attacks could be equal, depending on the constraints  $\Omega$  and  $\Gamma$ , and on the shape of the discriminant function  $h$ , which is also not necessarily convex (e.g., in deep learning [31]). In the case of *low-confidence* feature-space

attacks, projecting into the problem-space feasibility constraint may result in a positive or negative impact (not known a priori) on the value of the discriminant function. This can be seen from Figure 1, where the object  $\mathbf{x} + \delta^*$  would be found close to the center of the plot, where  $h(\mathbf{x}) = 0$ .

**Problem-Space Attack.** We now have all the components that are required to formalize a problem-space attack.

**DEFINITION 14 (PROBLEM-SPACE ATTACK).** *We define a problem-space attack as the problem of finding the sequence of valid transformations  $\mathbf{T}$  for which the object  $\mathbf{z} \in \mathcal{Z}$  with label  $y \in \mathcal{Y}$  is misclassified to a target class  $t \in \mathcal{Y}$  as follows:*

$$\operatorname{argmin}_{\mathbf{T} \in \mathcal{T}} f_t(\varphi(\mathbf{T}(\mathbf{z}))) = f_t(\mathbf{x} + \delta^* + \boldsymbol{\eta}) \quad (5)$$

$$\text{subject to: } \llbracket \mathbf{z} \rrbracket^\tau = \llbracket \mathbf{T}(\mathbf{z}) \rrbracket^\tau, \quad \forall \tau \in \Upsilon \quad (6)$$

$$\pi(\mathbf{T}(\mathbf{z})) = 1, \quad \forall \pi \in \Pi \quad (7)$$

$$\mathbf{A}(\mathbf{T}(\mathbf{z})) = \mathbf{T}(\mathbf{z}), \quad \forall \mathbf{A} \in \Lambda \quad (8)$$

where  $\boldsymbol{\eta}$  is a side-effect feature vector that separates the feature vector generated by  $\mathbf{T}(\mathbf{z})$  from the theoretical feature-space attack  $\mathbf{x} + \delta^*$  (under constraints  $\Omega$ ). An equivalent, more compact, formulation is as follows:

$$\operatorname{argmin}_{\mathbf{T} \in \mathcal{T}} f_t(\varphi(\mathbf{T}(\mathbf{z}))) = f_t(\mathbf{x} + \delta^* + \boldsymbol{\eta}) \quad (9)$$

$$\text{subject to: } \mathbf{T}(\mathbf{z}) \models \Gamma. \quad (10)$$

**Search Strategy.** The typical search strategy for adversarial perturbations in feature-space attacks is based on following the negative gradient of the objective function through some numerical optimization algorithm, such as stochastic gradient descent [49, 36, 50]. However, it is not possible to directly apply gradient descent in the general case of problem-space attacks, when the feature space is not invertible nor differentiable [70, 49]; and it is even more complicated if a transformation sequence  $\mathbf{T}$  produces side-effect features  $\boldsymbol{\eta} \neq \mathbf{0}$ . In the problem space, we identify two main types of search strategy: *problem-driven* and *feature-driven*. In the problem-driven approach, the search of the optimal  $\mathbf{T}$  proceeds heuristically by beginning with random mutations of the object  $\mathbf{z}$ , and then learning from experience how to appropriately mutate it further in order to misclassify it to the target class (e.g., using Genetic Programming [35] or variants of Monte Carlo tree search [70]). This approach iteratively uses local approximations of the negative gradient to mutate the objects. The feature-driven approach attempts to identify mutations that follow some signal from the feature space. This can be, for example, a target feature vector (e.g., a mimicry attack [24]) or the negative gradient by relying on an approximate inverse feature mapping (e.g., in PDF malware [20], in Android malware [46]). If a search strategy equally makes extensive use of both problem-driven and feature-driven methods, we call it a *hybrid* strategy. We note that search strategies may have different trade-offs in terms of *effectiveness* and *costs*, depending on the time and resources they require. While there are some promising avenues in this challenging but important line of research [54], it warrants further investigation in future work.

Feature-space attacks can still give us some useful information: before searching for a problem-space attack, we can verify whether a feature-space attack exists, which is a necessary condition for realizing the problem-space attack.

**THEOREM 1 (NECESSARY CONDITION FOR PROBLEM-SPACE ATTACKS).** *Given a problem-space object  $\mathbf{z} \in \mathcal{Z}$  of class  $y \in \mathcal{Y}$ , with features  $\varphi(\mathbf{z}) = \mathbf{x}$ , and a target class  $t \in \mathcal{Y}$ ,  $t \neq y$ , there exists a transformation sequence  $\mathbf{T}$  that causes  $\mathbf{T}(\mathbf{z})$  to be misclassified as  $t$  only if there is a solution for the*

feature-space attack under constraints  $\Omega$ . More formally, only if:

$$\exists \delta^* = \arg \min_{\delta \in \mathbb{R}^n: \delta \models \Omega} f_t(\mathbf{x} + \delta) : f_t(\mathbf{x} + \delta^*) < 0. \quad (11)$$

The proof of [Theorem 1](#) is in [Appendix C](#). We observe that [Theorem 1](#) is necessary but *not sufficient* because, although it is not required to be invertible or differentiable, some sort of “mapping” between problem- and feature-space perturbations needs to be known by the attacker. A *sufficient condition* for a problem-space attack, reflecting the attacker’s ideal scenario, is knowledge of a set of problem-space transformations which can alter feature values arbitrarily. This describes the scenario for some domains, such as images [36, 27], in which the attacker can modify any pixel value of an image independently.

**THEOREM 2 (SUFFICIENT CONDITION FOR PROBLEM-SPACE ATTACKS).** *Given a problem-space object  $z \in \mathcal{Z}$  of class  $y \in \mathcal{Y}$ , with features  $\varphi(z) = \mathbf{x}$ , and a target class  $t \in \mathcal{Y}$ ,  $t \neq y$ , there exists a transformation sequence  $\mathbf{T}$  that causes  $\mathbf{x}$  to be misclassified as  $t$  if [Equation 11](#) and [Equation 12](#) are satisfied:*

$$\exists \delta^* = \arg \min_{\delta \in \mathbb{R}^n: \delta \models \Omega} f_t(\mathbf{x} + \delta) : f_t(\mathbf{x} + \delta^*) < 0 \quad (11)$$

$$\forall \delta \in \mathbb{R}^n : \delta \models \Omega, \quad \exists \mathbf{T} : \mathbf{T}(z) \models \Gamma, \varphi(\mathbf{T}(z)) = \mathbf{x} + \delta \quad (12)$$

*Informally, an attacker is always able to find a problem-space attack if a feature-space attack exists (necessary condition) and they know problem-space transformations that can modify any feature by any value (sufficient condition).*

The proof of [Theorem 2](#) is in [Appendix C](#). In the general case, while there may exist an optimal feature-space perturbation  $\delta^*$ , there may *not* exist a problem-space transformation sequence  $\mathbf{T}$  that alters the feature space of  $\mathbf{T}(z)$  exactly so that  $\varphi(\mathbf{T}(z)) = \mathbf{x} + \delta^*$ . This is because, in practice, given a target feature-space perturbation  $\delta^*$ , a problem-space transformation may generate a vector  $\varphi(\mathbf{T}(z)) = \mathbf{x} + \delta^* + \boldsymbol{\eta}^*$ , where  $\boldsymbol{\eta}^* \neq \mathbf{0}$  (i.e., where there may exist at least one  $i$  for which  $\eta_i \neq 0$ ) due to the requirement that problem-space constraints  $\Gamma$  must be satisfied. This prevents easily finding a problem-space transformation that follows the negative gradient. Given this, the attacker is forced to apply some search strategy based on the available transformations.

**COROLLARY 2.1.** *If [Theorem 2](#) is satisfied only on a subset of feature dimensions  $X_i$  in  $\mathcal{X}$ , which collectively create a subspace  $\mathcal{X}_{eq} \subset \mathcal{X}$ , then the attacker can restrict the search space to  $\mathcal{X}_{eq}$ , for which they know that an equivalent problem/feature-space manipulation exists.*

### 2.3 Describing problem-space attacks in different domains

[Table 1](#) briefly illustrates the main parameters of problem-space attacks by considering a representative set of adversarial attacks in different domains: Android (our attack proposed in §3 and one recent proposal from the literature [83]), code attribution [70, 67], Windows [53, 80], text [65, 71], and PDF [35, 24]. The references contained in the conference version have been omitted from this document. For a deeper understanding, see also the table in the conference version.

This table shows how our formalization can be used to describe and compare problem-space attacks by highlighting their key concepts, strengths and weaknesses. For example, although transformations obviously depend on the domain, we observe a general trend that modifications adding features are more common than modifications changing or removing features. This can be attributed to the mere complexity of preserving semantics if content is rewritten, and motivates our approach in §3 to use automated software transplantation with opaque predicates. Moreover, various search strategies are used. The strategy is also often linked with the threat scenario. A feature-driven or hybrid strategy is rather used with more knowledge while a problem-driven strategy is chosen

Domain	Work	Problem-Space Constraints				Threat Model & Search Strategy
		Available Transformers $\mathcal{T}$	Plausibility $\Pi$	$\Upsilon$	$\Delta$	
Android	<i>Our Attack</i>	Code addition through automated software transplantation	● (i) Code is realistic by construction through use of exhaustive automated software transplantation. (ii) Mutated apps install and start on an emulator.	●	●	PK: Feature-driven
	[83]	Code addition of targeted Android components	● (i) Code is realistic by construction through use of limited set of perturbations. (ii) Mutated apps install and start on an emulator. (iii) Added intents could lead to unintended effects	●	●	ZK: Problem-driven
Code Attrib.	[70]	Rewrite source code with targeted code transformations (excl. layout on purpose)	● The code does not look suspicious and seems written by a human (survey with developers)	●	●	ZK: Problem-driven
	[67]	Rewrite code by exploiting layout, comments, and control-flow flattening	○ Modification visible by e.g. adding comments from target or using control-flow flattening	●	○	LK: Feature-driven
Windows	[53]	Append bytes at the end	● Appended bytes not visible by user, but easily by security expert	●	○	PK: Feature-driven
	[80]	Binary diversification: code is rewritten either through in-place randomization (IPR) or code displacement (DISP)	● Executed code is modified which might be difficult to spot	●	●	PK: Hybrid, ZK: Problem-driven
Text	[65]	Apply character-level (e.g., spelling mistakes) and word-level perturbations (e.g., synonyms)	● Modifications such as typos are visible, but might be overlooked	●	●	PK: Hybrid, ZK: Problem-driven
	[71]	Apply word-level perturbations (e.g., synonyms)	● Synonyms should preserve meaning, but no guarantee possible	●	●	ZK: Problem-driven
PDF	[35]	Addition/Removal of elements in the PDF tree structure	● PDF can be opened, but PDF elements are modified which might be visible	●	●	ZK: Problem-driven
	[24]	Add content in unused area between cross-reference table and trailer	● Modifications not visible by user, but easily by security expert	●	○	LK, PK: Feature-driven

Table 1. Problem-space attacks in different domains (with  $\mathcal{Z} \neq \mathcal{F}$ ), summarized with our formalization.

Recall notation:  $\Upsilon$ =preserved semantics;  $\Delta$ =robustness to preprocessing; ZK=zero knowledge, PK=perfect knowledge, and LK=limited knowledge

with zero knowledge. Finally, the semantics are usually preserved while plausibility and robustness to preprocessing are more often neglected. For instance, attacks that add content at the end or in unused areas of a file [e.g., 53, 24] can be reverted by considering the actual file payload. Likewise, layout modifications and added comments in source code [e.g., 67] can be easily reverted by using code formatting tools or even a simple regular expression, respectively.

### 3 ATTACK ON ANDROID

Our formalization of problem-space attacks reveals weaknesses in prior approaches. Hence, we propose—through our formalization—a novel problem-space attack in this domain that overcomes these limitations, especially in terms of preserved semantics and preprocessing robustness (see §2.3 and §7 for a detailed comparison).

