

When Does Machine Learning FAIL? Generalized Transferability for Evasion and Poisoning Attacks

Octavian Suci

Radu Mărginean

Yiğitcan Kaya

Hal Daumé III

Tudor Dumitraş

University of Maryland, College Park

Abstract

Recent results suggest that attacks against supervised machine learning systems are quite effective, while defenses are easily bypassed by new attacks. However, the specifications for machine learning systems currently lack precise adversary definitions, and the existing attacks make diverse, potentially unrealistic assumptions about the strength of the adversary who launches them. We propose the **FAIL** attacker model, which describes the adversary’s knowledge and control along four dimensions. Our model allows us to consider a wide range of weaker adversaries who have limited control and incomplete knowledge of the features, learning algorithms and training instances utilized.

To evaluate the utility of the **FAIL** model, we consider the problem of conducting targeted poisoning attacks in a realistic setting: the crafted poison samples must have clean labels, must be individually and collectively inconspicuous, and must exhibit a generalized form of transferability, defined by the **FAIL** model. By taking these constraints into account, we design StingRay, a targeted poisoning attack that is practical against 4 machine learning applications, which use 3 different learning algorithms, and can bypass 2 existing defenses. Conversely, we show that a prior evasion attack is less effective under generalized transferability. Such attack evaluations, under the **FAIL** adversary model, may also suggest promising directions for future defenses.

1 Introduction

Machine learning (ML) systems are widely deployed in safety-critical domains that carry incentives for potential adversaries, such as finance [14], medicine [18], the justice system [31], cybersecurity [1], or self-driving cars [6]. An ML classifier automatically learns classification models using labeled observations (samples) from a *training set*, without requiring predetermined rules for

mapping inputs to labels. It can then apply these models to predict labels for new samples in a *testing set*. An *adversary* knows some or all of the ML system’s parameters and uses this knowledge to craft training or testing samples that manipulate the decisions of the ML system according to the adversary’s goal—for example, to avoid being sentenced by an ML-enhanced judge.

Recent work has focused primarily on *evasion* attacks [4, 44, 17, 50, 35, 9], which can induce a *targeted* misclassification on a specific sample. As illustrated in Figures 1a and 1b, these test time attacks work by mutating the target sample to push it across the model’s decision boundary, without altering the training process or the decision boundary itself. They are not applicable in situations where the adversary does not control the target sample—for example, when she aims to influence a malware detector to block a benign app developed by a competitor. Prior research has also shown the feasibility of targeted *poisoning* attacks [34, 32]. As illustrated in Figure 1c, these attacks usually blend crafted instances into the training set to push the model’s boundary toward the target. In consequence, they enable misclassifications for instances that the adversary cannot modify.

These attacks appear to be very effective, and the defenses proposed against them are often bypassed in follow-on work [8]. However, to understand the actual security threat introduced by them, we must model the capabilities and limitations of realistic adversaries. Evaluating poisoning and evasion attacks under assumptions that overestimate the capabilities of the adversary would lead to an inaccurate picture of the security threat posed to real-world applications. For example, test time attacks often assume white-box access to the victim classifier [9]. As most security-critical ML systems use proprietary models [1], these attacks might not reflect actual capabilities of a potential adversary. Black-box attacks consider weaker adversaries, but they often make rigid assumptions about the adversary’s knowledge when investigating the *transferability* of an attack. Transferabil-




Figure 1: Targeted attacks against machine learning classifiers. (a) The pristine classifier would correctly classify the target. (b) An evasion attack would modify the target to cross the decision boundary. (c) Correctly labeled poisoning instances change the learned decision boundary. (d) At testing time, the target is misclassified but other instances are correctly classified.

ity is a property of attack samples crafted locally, on a surrogate model that reflects the adversary’s limited knowledge, allowing them to remain successful against the target model. Specifically, black-box attacks often investigate transferability in the case where the local and target models use different training *algorithms* [36]. In contrast, ML systems used in the security industry often resort to *feature* secrecy (rather than algorithmic secrecy) to protect themselves against attacks, e.g. by incorporating undisclosed features for malware detection [10].

In this paper, we make a first step towards modeling realistic adversaries who aim to conduct attacks against ML systems. To this end, we propose the **FAIL** model, a general framework for the analysis of ML attacks in settings with a variable amount of adversarial knowledge and control over the victim, along four tunable dimensions: **F**eatures, **A**lgorithms, **I**nstances, and **L**everage. By preventing any implicit assumptions about the adversarial capabilities, the model is able to accurately highlight the success rate of a wide range of attacks in realistic scenarios and forms a common ground for modeling adversaries. Furthermore, the **FAIL** framework generalizes the transferability of attacks by providing a multidimensional basis for surrogate models. This provides insights into the constraints of realistic adversaries, which could be explored in future research on defenses against these attacks. For example, our evaluation suggests that crafting transferable samples with an existing evasion attack is more challenging than previously believed.

To evaluate the utility of the **FAIL** model, we consider the problem of *conducting targeted poisoning attacks in a realistic setting*. Specifically, we impose four constraints on the adversary. First, the poison samples must have *clean labels*, as the adversary can inject them into the training set of the model under attack but cannot determine how they are labeled. Second, the samples must be *individually inconspicuous*, i.e. to be very similar to the existing training instances in order to prevent an easy detection, while collectively pushing the model’s boundary toward a target instance. Third, the samples must be *collectively inconspicuous* by bounding the col-

lateral damage on the victim (Figure 1d). Finally, the poison samples must exhibit a generalized form of transferability, as the adversary tests the samples on a surrogate model, trained with partial knowledge along multiple dimensions, defined by the **FAIL** model.

By taking into account the goals, capabilities, and limitations of realistic adversaries, we also design StingRay, a targeted poisoning attack that can be applied in a broad range of settings¹. Moreover, the StingRay attack is *model agnostic*: we describe concrete implementations against 4 ML systems, which use 3 different classification algorithms (convolutional neural network, linear SVM, and random forest). The instances crafted are able to bypass three anti-poisoning defenses, including one that we adapted to account for targeted attacks. By subjecting StingRay to the **FAIL** analysis, we obtain insights into the *transferability of targeted poison samples*, and we highlight promising directions for investigating defenses against this threat.

In summary, this paper makes three contributions:

- We introduce the **FAIL** model, a general framework for modeling realistic adversaries and evaluating their impact. The model generalizes the transferability of attacks against ML systems, across various levels of adversarial knowledge and control. We show that a previous black-box evasion attack is less effective under generalized transferability.
- We propose StingRay, a targeted poisoning attack that overcomes the limitations of prior attacks. StingRay is effective against 4 real-world classification tasks, even when launched by a range of weaker adversaries within the **FAIL** model. The attack also bypasses two existing anti-poisoning defenses.
- We systematically explore realistic adversarial scenarios and the effect of partial adversary knowledge and control on the resilience of ML models against a test-time attack and a training-time attack. Our

¹Our implementation code could be found at <https://github.com/sdsatumd>

results provide insights into the transferability of attacks across the FAIL dimensions and highlight potential directions for investigating defenses against these attacks.

This paper is organized as follows. In Section 2 we formalize the problem and our threat model. In Section 3 we introduce the FAIL attacker model. In Section 4 we describe the StingRay attack and its implementation. We present our experimental results in Section 5, review the related work in Section 6, and discuss the implications in Section 7.

2 Problem Statement

Lack of a unifying threat model to capture the dimensions of adversarial knowledge caused existing work to diverge in terms of adversary specifications. Prior work defined adversaries with inconsistent capabilities. For example, in [36] a *black-box* adversary possesses knowledge of the full feature representations, whereas its counterpart in [50] only assumes access to the raw data (i.e. before feature extraction).