#### 3.1 Threat Model

We assume an attacker with *perfect knowledge*  $\theta_{PK} = (\mathcal{D}, \mathcal{X}, g, \mathbf{w})$  (see Appendix B for details on threat models). This follows Kerckhoffs’ principle [1] and ensures a defense does not rely on “security by obscurity” by unreasonably assuming some properties of the defense can be kept secret [62]. Although deep learning has been extensively studied in adversarial attacks, recent research [e.g., 69] has shown that—if retrained frequently—the DREBIN classifier [21] achieves state-of-the-art performance for Android malware detection, which makes it a suitable target classifier for our attack. DREBIN relies on a linear SVM, and embeds apps in a *binary* feature-space  $\mathcal{X}$  which captures the presence/absence of components in Android applications in  $\mathcal{Z}$  (such as permissions, URLs, Activities, Services, strings). We assume to know classifier  $g$  and feature-space  $\mathcal{X}$ , and train the parameters  $\mathbf{w}$  with SVM hyperparameter  $C = 1$ , as in the original DREBIN paper [21]. Using DREBIN also enables us to evaluate the effectiveness of our problem-space attack against a recently proposed hardened variant, Sec-SVM [39]. Sec-SVM enforces more evenly distributed

feature weights, which require an attacker to modify more features to evade detection. Although the authors of [39] evaluated the robustness of Sec-SVM in the feature-space  $X$  for increasing  $l_0$  attack norms, in §4 we demonstrate that it is feasible to evade Sec-SVM in the problem-space.

We consider an attacker intending to evade detection based on *static analysis*, without relying on code obfuscation as it may increase suspiciousness of the apps [59, 28] (see §6).

### 3.2 Available Transformations

We use *automated software transplantation* [26] to extract slices of bytecode (i.e., *gadgets*) from benign *donor* applications and inject them into a malicious *host*, to mimic the appearance of benign apps and induce the learning algorithm to misclassify the malicious host as benign.<sup>1</sup> An advantage of this process is that we avoid relying on a hardcoded set of transformations [e.g., 70]; this ensures adaptability across different application types and time periods. In this work, we consider only *addition* of bytecode to the malware—which ensures that we do not hinder the malicious functionality of the app.

**Organ Harvesting.** In order to augment a malicious host with a given *benign feature*  $X_i$ , we must first extract a bytecode gadget  $\rho$  corresponding to  $X_i$  from some donor app. As we intend to produce realistic examples, we use *program slicing* [2] to extract a functional set of statements that includes a reference to  $X_i$ . The final gadget consists of the target reference (*entry point*  $L_o$ ), a forward slice (*organ*  $o$ ), and a backward slice (*vein*  $v$ ). We first search for  $L_o$ , corresponding to an appearance of code corresponding to the desired feature in the donor. Then, to obtain  $o$ , we perform a context-insensitive forward traversal over the donor’s System Dependency Graph (SDG), starting at the entry point, transitively including all of the functions called by any function whose definition is reached. Finally, we extract  $v$ , containing all statements needed to construct the parameters at the entry point. To do this, we compute a backward slice by traversing the SDG in reverse. Note that while there is only one organ, there are usually multiple veins to choose from, but only one is necessary for the transplantation. When traversing the SDG, class definitions that will certainly be already present in the host are excluded (e.g., system packages such as android and java). For example, for an Activity feature where the variable intent references the target Activity of interest, we might extract the invocation `startActivity(intent)` (entry point  $L_o$ ), the class implementation of the Activity itself along with any referenced classes (organ  $o$ ), and all statements necessary to construct intent with its parameters (vein  $v$ ). There is a special case for Activities that have no corresponding vein in the bytecode (e.g., a MainActivity or an Activity triggered by an intent filter declared in the Manifest); here, we provide an *adapted vein*, a minimal Intent creation and `startActivity()` call adapted from a previously mined benign app that will trigger the Activity. Note that organs with original veins are always prioritized above those without.

In this work we have extended the feature extraction capabilities modifying the existing implementation with the goal to amplify the total amount of available gadgets. Indeed, our framework currently supports the extraction of all types of features supported by the considered feature space [21]. Specifically, while in the conference paper we were supporting the extracting of Activity and URL components, now we also support the extraction of Suspicious calls, API calls, Provider, Receivers and Permissions. We do not focus on intent-filters due to possible disruptions of the original semantic §3.3.

With each additional component, we were required to devise a unique strategy for its extraction, thereby augmenting our existing capabilities. Take, for example, the process of extracting services,

<sup>1</sup>Our approach is generic and it would be immediate to do the opposite, i.e., transplant malicious code into a benign app. However, this would require a dataset with *annotated* lines of malicious code. For this practical reason and for the sake of clarity of this section, we consider only the scenario of adding benign code parts to a malicious app.



which necessitated the identification of 'startService' methods as a starting point, followed by tracking the subsequent connections from these methods. This technique was equally applied to other components, including Receivers and Providers. In the case of other categories such as APIs or permissions, our approach centered on pinpointing the specific functions tied to these elements and employing program slicing for identify and extract the whole set of dependencies. This systematic and tailored approach enabled us to effectively dissect and analyze the complexities inherent in each component and category, enhancing the overall scope and precision of our extraction process. Please refer to [Appendix G](#) for an in depth realistic example.

In order to preserve the final logic behind every type of feature we have personalized the extraction mechanism. As explained before, we need to extract pieces of code that reference or use the feature we are interested to extract in order to avoid potential dead code, e.g. for Activity we focus on extracting also the intent which invokes the target feature. This could be more difficult for permissions because we need to identify which piece of code may actually use that and it is not directly referenced. In that specific case, we need indeed a mapping between the actual permission and the possible system calls which would imply its usage.

**Organ Implantation.** In order to implant some gadget  $\rho$  into a host, it is necessary to identify an injection point  $L_H$  where  $v$  should be inserted. Implantation at  $L_H$  should fulfill two criteria: firstly, it should maintain the syntactic validity of the host; secondly, it should be as unnoticeable as possible so as not to contribute to any violation of plausibility. To maximize the probability of fulfilling the first criterion, we restrict  $L_H$  to be between two statements of a class definition in a non-system package. For the second criterion, we take a heuristic approach by using *Cyclomatic Complexity* (CC)—a software metric that quantifies the code complexity of components within the host—and choosing  $L_H$  such that we maintain existing homogeneity of CC across all components. Finally, the host entry point  $L_H$  is inserted into a *randomly chosen* function among those of the selected class, to avoid creating a pattern that might be identified by an analyst.

### 3.3 Preserved Semantics

Given an application  $z$  and its modified (adversarial) version  $z'$ , we aim to ensure that  $z$  and  $z'$  lead to the same dynamic execution, i.e., the malicious behavior of the application is preserved. We enforce this by construction by wrapping the newly injected execution paths in conditional statements that always return `False`. This guarantees the newly inserted code is never executed at runtime—so users will not notice anything odd while using the modified app. In [§3.4](#), we describe how we generate such conditionals without leaving artifacts.

To further preserve semantics, we also decide to omit `intent-filter` elements as transplantation candidates. For example, an `intent-filter` could declare the app as an eligible option for reading PDF files; consequently, whenever attempting to open a PDF file, the user would be able to choose the host app, which (if selected) would trigger an Activity defined in the transplanted benign bytecode—violating our constraint of preserving dynamic functionality.

### 3.4 Robustness to Preprocessing

Program analysis techniques that perform redundant code elimination would remove unreachable code. Our evasion attack relies on features associated with the transplanted code, and to preserve semantics we need conditional statements that always resolve to `False` at runtime; so, we must subvert static analysis techniques that may identify that this code is never executed. We achieve this by relying on *opaque predicates* [11], i.e., carefully constructed obfuscated conditions where the outcome is always known at design time (in our case, `False`), but the actual truth value is difficult or impossible to determine during a static analysis. We refer the reader to [Appendix D](#) for a detailed description of how we generate strong opaque predicates and make them look legitimate.

### 3.5 Plausibility

In our model, an example is satisfactorily plausible if it resembles a real, functioning Android application (i.e., is a valid member of the problem-space  $\mathcal{Z}$ ). Our methodology aims to maximize the plausibility of each generated object by injecting full slices of bytecode from *real* benign applications. There is only one case in which we inject artificial code: the opaque predicates that guard the entry point of each gadget (see Appendix D for an example). In general, we can conclude that plausibility is guaranteed *by construction* thanks to the use of automated software transplantation [26]. This contrasts with other approaches that inject *standalone* API calls and URLs or *no-op* operations [e.g., 57] that are completely orphaned and unsupported by the rest of the bytecode (e.g., an API call result that is never used).

We also practically assess that each mutated app still functions properly after modification by installing and running it on an Android emulator. Although we are unable to thoroughly explore every path of the app in this automated manner, it suffices as a smoke test to ensure that we have not fundamentally damaged the structure of the app.

### 3.6 Search Strategy

We propose a *feature-driven* search strategy based on a *greedy algorithm*, which aims to follow the gradient direction by transplanting a gadget with benign features into the malicious host. There are two main phases: *Initialization* (Ice-Box Creation) and *Attack* (Adversarial Program Generation). This section offers an overview of the proposed search strategy, and the detailed steps are reported in Appendix F.

**Initialization Phase (Ice-Box Creation).** We first harvest gadgets from potential donors and collect them in an *ice-box*  $G$ , which is used for transplantation at attack time. The main reason for this, instead of looking for gadgets on-the-fly, is to have an immediate estimate of the *side-effect features* when each gadget is considered for transplantation. Looking for gadgets on-the-fly is possible, but may lead to less optimal solutions and uncertain execution times.

For the initialization we aim to gather gadgets that move the score of an object towards the benign class (i.e., negative score), hence we consider the classifier's top  $n_f$  benign features (i.e., with negative weight). For each of the top- $n_f$  features, we extract  $n_d$  candidate gadgets, excluding those that lead to an overall positive (i.e., malicious) score. We recall that this may happen even for benign features since the context extracted through forward and backward slicing may contain many other features that are indicative of maliciousness. We empirically verify that with  $n_f = 500$  and  $n_d = 5$  we can create a successfully evasive app for all malware samples in our experiments. To estimate the side-effect feature vectors for the gadgets, we inject each into a *minimal app*, i.e., an Android app we developed with minimal functionality (see Appendix F). Notably, the ice-box can be expanded over time, as long as the target classifier does not change its weights significantly. Algorithm 1 in Appendix F outlines the steps of the initialization phase.

**Attack Phase.** We aim to automatically mutate  $z$  into  $z'$  so that it is misclassified as goodwill, i.e.,  $h(\varphi(z')) < 0$ , by transplanting harvested gadgets from the ice-box  $G$ . First we search for the list of ice-box gadgets that should be injected into  $z$ . Each gadget  $\rho_j$  in the ice-box  $G$  has feature vector  $\mathbf{r}_j$  which includes the desired feature and side-effect features. We consider the actual feature-space contribution of gadget  $i$  to the malicious host  $z$  with features  $\mathbf{x}$  by performing the set difference of the two binary vectors,  $\mathbf{r}_j \wedge \neg \mathbf{x}$ . We then sort the gadgets in order of decreasing negative contribution, which ideally leads to a faster convergence of  $z$ 's score to a benign value. Next we filter this candidate list to include gadgets *only if* they satisfy some practical feasibility criteria. We define a *check\_feasibility* function which implements some heuristics to limit the excessive increase of certain statistics which would raise suspiciousness of the app. Preliminary experiments revealed

a tendency to add too many permissions to the Android Manifest, hence, we empirically enforce that candidate gadgets add no more than 1 new permission to the host app. Moreover, we do not allow addition of permissions listed as *dangerous* in the Android documentation [78]. The other app statistics remain reasonably within the distribution of benign apps (more discussion in §4), and so we decide not to enforce a limit on them. The remaining candidate gadgets are iterated over and for each candidate  $\rho_j$ , we combine the gadget feature vector  $r_j$  with the input malware feature vector  $x$ , such that  $x' = x \vee r_j$ . We repeat this procedure until the updated  $x'$  is classified as *goodware* (for low-confidence attacks) or until an attacker-defined confidence level is achieved (for high-confidence attacks). Finally, we inject all the candidate gadgets at once via automated software transplantation, and verify that problem-space constraints are met and that the app is still classified as *goodware*. Algorithm 2 in Appendix F outlines the detailed steps of the attack phase.

## 4 EXPERIMENTAL EVALUATION

We evaluate the effectiveness of our novel problem-space Android attack, in terms of success rate and required time—and also when in the presence of feature-space defenses.

### 4.1 Experimental Settings

**Prototype.** We create a prototype of our novel problem-space attack (§3) using a combination of Python for the ML functionality and Java for the program analysis operations; in particular, to perform transplantations in the problem-space we rely on FlowDroid [22], which is based on Soot [12]. We release the code of our prototype to other academic researchers (see §8). We ran all experiments on an Ubuntu VM with 36 vCPUs, 252GB of RAM, and NVIDIA Tesla V100 GPU.