Compared to existing white-box or black-box models, in reality, things tend to be more nuanced. A commercial ML-based malware detector [1] can rely on a publicly known architecture with proprietary data collected from end hosts, and a mixture of known features (e.g. system calls of a binary), and undisclosed features (e.g. reputation scores of the binary). Existing adversary definitions are too rigid and cannot account for realistic adversaries against such applications. In this paper, we ask *how can we systematically model adversaries based on realistic assumptions about their capabilities?*

Some of the recent evasion attacks [28, 36] investigate the transferability property of their solutions. Proven transferability increases the strength of an attack as it allows adversaries with limited knowledge or access to the victim system to craft effective instances. Furthermore, transferability hinders defense strategies as it renders secrecy ineffective. However, existing work generally investigates transferability under single dimensions (e.g. limiting the adversarial knowledge about the victim algorithm). This weak notion of transferability limits the understanding of actual attack capabilities on real systems and fails to shed light on potential avenues for defenses. This paper aims to provide a means to define and evaluate a more *general transferability*, across a wide range of adversary models. The generalized view of threat models highlights limitations of existing training-time attacks. Existing attacks [51, 29, 20] often assume full control over the training process of victim classifiers and have similar shortcomings to white-box attacks. Those that do not assume full control generally omit im-

portant adversarial considerations. Targeted poisoning attacks [34, 32, 11] require control of the labeling process. However, an attacker is often unable to determine the labels assigned to the poison samples in the training set—consider a case where a malware creator may provide a poison sample for the training set of an ML-based malware detector, but its malicious/benign label will be assigned by the engineers who train the detector. These attacks risk being detected by existing defenses as they might craft samples that stand out from the rest of the training set. Moreover, they also risk causing collateral damage to the classifier; for example, in Figure 1c the attack can trigger the misclassification of additional samples from the target’s true class if the boundary is not molded to include only the target. Such collateral damage reduces the trust in the classifier’s predictions, and thus the potential impact of the attack. Therefore, we aim to observe whether an attack could address these limitations and discover *how realistic is the targeted poisoning threat?*

Machine learning background. For our purpose, a *classifier* (or hypothesis) is a function $h : X \rightarrow Y$ that maps instances to labels to perform *classification*. An *instance* $\mathbf{x} \in X$ is an entity (e.g., a binary program) that must receive a *label* $y \in Y = \{y_0, y_1, \dots, y_m\}$ (e.g., reflecting whether the binary is malicious). We represent an instance as a vector $\mathbf{x} = (x_1, \dots, x_n)$, where the features reflect attributes of the artifact (e.g. APIs invoked by the binary). A function $D(\mathbf{x}, \mathbf{x}')$ represents the distance in the feature space between two instances $\mathbf{x}, \mathbf{x}' \in X$. The function h can be viewed as a separator between the malicious and benign classes in the feature space X ; the plane of separation between classes is called *decision boundary*. The *training set* $S \subset X$ includes instances that have known labels $Y_S \subset Y$. The labels for instances in S are assigned using an *oracle*—for a malware classifier, an oracle could be an antivirus service such as VirusTotal, whereas for an image classifier it might be a human annotator. The *testing set* $T \subset X$ includes instances for which the labels are unknown to the learning algorithm.

Threat model. We focus on *targeted poisoning attacks* against machine learning classifiers. In this setting, we refer to the victim classifier as Alice, the owner of the target instance as Bob, and the attacker as Mallory. Bob and Mallory could also represent the same entity. Bob possesses an instance $\mathbf{t} \in T$ with label y_t , called the *target*, which will get classified by Alice. For example, Bob develops a benign application, and he ensures it is not flagged by an oracle antivirus such as VirusTotal. Bob’s expectation is that Alice would not flag the instance either. Indeed, the target would be correctly classified by Alice after learning a hypothesis using a pristine training set S^* (i.e. $h^* = A(S^*), h^*(\mathbf{t}) = y_t$). Mallory has partial

knowledge of Alice’s classifier and read-only access to the target’s feature representation, but they do not control either \mathbf{t} or the natural label y_t , which is assigned by the oracle. Mallory pursues two goals. The *first goal* is to introduce a targeted misclassification on the target by deriving a training set S from S^* : $h = A(S), h(\mathbf{t}) = y_d$, where y_d is Mallory’s desired label for \mathbf{t} . On binary classification, this translates to causing a false positive (FP) or false negative (FN). An example of FP would be a benign email message that would be classified as spam, while an FN might be a malicious sample that is not detected. In a multiclass setting, Mallory causes the target to be labeled as a class of choice. Mallory’s *second goal* is to minimize the effect of the attack on Alice’s overall classification performance. To quantify this collateral damage, we introduce the Performance Drop Ratio (PDR), a metric that reflects the performance hit suffered by a classifier after poisoning. This is defined as the ratio between the performance of the poisoned classifier and that of the pristine classifier: $PDR = \frac{\text{performance}(h)}{\text{performance}(h^*)}$. The metric encodes the fact that for a low-error classifier, Mallory could afford a smaller performance drop before raising suspicions.

3 Modeling Realistic Adversaries

Knowledge and Capabilities. Realistic adversaries conducting training time or testing time attacks are constrained by an imperfect *knowledge* about the model under attack and by limited *capabilities* in crafting adversarial samples. For an attack to be successful, samples crafted under these conditions must transfer to the original model. We formalize the adversary’s strength in the **FAIL** attacker model, which describes the adversary’s knowledge and capabilities along 4 dimensions:

- Feature knowledge $\mathcal{R} = \{x_i : x_i \in \mathbf{x}, x_i \text{ is readable}\}$: the subset of features known to the adversary.
- Algorithm knowledge A' : the learning algorithm that the adversary uses to craft poison samples.
- Instance knowledge S' : the labeled training instances available to the adversary.
- Leverage $\mathcal{W} = \{x_i : x_i \in \mathbf{x}, x_i \text{ is writable}\}$: the subset of features that the adversary can modify.

The **F** and **A** dimensions constrain the attacker’s understanding of the hypothesis space. Without knowing the victim classifier A , the attacker would have to select an alternative learning algorithm A' and hope that the evasion or poison samples crafted for models created by A' transfer to models from A . Similarly, if some features

are unknown (i.e., partial feature knowledge), the model used for crafting instances is an approximation of the original classifier. For classifiers that learn a representation of the input features (such as neural networks), limiting the **F** dimension results in a different, approximate internal representation that will affect the success rate of the attack. These limitations result in an inaccurate *assessment* of the impact that the crafted instances will have and affect the success rate of the attack. The **I** dimension affects the accuracy of the adversary’s view over the instance space. As S' might be a subset or an approximation of S^* , the poisoning and evasion samples might exploit gaps in the instance space that are not present in the victim’s model. This, in turn, could lead to an impact overestimation on the attacker side. Finally, the **L** dimension affects the adversary’s *ability* to craft attack instances. The set of modifiable features restricts the regions of the feature space where the crafted instances could lie. For poisoning attacks, this places an upper bound on the ability of samples to shift the decision boundary while for evasion it could affect their effectiveness. The read-only features can, in some cases, cancel out the effect of the modified ones. An adversary with partial leverage needs extra effort, e.g. to craft more instances (for poisoning) or to attack more of the modifiable features (for both poisoning and evasion).

Prior work has investigated transferability without modeling a full range of realistic adversaries across the FAIL dimensions. [36] focuses on the **A** dimension, and proposes a transferable evasion attack across different neural network architectures. Transferability of poisoning samples in [33] is partially evaluated on the **I** and **A** dimensions. The evasion attack in [25] considers **F**, **A** and **I** under a coarse granularity, but omits the **L** dimension. ML-based systems employed in the security industry [21, 10, 45, 39, 12] often combine undisclosed and known features to render attacks more difficult. In this context, the systematic evaluation of transferability along the **F** and **L** dimensions is still an open question.

Constraints. The attacker’s strategy is also influenced by a set of *constraints* that drive the attack design and implementation. While these are attack-dependent, we broadly classify them into three categories: *success*, *defense*, and *budget* constraints. *Success* constraints encode the attacker’s goals and considerations that directly affect the effectiveness of the attack, such as the assessment of the target instance classification. *Defense* constraints refer to the attack characteristics aimed to circumvent existing defenses (e.g. the post-attack performance drop on the victim). *Budget* considerations address the limitations in an attacker’s resources, such as the maximum number of poisoning instances or, for evasion attacks, the maximum number of queries to the victim model.

Implementing the FAIL dimensions. Performing empirical evaluations within the FAIL model requires further design choices that depend on the application domain and the attack surface of the system. To simulate weaker adversaries systematically, we formulate a questionnaire to guide the design of experiments focusing on each dimension of our model.

For the **F** dimension, we ask: *What features could be kept as a secret? Could the attacker access the exact feature values?* Feature subsets may not be publicly available (e.g. derived using a proprietary malware analysis tool, such as dynamic analysis in a contained environment), or they might be directly defined from instances not available to the attacker (e.g. low-frequency word features). Similarly, the exact feature values could be unknown (e.g. because of defensive feature squeezing [49]). Feature secrecy does not, however, imply the attacker’s inability to modify them through an indirect process [25] or extract surrogate ones.

The questions related to the **A** dimension are: *Is the algorithm class known? Is the training algorithm secret? Are the classifier parameters secret?* These questions define the spectrum for adversarial knowledge with respect to the learning algorithm: black-box access, if the information is public, gray-box, where the attacker has partial information about the algorithm class or the ensemble architecture, or white-box, for complete adversarial knowledge.

The **I** dimension controls the overlap between the instances available to the attacker and these used by the victim. Thus, here we ask: *Is the entire training set known? Is the training set partially known? Are the instances known to the attacker sufficient to train a robust classifier?* An application might use instances from the public domain (e.g. a vulnerability exploit predictor) and the attacker could leverage them to the full extent in order to derive their attack strategy. However, some applications, such as a malware detector, might rely on private or scarce instances that limit the attacker’s knowledge of the instance space. The scarcity of these instances drives the robustness of the attacker classifier which in turn defines the perceived attack effectiveness. In some cases, the attacker might not have access to any of the original training instances, being forced to train a surrogate classifier on independently collected samples [50, 29].

The **L** dimension encodes the practical capabilities of the attacker when crafting attack samples. These are tightly linked to the attack constraints. However, rather than being preconditions, they act as degrees of freedom on the attack. Here we ask: *Which features are modifiable by the attacker?* and *What side effects do the modifications have?* For some applications, the attacker may not be able to modify certain types of features, either because they do not control the generating process (e.g. an

Study	F	A	I	L
Test Time Attacks				
Genetic Evasion[50]	✓,✓	✓,✓	✓,X†	✓,✓
Black-box Evasion[37]	X,∅*	✓,✓	✓,✓	X,∅*
Model Stealing[46]	✓,✓	✓,✓	✓,✓	X,∅*
FGSM Evasion[17]	X,∅*	X,∅*	∅,∅	X,∅*
Carlini’s Evasion[9]	X,∅*	✓,✓	∅,∅	X,∅*
Training Time Attacks				
SVM Poisoning[5]	X,∅*	✓,X†	∅,∅	X,∅*
NN Poisoning[33]	✓,X†	✓,✓	✓,✓	X,∅*
NN Backdoor[20] ²	✓,X†	✓,✓	✓,X†	✓,✓
NN Trojan[29]	✓,X	✓,✓	✓,✓	✓,✓

Table 1: FAIL analysis of existing attacks. For each attack, we analyze the adversary model and evaluation of the proposed technique. Each cell contains the answers to our two questions, *AQ1* and *AQ2*: *yes* (✓), *omitted* (X) and *irrelevant* (∅). We also flag *implicit assumptions* (*) and a *missing evaluation* (†).

Study	F	A	I	L
Test Time Defenses				
Distillation[38]	X,✓	X,✓	X,X	X,X
Feature Squeezing[49]	✓,✓	X,X	X,X	✓,✓
Training Time Defenses				
RONI[34]	X,X	X,X	✓,X	X,X
Certified Defense[42]	X,X	X,X	✓,✓	X,X

Table 2: FAIL analysis of existing defenses. We analyze a defense’s approach to security: *DQ1* (secrecy) and *DQ2* (hardening). Each cell contains the answers to the two questions: *yes* (✓), and *no* (X).

exploit predictor that gathers features from multiple vulnerability databases) or when the modifications would compromise the instance integrity (e.g. a watermark on images that prevents the attacker from modifying certain features). In cases of dependence among features, targeting a specific set of features could have an indirect effect on others (e.g. an attacker injecting tweets to modify word feature distributions also changes features based on tweet counts).

3.1 Unifying Threat Model Assumptions

Discordant threat model definitions result in implicit assumptions about adversarial limitations, some of which might not be realistic. The FAIL model allows us to systematically reason about such assumptions. To demonstrate its utility, we evaluate a body of existing studies by means of answering two questions for each work.

²Gu et al.’s study investigates a scenario where the attacker performs the training on behalf of the victim. Consequently, the attacker has full access to the model architecture, parameters, training set and feature representation. However, with the emergence of frameworks such as [16], even in this threat model, it might be possible that the attacker does not know the training set or the features.

To categorize existing attacks, we first inspect a threat model and ask: *AQ1—Are bounds for attacker limitations specified along the dimension?*. The possible answers are: *yes*, *omitted* and *irrelevant*. For instance, the threat model in Carlini et al.’s evasion attack [9] specifies that the adversary requires complete knowledge of the model and its parameters, thus the answer is *yes* for the **A** dimension. In contrast, the analysis on the **I** dimension is *irrelevant* because the attack does not require access to the victim training set. However, the study does not discuss feature knowledge, therefore we mark the **F** dimension as *omitted*.

Our second question is: *AQ2—Is the proposed technique evaluated along the dimension?*. This question becomes *irrelevant* if the threat model specifications are *omitted* or *irrelevant*. For example, Carlini et al. evaluated transferability of their attack when the attacker does not know the target model parameters. This corresponds to the attacker algorithm knowledge, therefore the answer is *yes* for the **A** dimension.

Applying the FAIL model reveals implicit assumptions in existing attacks. An implicit assumption exists if the attack limitations are not specified along a dimension. Furthermore, even with explicit assumptions, some studies do not evaluate all relevant dimensions. We present these findings about previous attacks within the FAIL model in Table 1.

When looking at existing defenses through the FAIL model, we aim to observe how they achieve security: either by hiding information or limiting the attacker capabilities. For defenses that involve creating knowledge asymmetry between attackers and the defenders, i.e. secrecy, we ask: *DQ1—Is the dimension employed as a mechanism for secrecy?*. For example, feature squeezing [49] employs feature reduction techniques unknown to the attacker; therefore the answer is *yes* for the **F** dimension.

In order to identify hardening dimensions, which attempt to limit the attack capabilities, we ask: *DQ2—Is the dimension employed as a mechanism for hardening?*. For instance, the distillation defense [38] against evasion modifies the neural network weights to make the attack more difficult; therefore the answer is *yes* for the **A** dimension.

These defenses may come with inaccurate assessments for the adversarial capabilities and implicit assumptions. For example, distillation limits adversaries along the **F** and **A** dimensions but employing a different attack strategy could bypass it [9]. On poisoning attacks, the RONI [34] defense assumes training set secrecy, but does not evaluate the threat posed by attackers with sufficient knowledge along the other dimensions. As our results will demonstrate, this implicit assumption allows attackers to bypass the defense while remaining within the se-

crecy bounds.

The results for the evaluated defenses are found in Table 2. The detailed evaluation process for each of these studies can be found in our technical report [43].

4 The StingRay Attack

Reasoning about implicit and explicit assumptions in prior defenses allows us to design algorithms which exploit their weaknesses. In this section, we introduce StingRay, one such attack that achieves targeted poisoning while preserving overall classification performance. StingRay is a general framework for crafting poison samples.

At a high level, our attack builds a set of poison instances by starting from base instances that are close to the target in the feature space but are labeled as the desired target label y_d , as illustrated in the example from Figure 2. Here, the adversary has created a malicious Android app \mathbf{t} , which includes suspicious features (e.g. the WRITE_CONTACTS permission on the left side of the figure), and wishes to prevent a malware detector from flagging this app. The adversary, therefore, selects a benign app \mathbf{x}_b as a base instance. To craft each poison instance, StingRay alters a subset of a base instance’s features so that they resemble those of the target. As shown on the right side of Figure 2, these are not necessarily the most suspicious features, so that the crafted instance will likely be considered benign. Finally, StingRay filters crafted instances based on their negative impact on instances from S' , ensuring that their individual effect on the target classification performance is negligible. The sample crafting procedure is repeated until there are enough instances to trigger the misclassification of \mathbf{t} . Algorithm 1 shows the pseudocode of the attack’s two general-purpose procedures.