**Classifiers.** As defined in the threat model (§3.1), we consider the DREBIN classifier [21], based on a binary feature space and a linear SVM, and its recently proposed hardened variant, Sec-SVM [39], which requires the attacker to modify more features to perform an evasion. This last model is implemented in PyTorch. In this work, we apply a feature selection selecting the 10,000 most important features out of the  $\approx 1.5$  million features of the original representation, maintaining the same clean performance. We use hyperparameter  $C=1$  for the linear SVM as in [21] with the dual problem formulation, and identify the optimal Sec-SVM parameter  $k = 0.25$  (i.e., the maximum feature weight) in our setting by enforcing a maximum performance loss of 2% AUC.

**Attack Confidence.** We consider two attack settings: *low-confidence* (L) and *high-confidence* (H). The (L) attack merely overcomes the decision boundary (so that  $h(x) < 0$ ). The (H) attack maximizes the distance from the hyperplane into the *goodware* region; while generally this distance is unconstrained, here we set it to be  $\leq$  the negative scores of 25% of the benign apps (i.e., within their interquartile range). The percentage interval has been chosen to force the adversarial point in the distribution of *goodware*. If we were to enforce the adversarial points to be more negative than 100% of the *goodware*, an anomaly detection method would be able to detect our adversarial points; instead, by enforcing the first quartile we aim to be more inconspicuous. This avoids making superfluous modifications, which may only increase suspiciousness or the chance of transplantation errors, while being closer in nature to past mimicry attacks [18].

**Dataset.** We collect apps from AndroZoo [30], a large-scale dataset with timestamped Android apps crawled from different stores, and with VirusTotal summary reports. We use the same labeling criteria as Tesseract [69] (which is derived from [32]): an app is considered *goodware* if it has 0 VirusTotal detections, as *malware* if it has 4+ VirusTotal detections, and is discarded as *grayware* if it has between 1 and 3 VirusTotal detections. For the dataset composition, we follow the example of Tesseract and use an average of 10% malware [69]. The final dataset contains  $\approx 150K$  recent Android applications, dated between Jan 2016 and Dec 2018, specifically 135,708 *goodware* and 15,765 *malware*.




Fig. 2. Performance of SVM and Sec-SVM in absence of adversarial attacks.

**Dataset Split.** Tesseract [69] demonstrated that, in non-stationary contexts such as Android malware, if time-aware splits are not considered, then the results may be inflated due to *concept drift* (i.e., changes in the data distribution). However, here we aim to specifically evaluate the effectiveness of an adversarial attack. Although it likely exists, the relationship between adversarial and concept drift is still unknown and is outside the scope of this work. If we were to perform a time-aware split, it would be impossible to determine whether the success rate of our ML-driven adversarial attack was due to an intrinsic weakness of the classifier or due to natural evolution of malware (i.e., the introduction of new non-ML techniques malware developers rely on to evade detection). Hence, we perform a *random split* of the dataset to simulate *absence of concept drift* [69]; this also represents the most challenging scenario for an attacker, as they aim to mutate a test object coming from the same distribution as the training dataset (on which the classifier likely has higher confidence). In particular, we consider a 66% training and 34% testing random split.<sup>2</sup>

**Testing.** The test set contains a total of 5,203 malware. The statistics reported in the remainder of this section refer only to *true positive* malware (3,897 for SVM and 3,762 for Sec-SVM), i.e., we create adversarial variants only if the app is detected as malware by the classifier under evaluation. Intuitively, it is not necessary to make an adversarial example of a malware application that is already misclassified as goodware; hence, we avoid inflating results by removing false negative objects from the dataset. For the evaluation we apply a restriction on the harvested gadgets, selecting the stealthiest ones: we are indeed considering only the gadgets with maximum 20 classes in total and maximum one permission. During the transplantation phase of our problem-space attack some errors occur due to bugs and corner-case errors in the FlowDroid framework [22]. Since these errors are related on implementation limitations of the FlowDroid research prototype, and not conceptual errors, the success rates in the remainder of this section refer only to applications that did not throw FlowDroid exceptions during the transplantation phase.

## 4.2 Evaluation

We analyze the performance of our Android problem-space attack in terms of runtime cost and successful evasion rate. An attack is successful if an app  $z$ , originally classified as malware, is mutated into an app  $z'$  that is classified as goodware and satisfies the problem-space constraints.

Figure 2 reports the AUROC of SVM and Sec-SVM on the DREBIN feature space in absence of attacks. As expected [39], Sec-SVM sacrifices some detection performance in return for greater feature-space adversarial robustness.

<sup>2</sup>We consider only one split due to the overall time required to run the experiments. Including some prototype overhead, it requires about one month to run all configurations.

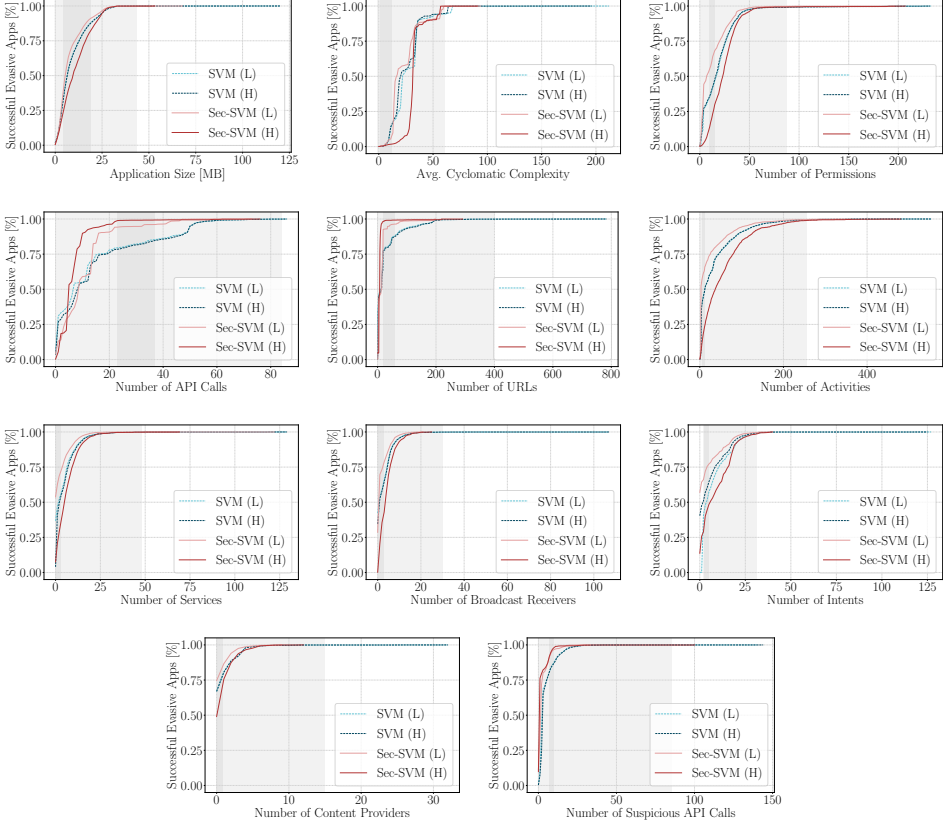


Fig. 3. Statistics of the evasive malware variants, compared with statistics of benign apps. The dark gray background highlights the area between first and third quartile of benign applications; the light gray background is based on the  $3\sigma$  rule and highlights values benign statistics between 0.15% and 99.85% of the distribution (i.e., spanning 99.7% of the distribution).

**Attack Success Rate.** We perform our attack using *true positive* malware from the test set, i.e., all malware objects correctly classified as malware. We consider four settings depending on the considered model and the attack confidence: SVM (L), SVM (H), Sec-SVM (L), and Sec-SVM (H). In absence of FlowDroid exceptions (see the conference version for more details), we are able to create an evasive variant for each malware in all four configurations. In other words, we achieve a misclassification rate of 100.0% on the successfully generated apps, where the problem-space constraints are satisfied by construction (as defined in §3). Figure 5 reports the cumulative distribution of features added when generating evasive apps for the four different configurations. As expected, Sec-SVM requires the attacker to modify more features, but here we are no longer interested in the feature-space properties, since we are performing a problem-space attack. This demonstrates that measuring attacker effort with  $l_p$  perturbations as in the original Sec-SVM evaluation [39] *overestimates* the robustness of the defense and is better assessed using our framework (§2). Interesting to note anyway, that compared to our previous version of the attack, the number of needed features has highly decrease due to the improved extraction capabilities of the new implementation.



While the plausibility problem-space constraint is satisfied by design by transplanting only realistic existing code, it is informative to analyze how the statistics of the evasive malware relate to the corresponding distributions in benign apps. Figure 3 reports the cumulative distribution of app statistics across the four settings: the  $X$ -axis reports the statistics values, whereas the  $Y$ -axis reports the cumulative percentage of evasive malware apps. We also shade two gray areas: a *dark gray area* between the first quartile  $q_1$  and third quartile  $q_3$  of the statistics for the benign applications; the *light gray area* refers to the  $3\sigma$  rule and reports the area within the 0.15% and 99.85% of the benign apps distribution.

Figure 3 shows that while evading Sec-SVM tends to cause a shift towards the higher percentiles of each statistic, the vast majority of apps falls within the gray regions in all configurations. However, thanks to the enhanced extraction capabilities of the framework this is less evident compared to our initial work, making the current attack even more efficient and powerful, especially against Sec-SVM. We note that this is just a qualitative analysis to verify that the statistics of the evasive apps roughly align with those of benign apps; it is not sufficient to have an anomaly in one of these statistics to determine that an app is malicious (otherwise, very trivial rules could be used for malware detection itself, and this is not the case). We also observe that there is little difference between the statistics generated by Sec-SVM and by traditional SVM; this means that greater feature-space perturbations do not necessarily correspond to greater perturbations in the problem-space, reinforcing the feasibility and practicality of evading Sec-SVM.

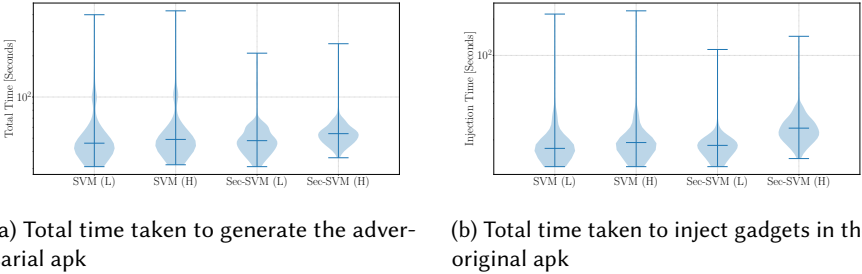


Fig. 4. Violin plots of times per adversarial app.

**Runtime Overhead.** The time to perform the search strategy occurring in the feature space is almost negligible; the most demanding operation is in the actual code modification. Figure 4 depicts the distribution of injection times for our test set malware which is the most expensive operation in our approach while the rest is mostly pipeline overhead. The time spent per app is low: in most cases, less than 100 seconds, and always less than 2,000 seconds ( $\sim 33$  mins). The low runtime cost suggests that it is feasible to perform this attack at scale and reinforces the need for new defenses in this domain.

### 4.3 Discussion

Through our evaluation we have shown that it is possible to create realistic adversarial samples and depending on the amount of available gadgets to implant the final outcome and difficulty varies. Indeed, with the new set of gadgets, which has far more comprehensive, we show that the average amount of injected features has been significantly reduced.

## 5 MODEL HARDENING

A further step we have done in this paper is to explore how feasible is adversarial training in order to harden the target classifiers against problem space generated adversarial samples. Our objective




Fig. 5. Breakdown of average number of features injected for each considered classifier.

is to assess the efficacy of various hardening methods in response to the attack strategy we have outlined, as referenced in [algorithm 2](#). In this scenario, the attacker faces no restrictions on the number of gadgets they can employ, with their only limitation being the total quantity of gadgets available for harvest. We show that different sets of hardening methodologies may lead to different robustness levels against the attack. We include experiments both focusing on linear SVM and on Sec-SVM[39]. To have full control of the training procedure, we approximate also the linear SVM as a *single-layer* neural network (NN) using PyTorch [45]. This enables us to evaluate the effectiveness of both hardening strategies (i.e., adversarial training and adversarial retraining) in this setting.

### 5.1 Hardening Strategies

As defenders, our goal is to explore and use different approaches in order to harden our considered classifier. In order to do so, we are considering two different approaches: adversarial training and adversarial retraining. Adversarial training [25] involves the creation of adversarial examples which are integrated into the model's training phase, where they are blended with the standard dataset throughout the training periods. The most widely accepted technique [43] involves consistently generating adversarial examples that result in the maximum possible loss. Subsequently, the model's adjustable parameters are fine-tuned to reduce the error rate when classifying these challenging adversarial examples. Then follows the unchanged equation and its explanation:

$$\min_{\theta} \mathbb{E}(x, y) \sim D [\max_{\delta \in S} L(\theta, x + \delta, y)] \quad (13)$$

where  $(x, y) \sim D$  represents training data sampled from the distribution  $D$ ,  $\delta \in S$  is the maximum allowed perturbation,  $L$  is the model loss and  $\theta$  is the model hyperparameters.

On the other hand, adversarial retraining as hardening strategy [79] consists in using attack  $A$  over the training dataset  $D$  to produce an adversarial training set  $D_a = \{(\hat{x}_i, \hat{y}_i)\}_{i=1}^N$ . We join the original and the adversarial training sets and retrain the classifier  $H(x)$  over this augmented set  $D^* = D \cup D_a$ . It is to be noted that this is different from classical adversarial training where new adversarial examples are generated and mixed with the original training set continuously. We have chosen to consider this strategy because it is suited for both traditional machine learning algorithms and neural networks, so can be equally applied to both our considered models.

We have hypothesised two different types of attacks that defender could use in order to create adversarial samples for hardening: a feature-space attack, and our proposed problem-space attack in [algorithm 2](#). The feature space attack is a variant of the latter in which we inject the most important features of the target classifier without considering  $\eta$ .

We consider and compare the following aspects:

- The amount of malware samples used during hardening
- The hardening technique adopted
- The transformation type for creating adversarial samples (PS vs. FS)
- The considered level of confidence both for evasion and for the hardening phase

More details about the considered settings can be found at [Appendix E](#).




Fig. 6. Percentage of feature space required to evade various configurations of the hardened classifier.

## 5.2 Evaluation

For this evaluation we assume that the defender is using an increasing percentage of the available malware in the training set (cf. §5.1), and we focus on the two classification models previously described, a LinearSVM and SecSVM. As we are focused on exploring scenarios in which attacker and defender has similar capabilities, we consider the following setting Our goal is to evaluate the robustness against our proposed white-box problem-space attack through the adversarial training [43] or adversarial retraining [79], comparing the impact of using realistic problem-space samples against unrealistic feature-space samples. We consider two different levels of confidence of the adversarial samples, as explained in §4.1, which can be set as a target both during hardening and evaluation. For example, a Low-Low configuration means that both during the hardening and during evaluation the target for the generated adversarial samples is low confidence.

Despite conducting a comprehensive analysis of all potential combinations in Figure 6, in this section we specifically report scenarios where attackers and defenders possess comparable capabilities, to investigate if it is feasible to develop hardened models that can effectively withstand sophisticated attacks from adversaries with the same capabilities as the defenders.

We omit visualizations of adversarial retraining results as our experiments consistently show this technique fails to enhance model resilience against the studied threat actors. Throughout our testing, attackers successfully generated high-confidence adversarial samples against models hardened with this approach, with ASR consistently exceeding 95%. These results conclusively demonstrate the ineffectiveness of this particular hardening strategy against the sophisticated adversaries examined in our research.

A notable observation across all our experimental evaluations is the uniformity of outcomes, regardless of the variations in experimental setups or models tested. This striking consistency highlights a fundamental limitation of adversarial retraining in the considered context. Indeed, its current form when does not defend against the generation of both low and high confidence adversarial samples by the considered threat actors.

On the other hand, hardening through adversarial training methodology has been proven more effective, especially on LinearSVM. Indeed, Figure 7 and Figure 8 show that adversarial training is able to significantly reduce the ASR against the generation of high confidence adversarial samples,

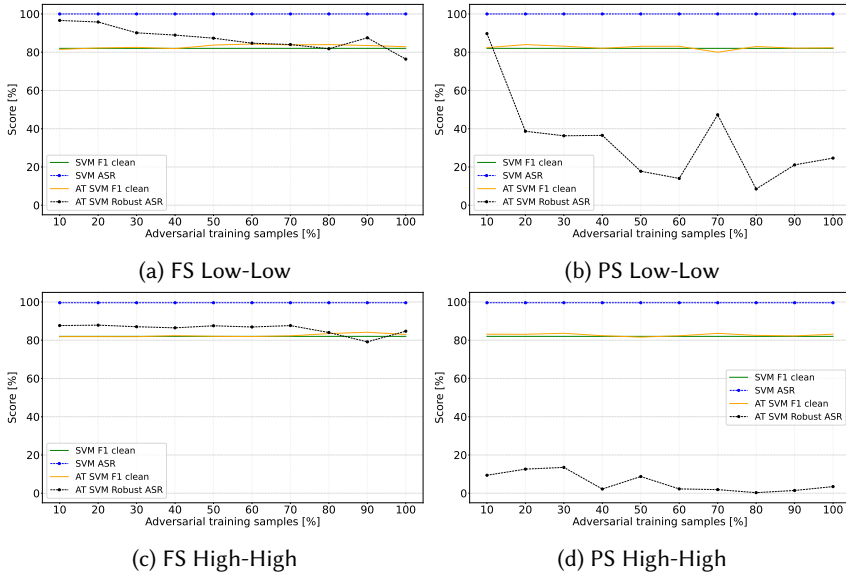


Fig. 7. Evaluating the effectiveness of both the original and the adversarially trained LinearSVM model against adversarial examples in the problem space, as well as assessing performance on clean data. The charts illustrate the varied performance outcomes resulting from different hardening approaches. "PS" indicates that the model underwent hardening through the use of adversarial samples in the Problem Space, whereas "FS" signifies hardening via adversarial samples in the Feature Space.

while against low confidence adversarial samples we can enforce some level of robustness but always higher than 20%. We can notice a huge difference between models trained using only feature space transformations and problem space: models hardened with feature space transformations even in the High-High scenario cannot reach a reasonable level of robustness while the ones hardened with problem space transformations can.

All the considered evasions are unbound, meaning that the attacker can use all the gadgets they were able to harvest, without any need of minimising the perturbation. Specifically for this work, we have been able to extract around  $\approx 800$  suitable gadgets from goodware. Looking at the number of features used in order to evade Figure 6, we can notice that in the majority of cases the model hardened using problem space samples significantly raise the bar in the amount of features (and gadgets) needed for evasion, especially for SVM.

### 5.3 Discussion

Our investigations underscore that adversarial retraining fails to adequately harden the model in question within the experimental frameworks we explored. This deficiency in strengthening is evident through the consistent and successful breaches by adversaries against models purportedly made more robust through adversarial retraining. The fundamental principle underlying this process is that each iteration of adversarial retraining explores a specific vulnerable region. However, this approach is inherently non-exhaustive, as it concentrates primarily on the considered decision boundary, which undergoes transformation subsequent to each retraining phase. This transformation exposes different blind spots within the model's decision-making process. Given that the feature space under consideration is characteristically high-dimensional and sparse, there exists a significant probability that this mechanism will not achieve convergence, or if convergence is attained, it may necessitate an extensive number of iterations. This is due to the vast and complex

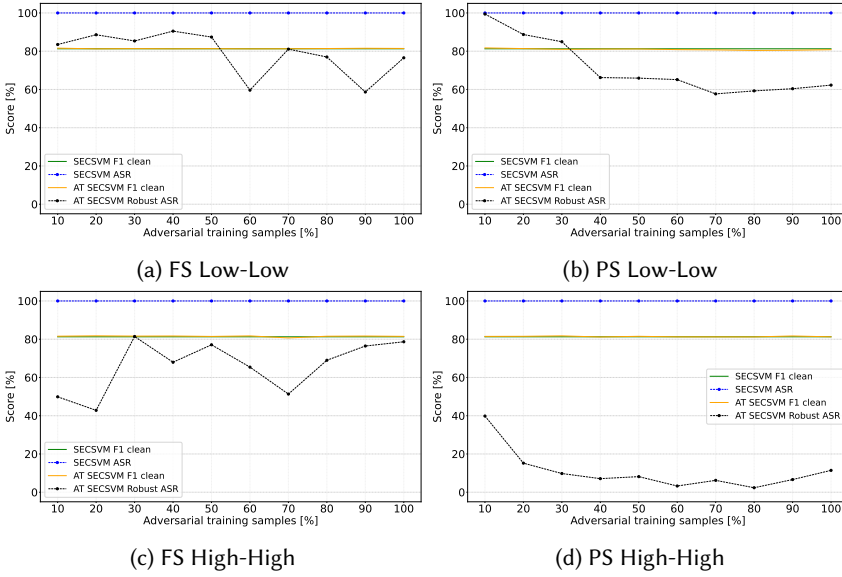


Fig. 8. Evaluating the effectiveness of both the original and the adversarially trained Sec-SVM model against adversarial examples in the problem space, as well as assessing performance on clean data. The charts illustrate the varied performance outcomes resulting from different hardening approaches. "PS" indicates that the model underwent hardening through the use of adversarial samples in the Problem Space, whereas "FS" indicates hardening via adversarial samples in the Feature Space.

nature of the feature space, which allows for numerous potential vulnerabilities to emerge as the decision boundary evolves. The iterative nature of this process, combined with the dynamic shifts in the decision boundary, creates a scenario where each reinforcement potentially unveils new areas of susceptibility. This phenomenon underscores the challenges inherent in achieving comprehensive robustness in high-dimensional machine learning models through adversarial retraining methods.

On the flip side, our research shows the beneficial impact of adversarial training on enhancing model robustness. Even if using feature space transformations for hardening brings some minor level of robustness, we discovered that adversarial samples created through problem-space transformations substantially increase the resilience of the target classifier. These samples prove to be more effective than feature spaced transformations, offering a comparative advantage for both classifiers subjected to our tests.

A noteworthy observation from our study is the superior potential for improvement exhibited by the SVM classifier over the Sec-SVM when subjected to adversarial training in though problem-space hardening. This difference in potential improvement is likely rooted in the inherent differences between their learning algorithms. The SVM classifier, tends to concentrate on a select set of essential features [55], while in contrast, the Sec-SVM employs a strategy of weight clamping that aims to distribute significance more uniformly across a broader spectrum of features by imposing upper and lower limits on their influence. While intended to foster a more balanced approach to feature importance, this strategy could inadvertently set a cap on the positive impacts of adversarial training by too strictly regulating the weights.

A consistent and significant theme emerging from our research is the distinct advantage offered by adversarial samples that are crafted through realistic transformations in the problem space, especially in terms of hardening models. This advantage manifests not only in the broadening of the total amount of budget needed in order to craft successful adversarial samples but also in the notable



enhancement of model robustness. Practically, we observed improvements in robustness ranging between 10 to 60% in classifiers that underwent hardening with realistic adversarial samples. This translates into an increased evasion budget requirement of 3 to 15%, indicating a substantial uplift in model defense capabilities.

Additionally, our analysis reveals the absence of uniform patterns in the ASR metrics, as evidenced in previous plots [Figure 8](#) and [Figure 7](#). This phenomenon may be attributed to the random selection process employed for each batch, where points for transformation are randomly chosen from a specified subset at a predetermined rate. This randomness underscores the critical importance of selecting representative samples for transformation, suggesting that the strategic alteration of the model's weights through these samples is more significant than the sheer number of points utilized. This insight serves to highlight the complex and nuanced nature of adversarial training, shedding light on the myriad of factors that play a crucial role in determining its effectiveness and success.

The impact of problem-space hardening on feature utilization for evasion is noteworthy, as illustrated in [Figure 6](#). This technique significantly increases the number of features required for successful evasion, reaching nearly 20% in the most robust scenarios. Interestingly, the Low-Low Problem Space hardening consistently demonstrates a higher threshold compared to other scenarios when testing for Low confidence evasions. This suggests that during the hardening phase, the process effectively targets the appropriate manifold area due to the perfect symmetry between attacker and defender strategies. In contrast, the High-Low (HL) scenario exhibits a lower feature requirement. This can be attributed to the fact that hardening against high confidence adversarial examples may lead to the exploration of different representation areas, which prove less effective in defending against attacks specifically targeting the Low confidence evasion domain.

## 6 DISCUSSION ON ATTACK AND DEFENSE RESULTS

We provide some deeper discussion on the results of our novel problem-space attack and the investigated hardening methods.

**Android Attack Effectiveness.** We conclude that it is practically *feasible* to evade the state-of-the-art Android malware classifier DREBIN [21] and its hardened variant, Sec-SVM [39], and that we are able to automatically generate realistic and inconspicuous evasive adversarial applications, often in less than 2 minutes. This shows for the first time that it is possible to create realistic adversarial applications at scale.