We describe concrete implementations of our attack against four existing applications: an image recognition system, an Android malware detector, a Twitter-based exploit predictor, and a data breach predictor. We reimplement the systems that are not publicly available, using the original classification algorithms and the original training sets to reproduce those systems as closely as possible. In total, our applications utilize three classification algorithms—convolutional neural network, linear SVM, and random forest—that have distinct characteristics. This spectrum illustrates the first challenge for our attack: identifying and encapsulating the application-specific steps in StingRay, to adopt a modular design with broad applicability. Making poisoning attacks practical raises additional challenges. For example, a naïve approach would be to inject the target with the desired label into the training set: $h(\mathbf{t}) = y_d$ (**S.I**). However, this is impractical because the adversary, under our threat

Algorithm 1 The StingRay attack.

```
1: procedure STINGRAY( $S^I, Y_{S^I}, \mathbf{t}, y_t, y_d$ )
2:    $I = \emptyset$ 
3:    $h = A^I(S^I)$ 
4:   repeat
5:      $\mathbf{x}_b = \text{GETBASEINSTANCE}(S^I, Y_{S^I}, \mathbf{t}, y_t, y_d)$ 
6:      $\mathbf{x}_c = \text{CRAFTINSTANCE}(\mathbf{x}_b, \mathbf{t})$ 
7:     if  $\text{GETNEGATIVEIMPACT}(S^I, \mathbf{x}_c) < \tau_{NI}$  then
8:        $I = I \cup \{\mathbf{x}_c\}$ 
9:        $h = A^I(S^I \cup I)$ 
10:    until ( $|I| > N_{min}$  and  $h(\mathbf{t}) = y_d$ ) or  $|I| > N_{max}$ 
11:     $PDR = \text{GETPDR}(S^I, Y_{S^I}, I, y_d)$ 
12:    if  $h(\mathbf{t}) \neq y_d$  or  $PDR < \tau_{PDR}$  then
13:      return  $\emptyset$ 
14:    return  $I$ 
15: procedure GETBASEINSTANCE( $S^I, Y_{S^I}, \mathbf{t}, y_t, y_d$ )
16:   for  $\mathbf{x}_b, y_b$  in SHUFFLE( $S^I, Y_{S^I}$ ) do
17:     if  $D(\mathbf{t}, \mathbf{x}_b) < \tau_D$  and  $y_b = y_d$  then
18:       return  $\mathbf{x}_b$ 
```

model, does not control the labeling function. Therefore, GETBASEINSTANCE works by selecting instances \mathbf{x}_b that already have the desired label and are close to the target in the feature space (S.II).

A more sophisticated approach would mutate these samples and use poison instances to push the model boundary toward the target’s class [32]. However, these instances might resemble the target class too much, and they might not receive the desired label from the oracle or even get flagged by an outlier detector. In CRAFTINSTANCE, we apply tiny perturbations to the instances (D.III) and by checking the negative impact NI of crafted poisoning instances on the classifier (D.IV) we ensure they remain *individually inconspicuous*.

Mutating these instances with respect to the target [34] (as illustrated in Figure 1c) may still reduce the overall performance of the classifier (e.g. by causing the misclassification of additional samples similar to the target). We overcome this via GETPDR by checking the performance drop of the attack samples (S.V), therefore ensuring that they remain *collectively inconspicuous*.

Even so, the StingRay attack adds robustness to the poison instances by crafting more instances than necessary, to overcome sampling-based defenses (D.VI). Nevertheless, the attack has a sampling budget that dictates the allowable number of crafted instances (B.VII). A detailed description of StingRay is found in Appendix A.

Attack Constraints. The attack presented above has a series of constraints that shape its effectiveness. Reasoning about them allows us to adapt StingRay to the specific restrictions on each application. These span all three categories identified in Section 3: Success(S.), De-

fense(D.) and Budget(B.):

S.I $h(\mathbf{t}) = y_d$: the desired class label for target

S.II $D(\mathbf{t}, \mathbf{x}_b) < \tau_D$: the inter-instance distance metric

D.III $\bar{s} = \frac{1}{|I|} \sum_{\mathbf{x}_c \in I} s(\mathbf{x}_c, \mathbf{t})$, where $s(\cdot, \cdot)$ is a *similarity* metric: crafting target resemblance

D.IV $NI < \tau_{NI}$: negative impact of poisoning instances

S.V $PDR < \tau_{PDR}$: the perceived performance drop

D.VI $|I| \geq N_{min}$: the minimum number of poison instances

B.VII $|I| \leq N_{max}$: maximum number of poisoning instances

The perceived success of the attacker goals (S.I and S.V) dictate whether the attack is triggered. If the PDR is large, the attack might become indiscriminate and the risk of degrading the overall classifier’s performance is high. The actual PDR could only be computed in the white-box setting. For scenarios with partial knowledge, it is approximated through the perceived PDR on the available classifier.

The impact of crafted instances is influenced by the distance metric and the feature space used to measure instance similarity (S.II). For applications that learn feature representations (e.g. neural nets), the similarity of learned features might be a better choice for minimizing the crafting effort.

The set of features that are actively modified by the attacker in the crafted instances (D.III) defines the *target resemblance* for the attacker, which imposes a trade-off between their inconspicuousness and the effectiveness of the sample. If this quantity is small, the crafted instances are less likely to be perceived as outliers, but a larger number of them is required to trigger the attack. A higher resemblance could also cause the oracle to assign crafted instances a different label than the one desired by the attacker.

The loss difference of a classifier trained with and without a crafted instance (D.IV) approximates the negative impact of that instance on the classifier. It may be easy for an attacker to craft instances with a high negative impact, but these instances may also be easy to detect using existing defenses.

In practice, the cost of injecting instances in the training set can be high (e.g. controlling a network of bots in order to send fake tweets) so the attacker aims to minimize the number of poison instances (D.VI) used in the attack. The adversary might also discard crafted instances that do not have the desired impact on the ML

model. Additionally, some poison instances might be filtered before being ingested by the victim classifier. However, if the number of crafted instances falls below a threshold N_{min} , the attack will not succeed. The maximum number of instances that can be crafted (**B.VII**) influences the outcome of the attack. If the attacker is unable to find sufficient poison samples after crafting N_{max} instances, they might conclude that the large fraction of poison instances in the training set would trigger suspicions or that they depleted the crafting budget.

Delivering Poisoning Instances. The mechanism through which poisoning instances are delivered to the victim classifier is dictated by the application characteristics and the adversarial knowledge. In the most general scenario, the attacker injects the crafted instances alongside existing ones, expecting that the victim classifier will be trained on them. For applications where models are updated over time or trained in mini-batches (such as an image classifier based on neural networks), the attacker only requires control over a subset of such batches and might choose to deliver poison instances through them. In cases where the attacker is unable to create new instances (such as a vulnerability exploit predictor), they will rely on modifying the features of existing ones by poisoning the feature extraction process. The applications we use to showcase StingRay highlight these scenarios and different attack design considerations.

4.1 Bypassing Anti-Poisoning Defenses

In this section, we discuss three defenses against poisoning attacks and how StingRay exploits their limitations.

The Micromodels defense was proposed for cleaning training data for network intrusion detectors [13]. The defense trains classifiers on non-overlapping epochs of the training set (*micromodels*) and evaluates them on the training set. By using a majority voting of the micromodels, training instances are marked as either safe or suspicious. Intuition is that attacks have relatively low duration and they could only affect a few micromodels. It also relies on the availability of accurate instance timestamps.

Reject on Negative Impact (RONI) was proposed against spam filter poisoning attacks [3]. It measures the incremental effect of each individual suspicious training instance and discards the ones with a relatively significant negative impact on the overall performance. RONI sets a threshold by observing the average negative impact of each instance in the training set and flags an instance when its performance impact exceeds the threshold. This threshold determines RONI’s ultimate effectiveness and ability to identify poisoning samples. The defense also requires a sizable clean set for testing instances. We

adapted RONI to a more realistic scenario, assuming no clean holdout set, implementing an iterative variant, as suggested in [41], that incrementally decreases the allowed performance degradation threshold. To the best of our knowledge, this version has not been implemented and evaluated before. However, RONI remains computationally inefficient as the number of trained classifiers scales linearly with the training set.

Target-aware RONI (tRONI) builds on the observation that RONI fails to mitigate *targeted* attacks [34] because the poison instances might not individually cause a significant performance drop. We propose a targeted variant which leverages prior knowledge about a test-time misclassification to determine training instances that might have caused it. While RONI estimates the negative impact of an instance on a holdout set, tRONI considers their effect on the target classification alone. Therefore tRONI is only capable of identifying instances that distort the target classification significantly. A detailed description of this defense is available in the technical report [43].

All these defenses aim to increase adversarial costs by forcing attackers to craft instances that result in a small loss difference (Cost **D.IV**). Therefore, they implicitly assume that poisoning instances stand out from the rest, and they negatively affect the victim classifier. However, attacks such as StingRay could exploit this assumption to evade detection by crafting a small number of inconspicuous poison samples.

4.2 Attack Implementation

We implement StingRay against four applications with distinct characteristics, each highlighting realistic constraints for the attacker. We omit certain technical details for space considerations, encouraging interested readers to consult the technical report [43].

Image classification. We first poison a neural-network (NN) based application for image classification, often used for demonstrating evasion attacks in the prior work. The input instances are images and the labels correspond to objects that are depicted in the image (e.g. airplane, dog, ship). We evaluate StingRay on our own implementation for CIFAR-10 [24]. 10,000 instances (1/6 of the data set) are used for validation and testing. In this scenario, the attacker has an image \mathbf{t} with true label y_t (e.g. a dog) and wishes to trick the model into classifying it as a specific class y_d (e.g. a cat).