**Obfuscation.** It could be argued that traditional obfuscation methods can be used to simply hide malicious functionality. The novel problem-space attack in this work evaluates the feasibility of an “adversarial-malware as a service” scenario, where the use of mass obfuscation may raise the suspicions of the defender; for example, antivirus companies often classify samples as malicious simply because they utilize obfuscation or packing [59, 28]. Moreover, some other analysis methods combine static and dynamic analysis to prioritize evaluation of code areas that are likely obfuscated [e.g., 42]. On the contrary, our transformations aim to be fully inconspicuous by adding only legitimate benign code and, to the best of our knowledge, we do not leave any relevant artifact in the process. While the effect on problem-space constraints may differ depending on the setting, attack methodologies such as ours and traditional obfuscation techniques naturally complement each other in aiding evasion and, in the program domain, code transplantation may be seen as a tool for developing new forms of inconspicuous obfuscation [63].

**Other Defense Directions.** A promising direction by [51] studies the use of *monotonic classifiers*, where adding features can only increase the decision score (i.e., an attacker cannot rely on adding more features to evade detection); however, such classifiers require non-negligible time towards manual feature selection (i.e., on features that are harder for an attacker to change), and—at least in the context of Windows malware [51]—they suffer from high false positives and an average

reduction in detection rate of 13%. Moreover, we remark that we decide to add goodwill parts to malware for practical reasons: the opposite transplantation would be immediate to do if a dataset with *annotated* malicious bytecode segments were available. As part of future work we aim to investigate whether it would still be possible to evade monotonic classifiers by adding only a minimal number of malicious slices to a benign application.

**Defenses Against Problem-Space Attacks.** Unlike settings where feature and problem space are closely related (e.g., images and audio), limitations on feature-space  $l_p$  perturbations are often insufficient to determine the risk and feasibility of an attack in the real world. Our novel problem-space formalization (§2) paves the way to the study of *practical* defenses that can be effective in settings which lack an inverse feature mapping. Simulating and evaluating attacker capabilities in the problem space helps define realistic threat models with more constrained modifications in the feature space—which may lead to more robust classifier design. Our Android evasion attack (§3) demonstrates for the first time that it is *feasible* to evade feature-space defenses such as Sec-SVM in the problem-space—and to do so *en masse*.

**Model hardening through adversarial training.** Our experiments demonstrate that depending on the considered strategies the result may highly vary. Adversarial retraining did not succeed in fortifying the classifier in any capacity. Conversely, the effect of adversarial training led to a reduction in the Attack Success Rate (ASR) by up to 80% when facing the creation of low-confidence adversarial samples, and by 90% for high-confidence ones. In our settings, problem space hardening revealed to be more effective in all considered scenario if compared to feature-space transformation adversarial samples, highlighting the exploration of more meaningful areas of the feature space. The difference between the two considered models' architectures suggest that the success of adversarial training is closely tied to the training methodology, the considered feature space and the classification model used, shown also by the opposite conclusions of [82] and [84]. Future research could delve into various classifiers and feature spaces to uncover the underlying relationship between representation and learning logic, aiming to boost the effectiveness of adversarial training. Another key aspect is related to the points selected during the hardening process: evidence shows that using all the available points does not always achieve the best robustness.

## 7 RELATED WORK

**Android Malware Detection** Recent research in Android malware detection has explored various deep learning approaches, each with unique characteristics and contributions to the field.

McLaughlin et al. [44] pioneered the use of convolutional neural networks (CNNs) for Android malware detection. Their approach operates directly on raw opcode sequences extracted from disassembled apps, eliminating the need for manual feature engineering. This method allows the model to automatically learn relevant features from low-level code representations, potentially capturing subtle patterns that might be overlooked in traditional feature-based approaches.

In contrast, Karbab et al. [52] developed MalDozer, which also uses CNNs but focuses on API call sequences as input. This higher-level representation captures the behavioral aspects of the app, as API calls directly reflect the app's interactions with the Android system. By using API sequences, MalDozer can potentially detect malware that employs code obfuscation techniques to evade lower-level analysis.

Xu et al. [60] proposed DeepRefiner, introducing a multi-layer approach that combines both static and dynamic analysis. Their system uses a shallow neural network for initial efficient filtering, followed by a deeper Long Short-Term Memory (LSTM) network for precise detection. This hybrid approach aims to balance computational efficiency with detection accuracy, addressing the scalability challenges often faced in malware detection systems.

Addressing this challenge, Wang et al. [74] proposed a hybrid model combining deep autoencoders and CNNs. The autoencoder component aims to learn a robust feature representation, potentially making the model more resilient to minor perturbations that might fool simpler models. This approach represents a step towards developing adversarially robust malware detection systems.

Oak et al. [68] took a different approach by exploring the use of pre-trained language models like BERT for malware detection. This method treats code sequences as a form of language, leveraging the power of large-scale pre-training to capture complex patterns in code. Their work also focused on addressing the challenge of imbalanced datasets, which is common in real-world malware detection scenarios where benign apps significantly outnumber malicious ones.

**Adversarial Machine Learning.** Adversarial ML attacks have been studied for more than a decade [49]. These attacks aim to modify objects either at training time (*poisoning* [58]) or at test time (*evasion* [18]) to compromise the confidentiality, integrity, or availability of a machine learning model. Many formalizations have been proposed in the literature to describe feature-space attacks, either as optimization problems [18, 36] (see also §2.1 for details) or game theoretic frameworks [5]. Recent work propose attacks under black-box assumptions, relying on random sampling with finite-difference strategies [86] or random search [85]

**Problem-Space Attacks.** Recently, research on adversarial ML has moved towards domains in which the feature mapping is not invertible or not differentiable. Here, the adversary needs to modify the objects in the problem space (i.e., input space) without knowing exactly how this will affect the feature space. This is known as the *inverse feature-mapping* problem [13, 18, 70]. Many works on problem-space attacks have been explored on different domains: text [65, 47], PDFs [20, 16, 15, 38, 35], Windows binaries [53, 56, 57, 72], Android apps [39, 41, 46], NIDS [8, 48, 61, 19], ICS [77], and Javascript source code [70]. Some work have also used reinforcement learning techniques to perform transformations [81], declaring an action space of problem-space compliant transformations. However, each of these studies has been conducted empirically and followed some inferred best practices: while they share many commonalities, it has been unclear how to compare them and what are the most relevant characteristics that should be taken into account while designing such attacks. Our formalization (§2) aims to close this gap, and we show how it can be used to describe representative feature-space and problem-space attacks from the literature (§2.3).

**Adversarial Android Malware.** This paper also proposes a novel adversarial problem-space attack in the Android domain (§3); our attack overcomes limitations of existing proposals, which are evidenced through our formalization. The most related approaches to our novel attack are on attribution [70], and on adversarial malware generation [46, 57, 41]. [70] do *not* consider malware detection, but design a set of simple mutations to change the programming style of an application to match the style of a target developer (e.g., replacing *for* loops with *while* loops). This strategy is effective for attribution, but is insufficient for malware detection as altering stylometric properties alone would not evade a malware classifier which captures program semantics. Moreover, it is not feasible to define a hardcoded set of transformations for all possible semantics—which may also leave artifacts in the mutated code. Conversely, our attack uses automated software transplantation to ensure the plausibility of the generated code and avoid hardcoded mutation artifacts.

Grosse et al. [41] perform minimal modifications that preserve semantics, and only modify single lines of code in the Manifest; but these may be easily detected and removed due to unused permissions or undeclared classes. Moreover, they limit their perturbation to 20 features, whereas our problem-space constraints represent a more realistic threat model.

[46] propose a method for adversarial Android malware generation. Similarly to us, they rely on *automated software transplantation* [26] and evaluate their adversarial attack against the DREBIN classifier [21]. However, they do not formally define which semantics are preserved by their

transformation, and their approach is extremely unstable, breaking the majority of apps they mutate (e.g., they report failures after 10+ modifications on average—which means they would likely not be able to evade Sec-SVM [39] which on average requires modifications of 50+ features). Moreover, the code is unavailable, and the paper lacks details required for reevaluating the approach, including any clear descriptions of preprocessing robustness. Conversely, our attack is resilient to the insertion of a large number of features (§4), preserves dynamic app semantics through opaque predicates (§3.3), and is resilient against static program analysis (§3.4).

[57] propose a black-box adversarial attack against Windows malware classifiers that rely on API sequence call analysis—an evasion strategy that is also applicable to similar Android classifiers. In addition to the limited focus on API-based sequence features, their problem-space transformation leaves two major artifacts which could be detected through program analysis: the addition of no-operation instructions (*no-ops*), and patching of the import address table (IAT). Firstly, the inserted API calls need to be executed at runtime and so contain individual no-ops hardcoded by the authors following a practice of “security by obscurity”, which is known to be ineffective [1, 62]; intuitively, they could be detected and removed by identifying the tricks used by attackers to perform no-op API calls (e.g., reading 0 bytes), or by filtering the “dead” API calls (i.e., which did not perform any real task) from the dynamic execution sequence before feeding it to the classifier. Secondly, to avoid requiring access to the source code, the new API calls are inserted and called using IAT patching. However, all of the new APIs must be included in a separate segment of the binary and, as IAT patching is a known malicious strategy used by malware authors [40], IAT calls to non-standard dynamic linkers or multiple jumps from the IAT to an internal segment of the binary would immediately be identified as suspicious. Conversely, our attack does not require hardcoding and by design is resilient against traditional non-ML program analysis techniques.

**Defenses against adversarial samples** A different approach from SecSVM which aims to create more robust models can be identified by the strategy described in [73], in which authors leverage on “conserved features” (those that cannot be unilaterally modified without compromising malicious functionality), in order to alter the used feature space. However the considered domain is different, they do evaluation on malicious PDF files, which implies that also the considered feature representation is really different from the one we are considering. While PDF feature space includes also structural features, Drebin feature space lack of such and so we could think as all features to be “conserved features”. For this reason, such reasoning is hardly applicable to the evaluated scenario of this paper.

**Impact of realistic adversarial samples on adversarial training** We are not the first researchers to investigate such research field, e.g. [82] [84]. Both work explore the impact of feature space generated adversarial samples compared to problem space ones in the malware domain and end to a different conclusion. Focusing on Adversarial Training, [82] make a wide scope research evaluating the different impact on adversarial training of realistic and unrealistic adversarial training samples on random forest classifiers in multiple domains: text classification, malicious traffic and windows malware domain. They evaluate different attack strategies and make an evaluation of the hardened model and verify if the adversarial training granted robustness against realistic attacks. They claim that unrealistic samples do not suffice to harden the classifier in the malware and text classification domain, while they are enough in the malicious traffic domain. [84] research the same problem focusing entirely in the windows malware domain training a DNN on raw-binary feature representation. They consider multiple state-of-the-art problem-space attacks - In-Place Replacement (IPR), Displacement (Disp) and Kreuk - and they conclude that it is not necessary to train on the highest-effort attacks to still gain substantial robustness against problem-space attacks, implying that feature space adversarial samples are enough in order to provide robustness. They also claim that early adversarial training harmed natural TPR, and training with more batches

is important to recover accuracy on the original data. This contrasting results can be explained analysing the differences between the two approaches. Indeed, they consider a different classifier and a different feature representation, which underlines the importance of the considered model and the representation used in order to evaluate the impact of adversarial training.

## 8 AVAILABILITY

We release the code and data of our approach to other researchers by responsibly sharing a private repository. The project website with instructions to request access is at: <https://s2lab.cs.ucl.ac.uk/projects/intriguing/>.

## 9 CONCLUSIONS

Since the seminal work that evidenced intriguing properties of neural networks [25], the community has become more widely aware of the brittleness of machine learning in adversarial settings [49].

To better understand real-world implications across different application domains, we propose a novel formalization of problem-space attacks as we know them today, that enables comparison between different proposals and lays the foundation for more principled designs in subsequent work. We uncover new relationships between feature space and problem space, and provide necessary and sufficient conditions for the existence of problem-space attacks. Our novel problem-space attack shows that automated generation of adversarial malware at scale is a realistic threat—taking on average less than 2 minutes to mutate a given malware example into a variant that can evade a hardened state-of-the-art classifier.

We also try to understand the effectiveness of adversarial learning and the impact of leveraging on realistic adversarial samples. The gathered evidence underline the benefits of the usage of problem-space adversarial samples for adversarial training for the considered models, highlighting that the impact may vary depending on the considered representation, model and adopted methodology.

Future work will explore this reasoning operating in different feature spaces, training different models and considering different feature space attacks, pursuing the underline relation between training and the representation taken into account.

## ACKNOWLEDGMENTS

This research was partially supported by: the UK EPSRC Grant EP/X015971/2; Google ASPIRE Awards; ); the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2092 CASA – 390781972; the IFI program of the German Academic Exchange Service (DAAD); and the Vienna Science and Technology Fund (WWTF) through the BREADS project (10.47379/VRG23011).

## REFERENCES

- [1] Auguste Kerckhoffs. “La cryptographie militaire”. In: *Journal des sciences militaires*. 1883.
- [2] Mark Weiser. “Program Slicing”. In: *Proceedings of the 5th International Conference on Software Engineering*. ICSE '81. San Diego, California, USA: IEEE Press, 1981, pp. 439–449. URL: <http://dl.acm.org/citation.cfm?id=800078.802557>.
- [3] Bart Selman, David G. Mitchell, and Hector J. Levesque. “Generating Hard Satisfiability Problems”. In: *Artif. Intell.* 81.1-2 (1996), pp. 17–29. DOI: [10.1016/0004-3702\(95\)00045-3](https://doi.org/10.1016/0004-3702(95)00045-3). URL: [https://doi.org/10.1016/0004-3702\(95\)00045-3](https://doi.org/10.1016/0004-3702(95)00045-3).
- [4] Benjamin C Pierce and C Benjamin. *Types and programming languages*. MIT press, 2002.
- [5] Nilesch Dalvi et al. “Adversarial classification”. In: *KDD*. ACM. 2004.
- [6] Daniel Lowd and Christopher Meek. “Good Word Attacks on Statistical Spam Filters.” In: *CEAS*. Vol. 2005. 2005.
- [7] Christopher M Bishop. *Pattern Recognition and Machine Learning*. 2006.
- [8] Prahlaad Fogla and Wenke Lee. “Evading network anomaly detection systems: formal reasoning and practical techniques”. In: *ACM Conference on Computer and Communications Security*. ACM, 2006, pp. 59–68.
- [9] Alfred V Aho, Ravi Sethi, and Jeffrey D Ullman. *Compilers, Principles, Techniques, and Tools (2nd Edition)*. Addison Wesley, 2007.



- [10] Engin Kirda Andreas Moser Christopher Kruegel. "Limits of Static Analysis for Malware Detection". In: (2007).
- [11] Andreas Moser, Christopher Kruegel, and Engin Kirda. "Limits of Static Analysis for Malware Detection". In: ACSAC. 2007.
- [12] Raja Vallée-Rai et al. "Soot: A Java bytecode optimization framework". In: *CASCON First Decade High Impact Papers*. IBM Corp. 2010.
- [13] Ling Huang et al. "Adversarial machine learning". In: *AISeC*. ACM. 2011.
- [14] Ling Huang et al. "Adversarial machine learning". In: *Proceedings of the 4th ACM workshop on Security and artificial intelligence*. ACM. 2011, pp. 43–58.
- [15] Pavel Laskov and Nedim Šrnić. "Static Detection of Malicious JavaScript-Bearing PDF Documents". In: ACSAC. ACM. 2011.
- [16] Davide Maiorca, Giorgio Giacinto, and Iginio Corona. "A Pattern Recognition System for Malicious PDF Files Detection". In: *Intl. Workshop on Machine Learning and Data Mining in Pattern Recognition*. Springer. 2012.
- [17] Battista Biggio, Giorgio Fumera, and Fabio Roli. "Security evaluation of pattern classifiers under attack". In: *IEEE TKDE* (2013).
- [18] Battista Biggio et al. "Evasion attacks against machine learning at test time". In: *ECML-PKDD*. Springer. 2013.
- [19] Iginio Corona, Giorgio Giacinto, and Fabio Roli. "Adversarial attacks against intrusion detection systems: Taxonomy, solutions and open issues". In: *Information Sciences* (2013).
- [20] Davide Maiorca, Iginio Corona, and Giorgio Giacinto. "Looking at the bag is not enough to find the bomb: an evasion of structural methods for malicious pdf files detection". In: *ASIACCS*. ACM. 2013.
- [21] Daniel Arp et al. "DREBIN: Effective and Explainable Detection of Android Malware in Your Pocket". In: *NDSS*. 2014.
- [22] Steven Arzt et al. "FlowDroid: precise context, flow, field, object-sensitive and lifecycle-aware taint analysis for Android apps". In: *PLDI*. ACM, 2014.
- [23] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. "Explaining and harnessing adversarial examples". In: *arXiv preprint arXiv:1412.6572* (2014).
- [24] Pavel Laskov et al. "Practical evasion of a learning-based classifier: A case study". In: *IEEE Symposium on Security and Privacy (S&P)*. IEEE. 2014, pp. 197–211.
- [25] Christian Szegedy et al. "Intriguing properties of neural networks". In: *ICLR* (2014).
- [26] Earl T Barr et al. "Automated Software Transplantation". In: *ISSTA*. ACM. 2015.
- [27] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. "Explaining and Harnessing Adversarial Examples". In: *ICLR (Poster)*. 2015.
- [28] Xabier Ugarte-Pedrero et al. "SoK: Deep packer inspection: A longitudinal study of the complexity of run-time packers". In: *IEEE Symposium on Security and Privacy*. 2015.
- [29] William Weiss and Cherie D'Mello. *Fundamentals of Model Theory*. University of Toronto, 2015.
- [30] Kevin Allix et al. "Androzoo: Collecting Millions of Android Apps for the Research Community". In: *ACM Mining Software Repositories (MSR)*. 2016.
- [31] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT press, 2016.
- [32] Brad Miller et al. "Reviewer Integration and Performance Measurement for Malware Detection". In: *DIMVA*. Springer. 2016.
- [33] Nicolas Papernot et al. "The limitations of deep learning in adversarial settings". In: *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE. 2016, pp. 372–387.
- [34] Mahmood Sharif et al. "Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition". In: *ACM CCS*. ACM. 2016.
- [35] Weilin Xu, Yanjun Qi, and David Evans. "Automatically Evading Classifiers". In: *NDSS*. 2016.
- [36] Nicholas Carlini and David Wagner. "Towards evaluating the robustness of neural networks". In: *IEEE Symp. S&P*. 2017.
- [37] Nicholas Carlini and David A. Wagner. "Adversarial Examples Are Not Easily Detected: Bypassing Ten Detection Methods". In: *AISeC@CCS*. ACM, 2017, pp. 3–14.
- [38] Hung Dang, Yue Huang, and Ee-Chien Chang. "Evading Classifiers by Morphing in the Dark". In: *ACM Conference on Computer and Communications Security*. ACM, 2017, pp. 119–133.
- [39] Ambra Demontis et al. "Yes, machine learning can be more secure! a case study on android malware detection". In: *IEEE Transactions on Dependable and Secure Computing* (2017).
- [40] Sebastian Eresheim, Robert Luh, and Sebastian Schrittwieser. "The evolution of process hiding techniques in malware-current threats and possible countermeasures". In: *Journal of Information Processing* (2017).
- [41] Kathrin Grosse et al. "Adversarial examples for malware detection". In: *ESORICS*. Springer. 2017.
- [42] Mourad Leslous et al. "GPFinder: tracking the invisible in android malware". In: *MALWARE*. IEEE. 2017.
- [43] Aleksander Madry et al. "Towards deep learning models resistant to adversarial attacks". In: *arXiv preprint arXiv:1706.06083* (2017).

- [44] Niall McLaughlin et al. "Deep android malware detection". In: *Proceedings of the seventh ACM on conference on data and application security and privacy*. 2017, pp. 301–308.
- [45] Adam Paszke et al. "Automatic Differentiation in PyTorch". In: *NIPS Autodiff Workshop*. 2017.
- [46] Wei Yang et al. "Malware detection in adversarial settings: Exploiting feature evolutions and confusions in android apps". In: *ACSAC*. ACM. 2017.
- [47] Moustafa Alzantot et al. "Generating natural language adversarial examples". In: *Empirical Methods in Natural Language Processing (EMNLP)*. 2018.
- [48] Giovanni Apruzzese and Michele Colajanni. "Evading Botnet Detectors Based on Flows and Random Forest with Adversarial Samples". In: *IEEE NCA*. 2018.
- [49] Battista Biggio and Fabio Roli. "Wild Patterns: Ten Years After The Rise of Adversarial Machine Learning". In: *Pattern Recognition* (2018).
- [50] Nicholas Carlini and David Wagner. "Audio adversarial examples: Targeted attacks on speech-to-text". In: *Deep Learning for Security (DLS) Workshop*. IEEE. 2018.
- [51] Inigo Incer et al. "Adversarially robust malware detection using monotonic classification". In: *Proc. Int. Workshop on Security and Privacy Analytics*. ACM. 2018.
- [52] ElMouatez Billah Karbab et al. "MalDozer: Automatic framework for android malware detection using deep learning". In: *Digital investigation* 24 (2018), S48–S59.
- [53] Bojan Kolosnjaji et al. "Adversarial malware binaries: Evading deep learning for malware detection in executables". In: *EUSIPCO*. IEEE. 2018.
- [54] Bogdan Kulynych et al. "Evading classifiers in discrete domains with provable optimality guarantees". In: *CoRR* abs/1810.10939 (2018).
- [55] Marco Melis et al. "Explaining black-box android malware detection". In: *EUSIPCO*. IEEE. 2018.
- [56] Edward Raff et al. "Malware detection by eating a whole exe". In: *AAAI Workshops*. 2018.
- [57] Ishai Rosenberg et al. "Generic black-box end-to-end attack against state of the art API call based malware classifiers". In: *RAID*. Springer. 2018.
- [58] Octavian Suciuc et al. "When Does Machine Learning FAIL? Generalized Transferability for Evasion and Poisoning Attacks". In: *USENIX Security Symposium* (2018).
- [59] Giovanni Vigna and Davide Balzarotti. "When Malware is Packin' Heat". In: *USENIX ENIGMA*. 2018.
- [60] Ke Xu et al. "Deeprefiner: Multi-layer android malware detection system applying deep neural networks". In: *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE. 2018, pp. 473–487.
- [61] Giovanni Apruzzese, Michele Colajanni, and Mirco Marchetti. "Evaluating the effectiveness of Adversarial Attacks against Botnet Detectors". In: *IEEE NCA*. 2019.
- [62] Nicholas Carlini et al. "On evaluating adversarial robustness". In: *arXiv preprint arXiv:1902.06705* (2019).
- [63] Aurore Fass, Michael Backes, and Ben Stock. "HideNoSeek: Camouflaging Malicious JavaScript in Benign ASTs". In: *ACM CCS*. 2019.
- [64] Samuel G Finlayson et al. "Adversarial attacks on medical machine learning". In: *Science* 363.6433 (2019), pp. 1287–1289.
- [65] Jinfeng Li et al. "TextBugger: Generating Adversarial Text Against Real-world Applications". In: *NDSS*. 2019.
- [66] Davide Maiorca, Battista Biggio, and Giorgio Giacinto. "Towards Robust Detection of Adversarial Infection Vectors: Lessons Learned in PDF Malware". In: *arXiv preprint* (2019).
- [67] Alina Matyukhina et al. "Adversarial Authorship Attribution in Open-Source Projects". In: *Proc. of ACM Conference on Data and Applications Security and Privacy (CODASPY)*. 2019.
- [68] Rajvardhan Oak et al. "Malware detection on highly imbalanced data through sequence modeling". In: *Proceedings of the 12th ACM Workshop on artificial intelligence and security*. 2019, pp. 37–48.
- [69] Feargus Pendlebury et al. "TESSERACT: Eliminating Experimental Bias in Malware Classification across Space and Time". In: *28th USENIX Security Symposium*. USENIX Sec. Santa Clara, CA: USENIX Association, 2019.
- [70] Erwin Quiring, Alwin Maier, and Konrad Rieck. "Misleading Authorship Attribution of Source Code using Adversarial Learning". In: *USENIX Security Symposium*. 2019.
- [71] Shuhuai Ren et al. "Generating Natural Language Adversarial Examples through Probability Weighted Word Saliency". In: *Annual Meeting of the Assoc. for Computational Linguistics (ACL)*. 2019.
- [72] Octavian Suciuc, Scott E Coull, and Jeffrey Johns. "Exploring adversarial examples in malware detection". In: *2019 IEEE Security and Privacy Workshops (SPW)*. IEEE. 2019, pp. 8–14.
- [73] Liang Tong et al. "Improving robustness of {ML} classifiers against realizable evasion attacks using conserved features". In: *28th USENIX Security Symposium (USENIX Security 19)*. 2019, pp. 285–302.
- [74] Wei Wang, Mengxue Zhao, and Jigang Wang. "Effective android malware detection with a hybrid model based on deep autoencoder and convolutional neural network". In: *Journal of Ambient Intelligence and Humanized Computing* 10 (2019), pp. 3035–3043.