We implement a neural network architecture that achieves a performance comparable to other studies [38, 9], obtaining a validation accuracy of 78%. Once the network is trained on the benign inputs, we proceed to poison the classifier. We generate and group poison in-

stances into batches alongside benign inputs. We define $\gamma \in [0, 1]$ to be the *mixing parameter* which controls the number of poison instances in a batch. In our experiments we varied γ over $\{0.125, 0.5, 1.0\}$ (i.e. 4, 16, and 32 instances of the batch are poison) and selected the value that provided the best attack success rate, keeping it fixed across successive updates. We then update³ the previously trained network using these batches until either the attack is perceived as successful or we exceed the number of available poisoning instances, dictated by the cut-off threshold of N_{max} . It is worth noting that if the learning rate is high and the batch contains too many poison instances, the attack could become indiscriminate. Conversely, too few crafted instances would not succeed in changing the target prediction, so the attacker needs to control more batches.

The main insight that motivates our method for generating adversarial samples is that there exist inputs to a network $\mathbf{x}_1, \mathbf{x}_2$ whose distance in pixel space $\|\mathbf{x}_1 - \mathbf{x}_2\|$ is much smaller than their distance in deep feature space $\|H_i(\mathbf{x}_1) - H_i(\mathbf{x}_2)\|$, where $H_i(\mathbf{x})$ is the value of the i^{th} hidden layer’s activation for the input \mathbf{x} . This insight is motivated by the very existence of test-time adversarial examples, where inputs to the classifier are very similar in pixel space, but are successfully misclassified by the neural network [4, 44, 17, 50, 37, 9]. Our attack consists of selecting *base* instances that are close to the target \mathbf{t} in *deep feature space*, but are labeled by the oracle as the attacker’s desired label y_d . CRAFTINSTANCE creates poison images such that the *distance to the target* \mathbf{t} in deep feature space is minimized and the resulting adversarial image is within τ_D distance in pixel space to \mathbf{t} . Recent observations suggest that features in the deeper layers of neural networks are not transferable [52]. This suggests that the selection of the layer index i in the objective function offers a trade-off between attack transferability and the magnitude of perturbations in crafted images (Cost **D.III**). In our experiments we choose H_i to be the third convolutional layer.

We pick 100 target instances uniformly distributed across the class labels. The desired label y_d is the one closest to the true label y_t from the *attacker’s classifier* point of view (i.e. it is the second best guess of the classifier). We set the cut-off threshold $N_{max} = 64$, equivalent to two mini-batches of 32 examples. The perturbation is upper-bounded at $\tau_D < 3.5\%$ resulting in a target resemblance $\bar{s} < 110$ pixels.

Android malware detection. The Drebin Android malware detector [2] uses a linear SVM classifier to predict if an application is malicious or benign. The Drebin data set consists of 123,453 Android apps, including

³ The update is performed on the entire network (i.e. all layers are updated).

target: \mathbf{t} (malicious)	poison: \mathbf{x}_c (benign)
api_call::setWifiEnabled	intent::LAUNCHER
permission::WRITE_CONTACTS	intent::MAIN
permission::ACCESS_WIFI_STATE	permission::ACCESS_WIFI_STATE
permission::READ_CONTACTS	activity::MainActivity
...	permission::READ_CONTACTS
	...

Legend:	Features tagged as suspicious by VT
	Features copied from \mathbf{t} to \mathbf{x}_c

Figure 2: The sample crafting process illustrated for the Drebin Android malware detector. Suspicious features are emphasized in VirusTotal reports using an opaque internal process, but the attacker is not constrained to copying them.

5,560 malware samples. These were labeled using 10 AV engines on VirusTotal [48], considering apps with at least two detections as malicious. The feature space has 545,333 dimensions. We use stratified sampling and split the data set into 60%-40% folds training and testing respectively, aiming to mimic the original classifier. Our implementation achieves 94% F1 on the testing set. The features are extracted from the application archives (APKs) using two techniques. First, from the *AndroidManifest* XML file, which contains meta information about the app, the authors extract the permission requested, the application components and the registered system callbacks. Second, after disassembling the *dex* file, which contains the app bytecode, the system extracts suspicious Android framework calls, actual permission usage and hardcoded URLs. The features are represented as bit vectors, where each index specifies whether the application contains a feature. The adversary aims to misclassify an Android app \mathbf{t} . Although the problems of inducing a targeted false positive (FP) and a targeted false negative (FN) are analogous from the perspective of our definitions, in practice the adversary is likely more interested in targeted FNs, so we focus on this case in our experiments. We evaluate this attack by selecting target instances from the testing set that would be correctly labeled as malicious by the classifier. We then craft instances by adding active features (permissions, API calls, URL requests) from the target to existing benign instances, as illustrated in Figure 2. Each of the crafted apps will have a subset of the features of \mathbf{t} , to remain individually inconspicuous. The poisoning instances are mixed with the pristine ones and used to train the victim classifier from scratch.

We craft 1,717 attacks to test the attack on the Drebin classifier. We use a cutoff threshold $N_{max} = 425$, which corresponds to 0.5% of the training set. The base instances are selected using the Manhattan distance $D = l_1$ and each poisoning instance has a target resemblance of $\bar{s} = 10$ features and a negative impact $\tau_{NI} < 50\%$.

Twitter-based exploit prediction. In [40], the authors present a system, based on a linear SVM, that predicts

which vulnerabilities are going to be exploited using features extracted from Twitter and public vulnerability databases. For each vulnerability, the predictor extracts word-based features (e.g. the number of tweets containing the word *code*), Twitter statistics (e.g. number of distinct users that tweeted about it), and domain-specific features for the vulnerability (e.g. CVSS score). The data set contains 4,140 instances out of which 268 are labeled as positive (a proof-of-concept exploit is publicly available). The classifier uses 72 features from 4 categories: CVSS Score, Vulnerability Database, Twitter traffic and Twitter word features. Due to the class imbalance, we use stratified samples of 60%–40% of the data set for training and testing respectively, obtaining a 40% testing F1.

The targeted attack selects a set I of vulnerabilities that are similar to \mathbf{t} (e.g. same product or vulnerability category), have no known exploits, and gathered fewer tweets. It then proceeds to post crafted tweets about *these* vulnerabilities that include terms normally found in the tweets about the *target* vulnerability. In this manner, the classifier gradually learns that these terms indicate vulnerabilities that are not exploited. However, the attacker’s leverage is limited since the features extracted from sources other than Twitter are not under the attacker’s control.

We simulate 1,932 attacks setting $N_{max} = 20$ and selecting the CVEs to be poisoned using the Euclidean distance $D = l_2$ with $\tau_{NI} < 50\%$.

Data breach prediction. The fourth application we analyze is a data breach predictor proposed in [30]. The system attempts to predict whether an organization is going to suffer a data breach, by using a random forest classifier. The features used in classification include indications of bad IT hygiene (e.g. misconfigured DNS servers) and malicious activity reports (e.g. blacklisting of IP addresses belonging to the organization). These features are absolute values (i.e. organization size), as well as time series based statistics (e.g. duration of attacks). The Data Breach Investigations Reports (DBIR) [47] provides the ground truth. The classifier uses 2,292 instances with 382 positive-labeled examples. The 74 existing features are extracted from externally observable network misconfiguration symptoms as well as blacklisting information about hosts in an organization’s network. A similar technique is used to compute the FICO Enterprise Security Score [15]. We use stratified sampling to build a training set containing 50% of the corpus and use the rest for testing and choosing targets for the attacks. The classifier achieves a 60% F1 score on the testing set.

In this case, the adversary plans to hack an organization \mathbf{t} , but wants to avoid triggering an incident pre-

diction despite the eventual blacklisting of the organization’s IPs. In our simulation, we choose \mathbf{t} from within organizations that were reported in DBIR and were not used at training time, being correctly classified at testing. The adversary chooses a set I of organizations that do not appear in the DBIR and modifies their feature representation. The attacker has limited leverage and is only able to influence time series based features indirectly, by injecting information in various blacklists.

We generate 2,002 attacks under two scenarios: the attacker has compromised a blacklist and is able to influence the features of many organizations, or the attacker has infiltrated a few organizations and it uses them to modify their reputation on all the blacklists. We set $N_{max} = 50$ and the instances to be poisoned are selected using the Euclidean distance $D = l_2$ with $\tau_{NI} < 50\%$.