- [75] Qixue Xiao et al. “Seeing is Not Believing: Camouflage Attacks on Image Scaling Algorithms”. In: *USENIX Security Symposium*. USENIX Association, 2019, pp. 443–460.
- [76] Hongyang Zhang et al. “Theoretically principled trade-off between robustness and accuracy”. In: *International conference on machine learning*. PMLR. 2019, pp. 7472–7482.
- [77] Giulio Zizzo et al. “Adversarial Machine Learning Beyond the Image Domain”. In: *ACM DAC*. 2019.
- [78] Android. *Permissions overview - Dangerous Permissions*. 2020. URL: [https://developer.android.com/guide/topics/permissions/overview%5C#dangerous%5C\\_permissions](https://developer.android.com/guide/topics/permissions/overview%5C#dangerous%5C_permissions).
- [79] Jiyu Chen, David Wang, and Hao Chen. “Explore the transformation space for adversarial images”. In: *Proceedings of the Tenth ACM Conference on Data and Application Security and Privacy*. 2020, pp. 109–120.
- [80] Keane Lucas et al. “Malware Makeover: Breaking ML-Based Static Analysis by Modifying Executable Bytes”. In: *Proc. of ACM Asia Conference on Computer and Communications Security (ASIA CCS)*. 2021.
- [81] Wei Song et al. “Mab-malware: A reinforcement learning framework for blackbox generation of adversarial malware”. In: *Proceedings of the 2022 ACM on Asia conference on computer and communications security*. 2022, pp. 990–1003.
- [82] Salijona Dyrnishi et al. “On the empirical effectiveness of unrealistic adversarial hardening against realistic adversarial attacks”. In: *2023 IEEE symposium on security and privacy (SP)*. IEEE. 2023, pp. 1384–1400.
- [83] Ping He et al. “Efficient Query-Based Attack against ML-Based Android Malware Detection under Zero Knowledge Setting”. In: *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. 2023, pp. 90–104.
- [84] Keane Lucas et al. “Adversarial Training for {Raw-Binary} Malware Classifiers”. In: *32nd USENIX Security Symposium (USENIX Security 23)*. 2023, pp. 1163–1180.
- [85] Hamid Bostani and Veelasha Moonsamy. “Evadedroid: A practical evasion attack on machine learning for black-box android malware detection”. In: *Computers & Security* 139 (2024), p. 103676.
- [86] Marco Rando et al. “A New Formulation for Zeroth-Order Optimization of Adversarial EXamples in Malware Detection”. In: *arXiv preprint arXiv:2405.14519* (2024).

## Appendices

### A SYMBOL TABLE

Table 2 provides a reference for notation and major symbols used throughout the paper.

### B THREAT MODEL

The threat model must be defined in terms of attacker *knowledge* and *capability*, as in related literature [49, 58, 62]. While the attacker knowledge is represented in the same way as in the traditional feature-space attacks, their capability also includes the problem-space constraints  $\Gamma$ . For completeness, we report the threat model formalization proposed in Biggio and Roli [49].

**Attacker Knowledge.** We represent the knowledge as a set  $\theta \in \Theta$  which may contain (i) training data  $\mathcal{D}$ , (ii) the feature set  $\mathcal{X}$ , (iii) the learning algorithm  $g$ , along with the loss function  $\mathcal{L}$  minimized during training, (iv) the model parameters/hyperparameters  $\mathbf{w}$ . A parameter is marked with a *hat* symbol if the attacker knowledge of it is limited or only an estimate (i.e.,  $\hat{\mathcal{D}}$ ,  $\hat{\mathcal{X}}$ ,  $\hat{g}$ ,  $\hat{\mathbf{w}}$ ). There are three major scenarios [49]:

- *Perfect Knowledge (PK) white-box attacks*, in which the attacker knows all parameters and  $\theta_{PK} = (\mathcal{D}, \mathcal{X}, g, \mathbf{w})$ .
- *Limited Knowledge (LK) gray-box attacks*, in which the attacker has some knowledge on the target system. Two common settings are LK with Surrogate Data (LK-SD), where  $\theta_{LK-SD} = (\hat{\mathcal{D}}, \mathcal{X}, g, \hat{\mathbf{w}})$ , and LK with Surrogate Learners, where  $\theta_{LK-SL} = (\hat{\mathcal{D}}, \mathcal{X}, \hat{g}, \hat{\mathbf{w}})$ . Knowledge of the feature space and the ability to collect surrogate data,  $\theta \supseteq (\hat{\mathcal{D}}, \mathcal{X})$ , enables the attacker to perform *mimicry attacks* in which the attacker manipulates examples to resemble the high density region of the target class [18, 8].

Table 2. Table of symbols.

SYMBOL	DESCRIPTION
$\mathcal{Z}$	Problem space (i.e., input space).
$\mathcal{X}$	Feature space $\mathcal{X} \subseteq \mathbb{R}^n$ .
$\mathcal{Y}$	Label space.
$\varphi$	Feature mapping function $\varphi : \mathcal{Z} \rightarrow \mathcal{X}$ .
$h_i$	Discriminant function $h_i : \mathcal{X} \rightarrow \mathbb{R}$ that assigns object $\mathbf{x} \in \mathcal{X}$ a score in $\mathbb{R}$ (e.g., distance from hyperplane) that represents fitness to class $i \in \mathcal{Y}$ .
$g$	Classifier $g : \mathcal{X} \rightarrow \mathcal{Y}$ that assigns object $\mathbf{x} \in \mathcal{X}$ to class $y \in \mathcal{Y}$ . Also known as <i>decision function</i> . It is defined based on the output of the discriminant functions $h_i, \forall i \in \mathcal{Y}$ .
$\mathcal{L}_y$	Loss function $\mathcal{L}_y : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ of object $\mathbf{x} \in \mathcal{X}$ with respect to class $y \in \mathcal{Y}$ .
$f_{y,\kappa}$	Attack objective function $f_{y,\kappa} : \mathcal{X} \times \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}$ of object $\mathbf{x} \in \mathcal{X}$ with respect to class $y \in \mathcal{Y}$ with maximum confidence $\kappa \in \mathbb{R}$ .
$f_y$	Compact notation for $f_{y,0}$ .
$\Omega$	Feature-space constraints.
$\delta$	$\delta \in \mathbb{R}^n$ is a symbol used to denote a feature-space perturbation vector.
$\eta$	Side-effect feature vector.
$T$	Transformation $T : \mathcal{Z} \rightarrow \mathcal{Z}$ .
$\mathbf{T}$	Transformation sequence $\mathbf{T} = T_n \circ T_{n-1} \circ \dots \circ T_1$ .
$\mathcal{T}$	Space of available transformations.
$\Upsilon$	Suite of automated tests $\tau \in \Upsilon$ to verify preserved semantics.
$\Pi$	Suite of manual tests $\pi \in \Pi$ to verify plausibility. In particular, $\pi(z) = 1$ if $z \in \mathcal{Z}$ is plausible, else $\pi(z) = 0$ .
$\Lambda$	Set of preprocessing operators $\mathbf{A} \in \Lambda$ for which $z \in \mathcal{Z}$ should be resistant (i.e., $\mathbf{A}(\mathbf{T}(z)) = \mathbf{T}(z)$ ).
$\Gamma$	Problem-space constraints $\Gamma$ , consisting of $\{\Pi, \Upsilon, \mathcal{T}, \Lambda\}$ .
$\mathcal{D}$	Training dataset.
$\mathbf{w}$	Model hyper-parameters.
$\Theta$	Knowledge space.
$\theta$	Threat model assumptions $\theta \in \Theta$ ; more specifically, $\theta = (\mathcal{D}, \mathcal{X}, g, \mathbf{w})$ . A <i>hat</i> symbol is used if only estimates of parameters are known. See <a href="#">Appendix B</a> for details.

- *Zero Knowledge (ZK) black-box attacks*, where the attacker has no information on the target system, but has some information on which kind of feature extraction is performed (e.g., only static analysis in programs, or structural features in PDFs). In this case,  $\theta_{LK} = (\hat{\mathcal{D}}, \hat{\mathcal{X}}, \hat{g}, \hat{\mathbf{w}})$ .

Note that  $\theta_{PK}$  and  $\theta_{LK}$  imply knowledge of any defenses used to secure the target system against adversarial examples, depending on the degree to which each element is known [37].

**Attacker Capability.** The capability of an attacker is expressed in terms of his ability to modify feature space and problem space, i.e., the attacker capability is described through feature-space constraints  $\Omega$  and problem-space constraints  $\Gamma$ .

We observe that the attacker's knowledge and capability can also be expressed according to the FAIL [58] model as follows: knowledge of *Features*  $\mathcal{X}$  (F), the learning *Algorithm*  $g$  (A), *Instances* in training  $\mathcal{D}$  (I), *Leverage* on feature space and problem space with  $\Omega$  and  $\Gamma$  (L).

More details on the threat models can be found in [49, 58].

## C THEOREM PROOFS

**Proof of Theorem 1.** We proceed with a proof by contradiction. Let us consider a problem-space object  $z \in \mathcal{Z}$  with features  $\mathbf{x} \in \mathcal{X}$ , which we want to misclassify as a target class  $t \in \mathcal{Y}$ . Without loss of generality, we consider a low-confidence attack, with desired attack confidence  $\kappa = 0$  (see Equation 3). We assume by contradiction that there is no solution to the feature-space

attack; more formally, that there is no solution  $\delta^* = \arg \min_{\delta \in \mathbb{R}^n: \delta \models \Omega} f_t(\mathbf{x} + \delta)$  that satisfies  $f_t(\mathbf{x} + \delta^*) < 0$ . We now try to find a transformation sequence  $\mathbf{T}$  such that  $f_t(\varphi(\mathbf{T}(z))) < 0$ . Let us assume that  $\mathbf{T}^*$  is a transformation sequence that corresponds to a successful problem-space attack. By definition,  $\mathbf{T}^*$  is composed by individual transformations: a first transformation  $T_1$ , such that  $\varphi(T_1(z)) = \mathbf{x} + \delta_1$ ; a second transformation  $T_2$  such that  $\varphi(T_2(T_1(z))) = \mathbf{x} + \delta_1 + \delta_2$ ; a  $k$ -th transformation  $\varphi(T_k(\dots T_2(T_1(z)))) = \mathbf{x} + \sum_k \delta_k$ . We recall that the feature-space constraints are determined by the problem-space constraints, i.e.,  $\Gamma \vdash \Omega$ , and that, with slight abuse of notation, we can write that  $\Omega \subseteq \Gamma$ ; this means that the search space allowed by  $\Gamma$  is smaller or equal than that allowed by  $\Omega$ . Let us now replace  $\sum_k \delta_k$  with  $\delta^\dagger$ , which is a feature-space perturbation corresponding to the problem-space transformation sequence  $\mathbf{T}$ , such that  $f_t(\mathbf{x} + \delta^\dagger) < 0$  (i.e., the sample is misclassified). However, since the constraints imposed by  $\Gamma$  are stricter or equal than those imposed by  $\Omega$ , this means that  $\delta^\dagger$  must be a solution to  $\arg \min_{\delta \in \Omega} f_t(\mathbf{x} + \delta)$  such that  $f_t(\mathbf{x} + \delta^\dagger) < 0$ . However, this is impossible, because we hypothesized that there was no solution for the feature-space attack under the constraints  $\Omega$ . Hence, having a solution in the feature-space attack is a *necessary condition* for finding a solution for the problem-space attack.

**Proof of Theorem 2.** The existence of a feature-space attack (Equation 11) is the necessary condition, which has been already proved for Theorem 1. Here, we need to prove that, with Equation 12, the condition is sufficient for the attacker to find a problem-space transformation that misclassifies the object. Another way to write Equation 12 is to consider that the attacker knows transformations that affect individual features only (modifying more than one feature will result as a composition of such transformations). Formally, for any object  $z \in \mathcal{Z}$  with features  $\varphi(z) = \mathbf{x} \in X$ , for any feature-space dimension  $X_i$  of  $X$ , and for any value  $v \in \text{domain}(X_i)$ , let us assume the attacker knows a valid problem-space transformation sequence  $\mathbf{T} : \mathbf{T}(z) \models \Gamma, \varphi(\mathbf{T}(z)) = \mathbf{x}'$ , such that:

$$x'_i = x_i + v, \quad x_i \in \mathbf{x}, x'_i \in \mathbf{x}' \quad (14)$$

$$x'_j = x_j, \quad \forall j \neq i, x_j \in \mathbf{x}, x'_j \in \mathbf{x}' \quad (15)$$

Intuitively, these two equations refer to the existence of a problem-space transformation  $\mathbf{T}$  that affects only one feature  $X_i$  in  $X$  by any amount  $v \in \text{domain}(X_i)$ . In this way, given *any* adversarial feature-space perturbation  $\delta^*$ , the attacker is sure to find a transformation sequence that modifies each individual feature step-by-step. In particular, let us consider  $idx_0, \dots, idx_{q-1}$  corresponding to the  $q > 0$  values in  $\delta^*$  that are different from 0 (i.e., values corresponding to an actual feature-space perturbation). Then, a transformation sequence  $\mathbf{T} : \mathbf{T}(z) \models \Gamma, \mathbf{T} = \mathbf{T}^{idx_q-1} \circ \mathbf{T}^{idx_{q-2}} \circ \dots \circ \mathbf{T}^{idx_0}$  can always be constructed by the attacker to satisfy  $\varphi(\mathbf{T}(z)) = \mathbf{x} + \delta^*$ . We highlight that we do not consider the existence of a specific transformation in  $\mathcal{Z}$  that maps to  $\mathbf{x} + \delta^*$  because that may not be known by the attacker; hence, the attacker may never learn such a specific transformation. Thus, Equation 12 must be valid for all possible perturbations within the considered feature space.

## D OPAQUE PREDICATES GENERATION

We use opaque predicates [10] as inconspicuous conditional statements always resolving to False to preserve dynamic semantics of the Android applications.

To ensure the intractability of such an analysis, we follow the work of [11] and build opaque predicates using a formulation of the 3-SAT problem such that resolving the truth value of the predicate is equivalent to solving the NP-complete 3-SAT problem.

The  $k$ -satisfiability ( $k$ -SAT) problem asks whether the variables of a Boolean logic formula can be consistently replaced with True or False so that the entire formula evaluates to True; if so the

formula is *satisfiable*. Such a formula is easily expressed in its conjunctive normal form:

$$\bigwedge_{i=1}^m (V_{i1} \vee V_{i2} \vee \dots \vee V_{ik}),$$

where  $V_{ij} \in \{v_1, v_2, \dots, v_n\}$  are Boolean variables and  $k$  is the number of variables per clause.

To consistently generate NP-Hard  $k$ -SAT problems we use *Random  $k$ -SAT* [3] in which there are 3 parameters: the number of variables  $n$ , the number of clauses  $m$ , and the number of literals per clause  $k$ .

To construct a 3-SAT formula,  $m$  clauses of length 3 are generated by randomly choosing a set of 3 variables from the  $n$  available, and negating each with probability 50%. An empirical study by [3] showed that  $n$  should be at least 40 to ensure the formulas are hard to resolve. Additionally, they show that formulas with too few clauses are *under-constrained* while formulas with too many clauses are *over-constrained*, both of which reduce the search time. These experiments led to the following conjecture. For further details please refer to the conference version.

## E DEFENDER EVALUATION SETTINGS

In §5.1, we examine different strategies for enhancing classifier robustness, which include varying approaches to generating adversarial examples, levels of confidence, and hardening methods. We adjust the proportion of adversarial examples from the training set, in 10% increments from 10 to 100%, to see how it affects robustness. The study also considers other variables such as the model used, hardening approach, and transformation type. In our experimental framework, we explore several key dimensions to comprehensively evaluate model hardening effectiveness. Regarding confidence levels, we employ the two distinct levels defined in Section §4.1, applying them across both hardening and testing phases. This creates four possible confidence configurations: High-Low, where models are hardened using high confidence adversarial points but tested against low confidence evasion attempts; Low-Low, utilizing low confidence adversarial points for both hardening and testing; High-High, employing high confidence points for both processes; and Low-High, where hardening occurs with low confidence points while testing uses high confidence evasion techniques. For our model selection, we conduct parallel analyses on both LinearSVM and Sec-SVM implementations to compare their respective resilience characteristics. The hardening methodologies under investigation include both adversarial training and adversarial retraining approaches, allowing us to contrast their effectiveness under various conditions. Furthermore, our transformation evaluation framework encompasses both feature space transformations and problem space transformations during the hardening process, while consistently evaluating against adversarial points generated through realistic transformations. This comprehensive evaluation strategy enables us to assess hardening efficacy across multiple dimensions of the adversarial machine learning landscape.

## F ATTACK ALGORITHMS

As previously described in the paper, we adopt a greedy approach in order to create adversarial samples, both in the feature space and in the problem space. Algorithm 1 and Algorithm 2 describe in detail the two main phases of our search strategy: organ harvesting and adversarial program generation. For the sake of simplicity, we describe a low-confidence attack, i.e., the attack is considered successful as soon as the classification score is below zero. It is immediate to consider high-confidence variations (as we evaluate in §4).



**Input:** Discriminant function  $h(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + b$ , which classifies  $\mathbf{x}$  as malware if  $h(\mathbf{x}) > 0$ , otherwise as goodwill. Minimal app  $z_{min} \in \mathcal{Z}$  with features  $\varphi(z_{min}) = \mathbf{x}_{min}$ .

**Parameters:** Number of features to consider  $n_f$ ; number of donors per-feature  $n_d$ .

**Output:** Ice-box of harvested organs with feature vectors.

---

ice-box  $\leftarrow \{\}$  Empty key-value dictionary.  
 $L \leftarrow$  List of pairs  $(w_i, i)$ , sorted by increasing value of  $w_i$ .  
 $L' \leftarrow$  First  $n_f$  elements of  $L$ , then remove any entry with  $w_i \geq 0$ .

**for**  $(w_i, i)$  **in**  $L'$  **do**  
  ice-box[i]  $\leftarrow \{\}$  Empty list for gadgets with feature  $i$ .  
  **while**  $\text{length}(\text{ice-box}[i]) < n_d$  **do**  
     $z_j \leftarrow$  Randomly sample a benign app with feature  $x_i = 1$ .  
    Extract gadget  $\rho_j \in \mathcal{Z}$  with feature  $x_i = 1$  from  $z_j$ .  
     $s \leftarrow$  Software stats of  $\rho_j$   
     $z' \leftarrow$  Inject gadget  $\rho_j$  in app  $z_{min}$ .  
     $(\mathbf{x}_{min} \vee \mathbf{e}_i \vee \eta_j) \leftarrow \varphi(z')$   $\mathbf{e}_i$  is a one-hot vector.  
     $r_j \leftarrow (\mathbf{e}_i \vee \eta_j) \leftarrow \varphi(z') \wedge \neg \mathbf{x}_{min}$  Gadget features obtained through set difference.  
    **if**  $h(r_j) > 0$  **then**  
      Discard the gadget;  
    **else**  
      Append  $(\rho_j, r_j, s)$  to ice-box[i]. Store gadget  
**return** ice-box;

### Algorithm 1: Initialization (Ice-Box Creation)

**Input:** Discriminant function  $h(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + b$ , which classifies  $\mathbf{x}$  as malware if  $h(\mathbf{x}) > 0$ , otherwise as goodwill. Malware app  $z \in \mathcal{Z}$ . Ice-box  $G$ .

**Parameters:** Problem-space constraints.

**Output:** Adversarial app  $z' \in \mathcal{Z}$  such that  $h(\varphi(z')) < 0$ .

---

$\mathcal{T} \leftarrow$  Transplantation through gadget addition.  
 $\Upsilon \leftarrow$  Smoke test through app installation and execution in emulator.  
 $\Pi \leftarrow$  Plausibility by-design through code consolidation.  
 $\Lambda \leftarrow$  Artifacts from last column of Table 1.  
 $\Gamma \leftarrow \{\mathcal{T}, \Upsilon, \Pi, \Lambda\}$   
 $s_z \leftarrow$  Software stats of  $z$   
 $\mathbf{x} \leftarrow \varphi(z)$   
 $L \leftarrow \{\}$  Empty list.  
 $T(z) \leftarrow$  Empty sequence of problem-space transformations.

**for**  $(\rho_j, r_j, s)$  **in**  $G$  **do**  
   $d_j \leftarrow r_j \wedge \neg \mathbf{x}$  Feature-space contribution of gadget  $j$ .  
   $score_j \leftarrow h(d_j)$  Impact on decision score.  
  Append the pair  $(score_j, i, j)$  to  $L$  Feature  $i$ , Gadget  $j$ .

**for**  $(score_j, i, j)$  **in**  $L'$  **do**  
  **if**  $z$  has  $x_i = 1$  **then**  
    Do nothing; Feature  $i$  already present.  
  **else if**  $z$  has  $x_i = 0$  **then**  
     $(\rho_j, r_j, s) \leftarrow$  element  $j$  in ice-box  $G$   
    **if**  $\text{check\_feasibility}(s_z, s)$  is *True* **then**  
       $\mathbf{x} \leftarrow (\mathbf{x} \vee \mathbf{e}_i \vee \eta_j)$  Update features of  $z$ .  
      Append transplantation  $T \in \mathcal{T}$  of gadget  $\rho_j$  in  $T(z)$ .  
      **if**  $h(\mathbf{x}) < 0$  **then**  
        Exit from cycle; Attack gadgets found.  
   $z' \leftarrow$  Apply transformation sequence  $T(z)$  Inject chosen gadgets.

**if**  $h(\varphi(z')) < 0$  **and**  $T(z) \models \Gamma$  **then**  
  **return**  $z'$ ; Attack successful.  
**else**  
  **return** Failure;

### Algorithm 2: Attack (Adv. Program Generation)

## G SIDE EFFECT FEATURES EXAMPLE

This section provides a concrete example to illustrate the impact of side-effect features within our domain of study. Specifically, [algorithm 2](#) identifies the URL <http://intradaystatergies.blogspot.com/2018/07/blog-post.html> as the most benign feature in the initial analysis. Further investigation of the corresponding code slice in the icebox reveals a total of 1098 Jimple classes [12] and eight distinct side-effect features:

- API call: `api_calls::android/content/Context;->startActivity`
- API call: `api_calls::android/location/LocationManager;->getLastKnownLocation`
- Intent: `intents::android_intent_action_MAIN`
- Intent: `intents::android_intent_category_LAUNCHER`
- API call: `api_calls::java/net/URLConnection;->getInputStream`
- API permission: `api_permissions::android_permission_ACCESS_FINE_LOCATION`
- API permission: `api_permissions::android_permission_ACCESS_INTERNET`

Due to space constraints, we are only reporting the code showing the presence of the side-effect features in the bytecode of one of the extracted classes.

Listing 1. Bytecode extracted from a benign donor app containing a side effect feature. The target feature is `startActivity` as highlighted.

```
public void onStartActivityFromFragment(android.support.v4.app.Fragment, android.content.
    Intent, int, android.os.Bundle)
{
    android.content.Context $r4;
    java.lang.IllegalStateException $r5;
    android.os.Bundle $r3;
    android.support.v4.app.Fragment $r1;
    int $i0;
    android.content.Intent $r2;
    android.support.v4.app.FragmentHostCallback r0;
    r0 := @this: android.support.v4.app.FragmentHostCallback;
    $r1 := @parameter0: android.support.v4.app.Fragment;
    $r2 := @parameter1: android.content.Intent;
    $i0 := @parameter2: int;
    $r3 := @parameter3: android.os.Bundle;
    if $i0 == -1 goto label1;

    ...
label1:
    $r4 = r0.<android.support.v4.app.FragmentHostCallback: android.content.Context
        mContext>;
    virtualinvoke $r4.<android.content.Context: void StartActivity(android.content.Intent)>($r2);
    return;
}
```

The injection process entails integrating the complete collection of Jimple classes along with their associated features—comprising eight side-effect features and one primary feature—into the target malware. Our examination of 18,000 extracted gadgets indicates an average of six side-effect features for each harvested feature. These statistics are derived from the entire set of extracted gadgets, without accounting for the imposed limitations of injecting 20 classes and one permission. However, this example clearly illustrates that there is no direct relationship between the number of injected Jimple classes and the corresponding feature set. In fact, the quantity of Jimple classes is often significantly greater than the number of features they contain. This discrepancy suggests that, for a stealthy injection process, certain constraints must be considered to prevent the creation of a final adversarial sample that could raise suspicions due to the excessive number of injected classes.