4.3 Practical Considerations

Running time of StingRay. The main computational expenses of StingRay are: crafting the instances in CRAFTINSTANCE, computing the distances to the target in GETBASEINSTANCE, and measuring the negative impact of the crafted instances in GETNEGATIVEIMPACT.

CRAFTINSTANCE depends on the crafting strategy and its complexity in searching for features to perturb. For the image classifier, we adapt an existing evasion attack, showing that we could reduce the computational cost by finding adversarial examples on hidden layers instead of the output layer. For all other applications we evaluated, the choice of features is determined in constant time.

The GETBASEINSTANCE procedure computes inter-instance distances once per attack, and it is linear in terms of the attackers training set size for a particular label. For larger data sets the distance computation could be approximated (e.g. using a low-rank approximation).

In GETNEGATIVEIMPACT, we obtain a good approximation of the negative impact (NI) by training locally-accurate classifiers on small instance subsets and performing the impact test on batches of crafted instances.

Labeling poisoning instances. Our attacker model assumes that the adversary does not control the oracle used for labeling poisoning instances. Although the attacker could craft poisoning instances that closely resemble the target \mathbf{t} to make them more powerful, they could be flagged as outliers or the oracle could assign them a label that is detrimental for the attack. It is therefore beneficial to reason about the oracles specific to all applications and the mechanisms used by StingRay to obtain the desired labels.

For the image classifier, the most common oracle is a consensus of human analysts. In an attempt to map

the effect of adversarial perturbations on human perception, the authors of [35] found through a user study that the maximum fraction of perturbed pixels at which humans will correctly label an image is 14%. We, therefore, designed our experiments to remain within these bounds. Specifically, we measure the pixel space perturbation as the l_∞ distance and discard poison samples with $\tau_D > 0.14$ prior to adding them to I .

The Drebin classifier uses VirusTotal as the oracle. In our experiments, the poison instances would need to maintain the benign label. We systematically create over 19,000 Android applications that correspond to attack instances and utilize VirusTotal, in the same way as Drebin does, to label them. To modify selected features of the Android apps, we reverse-engineer Drebin’s feature extraction process to generate apps that would have the desired feature representation. We generate these applications for the scenario where only the subset of features extracted from the *AndroidManifest* are modifiable by the attacker, similar to prior work [19]. In 89% of these cases, the crafted apps bypassed detection, demonstrating the feasibility of our strategy in obtaining negatively labeled instances. However, in our attack scenario, we assume that the attacker is not consulting the oracle, releasing all crafted instances as part of the attack.

For the exploit predictor, labeling is performed independently of the feature representations of instances used by the system. The adversary manipulates the public discourse around existing vulnerabilities, but the label exists with respect to the availability of an exploit. Therefore the attacker has more degrees of freedom in modifying the features of instances in I , knowing that their desired labels will be preserved.

In case of the data breach predictor, the attacker utilizes organizations with no known breach and aims to poison the blacklists that measure their hygiene, or hacks them directly. In the first scenario, the attacker does not require access to an organization’s networks, therefore the label will remain intact. The second scenario would be more challenging, as the adversary would require extra capabilities to ensure they remain stealthy while conducting the attack.

5 Evaluation

We start by evaluating weaker evasion and poisoning adversaries, within the **FAIL** model, on the image and malware classifiers (Section 5.1). Then, we evaluate the effectiveness of existing defenses against StingRay (Section 5.2) and its applicability to a larger range of classifiers. Our evaluation seeks to answer four research questions: *How could we systematically evaluate the transferability of existing evasion attacks? What are the limitations of realistic poisoning adversaries? When are tar-*

geted poison samples transferable? Is StingRay effective against multiple applications and defenses? We quantify the effectiveness of the evasion attack using the percentage of successful attacks (SR), while for StingRay we also measure the Performance Drop Ratio (PDR). We measure the PDR on holdout testing sets by considering either the average accuracy, on applications with balanced data sets, or the average F1 score (the harmonic mean between precision and recall), which is more appropriate for highly imbalanced data sets.

5.1 FAIL Analysis

In this section, we evaluate the image classifier and the malware detector using the **FAIL** framework. The model allows us to utilize both a state of the art evasion attack as well as StingRay for the task. To control for additional confounding factors when evaluating StingRay, in this analysis we purposely omit the negative impact-based pruning phase of the attack. We chose to implement the FAIL analysis on the two applications as they do not present built-in leverage limitations and they have distinct characteristics.

Evasion attack on the image classifier. The first attack subjected to the FAIL analysis is JSMA [35], a well-known targeted evasion attack. Transferability of this attack has previously been studied only for an adversary with limited knowledge along the **A** and **I** dimensions [37]. We attempt to reuse an application configuration similar in prior work, implementing our own 3-layer convolutional neural network architecture for the MNIST handwritten digit data set [26]. The validation accuracy of our model is 98.95%. In table 3, we present the average results of our 11 experiments, each involving 100 attacks.

For each experiment, the table reports the Δ variation of the **FAIL** dimension investigated, two SR statistics: *perceived* (as observed by the attacker on their classifier) and *potential* (the effect on the victim if all attempts are triggered by the attacker) as well as the mean perturbation $\bar{\tau}_D$ introduced to the evasion instances.

Experiment #6 corresponds to the white-box adversary, where we observe that the white-box attacker could reach 80% SR.

Experiments #1–2 model the scenario in which the attacker has limited **Feature knowledge**. Realistically, these scenarios can simulate an evasion or poisoning attack against a self-driving system, conducted without knowing the vehicle’s camera angles—wide or narrow. We simulate this by cropping a frame of 3 and 6 pixels from the images, decreasing the available features by 32% and 62%, respectively. The attacker uses the cropped images for training and testing the classifier, as

#	Δ	SR %	$\bar{\epsilon}_D$	Δ	SR %	PDR	Instances
FAIL:Unknown features							
1	32%	67/3	0.070	39%	87/63/67	0.93/0.96/0.96	8/4/10
2	62%	86/7	0.054	66%	84/71/74	0.94/0.95/0.95	8/4/9
FAIL:Unknown algorithm							
3	shallow	99/10	0.035	shallow	83/65/68	0.97/0.97/0.96	17/14/15
4	narrow	82/20	0.027	narrow	75/67/72	0.96/0.97/0.96	20/16/17
FAIL:Unavailable training set							
5	35000	93/18	0.032	35000	73/68/76	0.97/0.96/0.96	17/16/14
6	50000	80/80	0.026	50000	78/70/74	0.97/0.97/0.97	18/16/15
FAIL:Unknown training set							
7	45000	90/18	0.029	45000	82/69/74	0.98/0.96/0.96	16/10/15
8	50000	96/19	0.034	50000	70/62/68	0.95/0.96/0.96	17/8/17
FAIL:Read-only features							
9	18%	80/4	0.011	25%	80/70/72	0.97/0.97/0.97	19/16/15
10	41%	80/34	0.022	50%	80/71/76	0.97/0.97/0.97	18/16/13
11	62%	80/80	0.026	75%	83/78/79	0.97/0.97/0.96	16/16/12

Table 3: JSMA on the image classifier

Table 4: StingRay on the image classifier

Δ	SR %	PDR	Instances
109066	79/3/5	0.99/0.99/1.00	73/50/53
327199	77/12/13	0.99/0.99/1.00	51/50/15
SGD	42/33/42	0.99/0.99/0.99	65/50/31
dSVM	38/35/48	0.99/0.99/0.99	78/50/61
8514	69/27/27	0.90/0.99/0.99	57/50/42
85148	50/50/50	0.99/0.99/0.99	77/50/61
8514	53/21/24	0.93/0.99/1.00	62/50/49
43865	36/29/39	1.04/0.99/0.99	100/50/87
851	73/12/13	0.67/0.99/1.00	50/50/10
8514	49/16/17	0.90/0.99/1.00	61/50/47
85148	32/32/32	0.99/0.99/0.99	79/50/57

Table 5: StingRay on the malware classifier

Tables 3, 4, 5: FAIL analysis of the two applications. For each JSMA experiment, we report the attack SR (perceived/potential), as well as the mean perturbation $\bar{\epsilon}_D$ introduced to the evasion instances. For each StingRay experiment, we report the SR and PDR (perceived/actual/potential), as well as statistics for the crafted instances on successful attacks (mean/median/standard deviation). Δ represents the variation of the **FAIL** dimension investigated.

well as for crafting instances. On the victim classifier, the cropped part of the images is added back without altering the perturbations.

With limited knowledge along this dimension (#1-2) the perceived success remains high, but the actual SR is very low. This suggests that the evasion attacks are very sensitive in such scenarios, highlighting a potential direction for future defenses.

We then model an attacker with limited **Algorithm knowledge**, possessing a similar architecture, but with smaller network capacity. For the shallow network (#3) the attacker network has one less hidden layer; the narrow architecture (#4) has half of the original number of neurons in the fully connected hidden layers. Here we observe that the shallow architecture (#3) renders almost all attacks as successful on the attacker. However, the potential SR on the victim is higher for the narrow setup (#4). This contradicts claims in prior work [37], which state that the used architecture is not a factor for success. **Instance knowledge.** In #5 we simulate a scenario in which the attacker only knows 70% of the victim training set, while #7-8 model an attacker with 80% of the training set available and an additional subset of instances sampled from the same distribution.


These results might help us explain the contradiction with prior work. Indeed, we observe that a robust attacker classifier, trained on a sizable data set, reduces the SR to 19%, suggesting that the attack success sharply declines with fewer victim training instances available. In contrast, in [37] the SR remains at over 80% because of the non-random data-augmentation technique used to build the attacker training set. As a result, the attacker model is a closer approximation of the victim one, impacting the analysis along the **A** dimension.

Experiments #9–11 model the case where the attacker has limited **Leverage** and is unable to modify some of the instance features. This could represent a region where watermarks are added to images to check their integrity. We simulate it by considering a border in the image from which the modified pixels would be discarded, corresponding to the attacker being able to modify to 18%, 41% and 62% of an image respectively. We observe a significant drop in transferability, although #11 shows that the SR is not reduced with leverage above a certain threshold.

StingRay on the image classifier. We now evaluate the poisoning attack described in 4.2 under the same scenarios defined above. Table 4 summarizes our results. In contrast to evasion, the table reports the SR, PDR, and the number of poison instances needed. Here, besides the *perceived* and *potential* statistics, we also report the *actual* SR and PDR (as reflected on the victim when triggering only the attacks perceived successful).

For limited **Feature knowledge**, we observe that the perceived SR is over 84% but the actual success rate drops significantly on the victim. However, the actual SR for #2 is similar to the white-box attacker (#6), showing that features derived from the exterior regions of an image are less specific to an instance. This suggests that although reducing feature knowledge decreases the effectiveness of StingRay, the specificity of some known features may still enable successful attacks.

Along the **A** dimension, we observe that both architectures allow the attacker to accurately approximate the deep space distance between instances. While the perceived SR is overestimated, the actual SR of these attacks is comparable to the white-box attack, showing that ar-



(a) Limited **Feature** knowledge. (b) Limited **Leverage**.

Figure 3: Example of original and crafted images. Images in the left panel are crafted with 39% and 66% of features unknown. In the right panel, the images are crafted with 100% and 50% leverage.

chitecture secrecy does not significantly increase the resilience against these attacks. The open-source neural network architectures readily available for many of classification tasks would aid the adversary. Along the **I** dimension, in #5, the PDR is increased because the smaller available training set size prevents them from training a robust classifier. In the white-box attack #6 we observe that the perceived, actual and potential SRs are different. We determined that this discrepancy is caused by documented nondeterminism in the implementation framework. This affects the order in which instances are processed, causing variance on the model parameters, which in turns reflects on the effectiveness of poisoning instances. Nevertheless, we observe that the potential SR is higher in #5, even though the amount of available information is larger in #6. This highlights the benefit of a fine-grained analysis along all dimensions, since the attack success rate may not be monotonic in terms of knowledge levels.

Surprisingly, we observe that the actual SR for #8, where the attacker has more training instances at their disposal, is lower than for #7. This is likely caused by the fact that, with a larger discrepancy between the training sets of the victim and the attacker classifier, the attacker is more likely to select base instances that would not be present in the victim training set. After poisoning the victim, the effect of crafted instances would not be bootstrapped by the base instances, and the attacker fails. The results suggest that the attack is sensitive to the presence of specific pristine instances in the training set, and variance in the model parameters could mitigate the threat. However, determining which instances should be kept secret is subject for future research.

Limited **Leverage** increases the actual SR beyond the white-box attack. When discarding modified pixels, the overall perturbation is reduced. Thus, it is more likely that the poison samples will become collectively inconspicuous, increasing the attack effectiveness. Figure 3 illustrates some images crafted by constrained adversaries.

The FAIL analysis results show that the perceived PDR is generally an accurate representation of the actual value, making it easy for the adversary to assess the instance inconspicuousness and indiscriminate damage caused by the attack. The attacks transfer surprisingly well from the attacker to the victim, and a significant

number of failed attacks would potentially be successful if triggered on the victim. We observe that limited leverage allows the attacker to localize their strategy, crafting attack instances that are even *more successful* than the white-box attack.

StingRay on the malware classifier. In order to evaluate StingRay in the FAIL setting on the malware classifier, we trigger all 1,717 attacks described in 4.2 along 11 dimensions. Table 5 summarizes the results. Experiment #6 corresponds to the white-box attacker.

Experiments #1–2 look at the case where **Features** are unknown to the adversary. In this case, the surrogate model used to craft poison instances includes only 20% and 60% of the features respectively. Surprisingly, the attack is highly ineffective. Although the attacker perceives the attack as successful in some cases, the classifier trained on the available feature subspace is a very inaccurate approximation of the original one, resulting in an actual SR of at most 12%. These results echo these from evasion, indicating that features secrecy might prove a viable lead towards effective defenses. We also investigate adversaries with various degrees of knowledge about the classification **Algorithm**. Experiment #3 trains a linear model using the Stochastic Gradient Descent (SGD) algorithm, and in #4 (dSVM), the hyperparameters of the SVM classifier are not known by the attacker. Compared with the original Drebin SVM classifier, the default setting in #4 uses a larger regularization parameter. This suggests that regularization can help mitigate the impact of individual poison instances, but the adversary may nevertheless be successful by injecting more crafted instances in the training set.

Instance knowledge. Experiments #5–6 look at a scenario in which the known instances are subsets of S^* . Unsurprisingly, the attack is more effective as more instances from S^* become available. The attacker’s inability to train a robust surrogate classifier is reflected through the large perceived PDR. For experiments #7–8, victim training instances are not available to the attacker, their classifier being trained on samples from the same underlying distribution as S^* . Under these constraints, the adversary could only approximate the effect of the attack on the targeted classifier. Additionally, the training instances might be significantly different than the base instances available to the adversary, canceling the effect of crafted instances. The results show, as in the case of the image classifier, that poison instances are highly dependent on other instances present in the training set to bootstrap their effect on target misclassification. We further look at the impact of limited **Leverage** on the attack effectiveness. Experiments #9–11 look at various training set sizes for the case where only the features extracted from *AndroidManifest.xml* are modifi-

	StingRay	RONI	tRONI	MM
	I /SR%/PDR	Fix%/PDR		
Images	16/70/0.97	-/-	-/-	-/-
Malware	77/50/0.99	0/0.98	15/0.98	-/-
Exploits	7/6/1.00	0/0.97	40/0.67	0/0.33
Breach	18/34/0.98	-/-	20/0.96	55/0.91

Table 6: Effectiveness of StingRay and of existing defenses against it on all applications. Each attack cell reports the average number of poison instances $|I|$, the SR and actual PDR. Each defense cell reports the percentage of fixed attacks and the PDR after applying it.

able. These features correspond to approximately 40% of the 545,333 existing features. Once again, we observe that the effectiveness of a constrained attacker is reduced. This signals that a viable defense could be to extract features from uncorrelated sources, which would limit the leverage of such an attacker.

The FAIL analysis on the malware classifier reveals that the actual drop in performance of the attacks is insignificant on all dimensions, but the attack effectiveness is generally decreased for weaker adversaries. However, feature secrecy and limited leverage appear to have the most significant effect on decreasing the success rate, hinting that they might be a viable defense.

5.2 Effectiveness of StingRay

In this section we explore the effectiveness of StingRay across all applications described in 4.2 and compare existing defense mechanisms in terms of their ability to prevent the targeted mispredictions. Table 6 summarizes our findings. Here we only consider the strongest (white-box) adversary to determine upper bounds for the resilience against attacks, without assuming any degree of secrecy.

Image classifier. We observe that the attack is successful in 70% of the cases and yields an average PDR of 0.97, requiring an average of 16 instances. Upon further analysis, we discovered that the performance drop is due to other testing instances similar to the target being misclassified as y_d . By tuning the attack parameters (e.g. the layer used for comparing features or the degree of allowed perturbation) to generate poison instances that are more specific to the target, the performance drop on the victim could be further reduced at the expense of requiring more poisoning instances. Nevertheless, this shows that neural nets define a fine-grained boundary between class-targeted and instance-targeted poisoning attacks and that it is not straightforward to discover it, even with complete adversarial knowledge.

None of the three poisoning defenses are applicable on this task. RONI and tRONI require training over 50,000 classifiers for each level of inspected negative impact.

This is prohibitive for neural networks which are known to be computationally intensive to train. Since we could not determine reliable timestamps for the images in the data set, MM was not applicable either.

Malware classifier. StingRay succeeds in half of the cases and yields a negligible performance drop on the victim. The attack being cut off by the crafting budget on most failures (Cost **B.VII**) suggests that some targets might be too "far" from the class separator and that moving this separator becomes difficult. Nevertheless, understanding what causes this hardness remains an open question.

On defenses, we observe that RONI often fails to build correctly-predicting folds on Drebin and times out. Hence, we investigate the defenses against only 97 successful attacks for which RONI did not timeout. MM rejects all training instances while RONI fails to detect any attack instances. tRONI detects very few poison instances, fixing only 15% of attacks, as they do not have a large negative impact, individually, on the misclassification of the target. None of these defenses are able to fix a large fraction of the induced mispredictions.

Exploit predictor. While poisoning a small number of instances, the attack has a very low success rate. This is due to the fact that the non-Twitter features are not modifiable; if the data set does not contain other vulnerabilities similar to the target (e.g. similar product or type), the attack would need to poison more CVEs, reaching N_{max} before succeeding. The result, backed by our FAIL analysis of the other linear classifier in Section 5.1, highlights the benefits of built-in leverage limitations in protecting against such attacks.

MM correctly identifies the crafted instances but also marks a large fraction of positively-labeled instances as suspicious. Consequently, the PDR on the classifier is severely affected. In instances where it does not timeout, RONI fails to mark any instance. Interestingly, tRONI marks a small fraction of attack instances which helps correct 40% of the predictions but still hurting the PDR. The partial success of tRONI is due to two factors: the small number of instances used in the attack and the limited leverage for the attacker, which boosts the negative impact of attack instances through resampling. We observed that due to variance, the negative impact computed by tRONI is larger than the one perceived by the attacker for discovered instances. The adversary could adapt by increasing the confidence level of the statistic that reflects negative impact in the StingRay algorithm.

Data breach predictor: The attacks for this application correspond to two scenarios, one with limited leverage over the number of time series features. Indeed, the one in which the attacker has limited leverage has an SR of 5%, while the other one has an SR of 63%. This corroborates our observation of the impact of adversarial

leverage for the exploit prediction. RONI fails due to consistent timeouts in training the random forest classifier. tRONI fixes 20% of the attacks while decreasing the PDR slightly. MM is a natural fit for the features based on time series and is able to build more balanced voting folds. The defense fixes 55% of mispredictions, at the expense of lowering the PDR to 0.91.

Our results suggest that StingRay is practical against a variety of classification tasks—even with limited degrees of leverage. Existing defenses, where applicable, are easily bypassed by lowering the required negative impact of crafted instances. However, the reduced attack success rate on applications with limited leverage suggests new directions for future defenses.

6 Related Work

Several studies proposed ways to model adversaries against machine learning systems. [25] proposes *FTC*—*features, training set, and classifier*, a model to define an attacker’s knowledge and capabilities in the case of a practical evasion attack. Unlike the FTC model, the FAIL model is evaluated on both test- and training-time attacks, enables a fine-grained analysis of the dimensions and includes Leverage. These characteristics allow us to better understand how the **F** and **L** dimensions influence the attack success. Furthermore, [27, 7] introduce game theoretical Stackelberg formulations for the interaction between the adversary and the data miner in the case of data manipulations. Adversarial limitations are also discussed in [22]. Several attacks against machine learning consider adversaries with varying degrees of knowledge, but they do not cover the whole spectrum [4, 35, 37]. Recent studies investigate transferability, in attack scenarios with limited knowledge about the target model [36, 28, 9]. The FAIL model unifies these dimensions and can be used to model these capabilities systematically across multiple attacks under realistic assumptions about adversaries. Unlike game theoretical approaches, FAIL does not assume perfect knowledge on either the attacker or the defender. By defining a wider spectrum of adversarial knowledge, FAIL generalizes the notion of transferability.

Prior work introduced indiscriminate and targeted poisoning attacks. For indiscriminate poisoning, a spammer can force a Bayesian filter to misclassify legitimate emails by including a large number of dictionary words in spam emails, causing the classifier to learn that all tokens are indicative of spam [3]. An attacker can degrade the performance of a Twitter-based exploit predictor by posting fraudulent tweets that mimic most of the features of informative posts [40]. One could also the damage overall performance of an SVM classifier by injecting a small volume of crafted attack points [5]. For targeted

poisoning, a spammer can trigger the filter against a specific legitimate email by crafting spam emails resembling the target [34]. This was also studied in the healthcare field, where an adversary can subvert the predictions for a whole target class of patients by injecting fake patient data that resembles the target class [32]. StingRay is a model-agnostic targeted poisoning attack and works on a broad range of applications. Unlike existing targeted poisoning attacks, StingRay aims to bound indiscriminate damage to preserve the overall performance.

On neural networks, [23] proposes a targeted poisoning attack that modifies training instances which have a strong influence on the target loss. In [51], the poisoning attack is a white-box indiscriminate attack adapted from existing evasion work. Furthermore, [29] and [20] introduce backdoor and trojan attacks where adversaries cause the classifiers to misbehave when a trigger is present in the input. The targeted poisoning attack proposed in [11] requires the attacker to assign labels to crafted instances. Unlike these attacks, StingRay does not require white-box or query access the original model. Our attack does not require control over the labeling function or modifications to the target instance.

7 Discussion

The vulnerability of ML systems to evasion and poisoning attacks leads to an arms race, where defenses that seem promising are quickly thwarted by new attacks [17, 37, 38, 9]. Previous defenses make implicit assumptions about how the adversary’s capabilities should be constrained to improve the system’s resilience to attacks. The FAIL adversary model provides a framework for exposing and systematizing these assumptions. For example, the feature squeezing defense [49] constrains the adversary along the **A** and **F** dimensions by modifying the input features and adding an adversarial example detector. Similarly, RONI constrains the adversary along the **I** dimension by sanitizing the training data. The ML-based systems employed in the security industry [21, 10, 39, 12], often rely on undisclosed features to render attacks more difficult, thus constraining the **F** dimension. In Table 2 we highlight implicit and explicit assumptions of previous defenses against poisoning and evasion attacks.

Through our systematic exploration of the FAIL dimensions, we provide the first experimental comparison of the importance of these dimensions for the adversary’s goals, in the context of targeted poisoning and evasion attacks. For a linear classifier, our results suggest that feature secrecy is the most promising direction for achieving attack resilience. Additionally, reducing leverage can increase the cost for the attacker. For a neural network based image recognition system, our results suggest that

StingRay’s samples are transferable across all dimensions. Interestingly, limiting the leverage causes the attacker to craft instances that are more potent in triggering the attack. We also observed that secrecy of training instances provides limited resilience.

Furthermore, we demonstrated that the **FAIL** adversary model is applicable to targeted evasion attacks as well. By systemically capturing an adversary’s knowledge and capabilities, the FAIL model also defines a more general notion of attack transferability. In addition to investigating transferability under certain dimensions, such as the **A** dimension in [9] or **A** and **I** dimensions in [37], generalized transferability covers a broader range of adversaries. At odds with the original findings in [37], our results suggest a lack of generalized-transferability for a state of the art evasion attack; while highlighting feature secrecy as the most prominent factor in reducing the attack success rate. Future research may utilize this framework as a vehicle for reasoning about the most promising directions for defending against other attacks.

Our results also provide new insights for the broader debate about the generalization capabilities of neural networks. While neural networks have dramatically reduced test-time errors for many applications, which suggests they are capable of generalization (e.g. by learning meaningful features from the training data), recent work [53] has shown that neural networks can also memorize randomly-labeled training data (which lack meaningful features). We provide a first step toward understanding the extent to which an adversary can exploit this behavior through targeted poisoning attacks. Our results are consistent with the hypothesis that an attack, such as StingRay, can force selective memorization for a target instance while preserving the generalization capabilities of the model. We leave testing this hypothesis rigorously for future work.

8 Conclusions

We introduce the **FAIL** model, a general framework for evaluating realistic attacks against machine learning systems. We also propose StingRay, a targeted poisoning attack designed to bypass existing defenses. We show that our attack is practical for 4 classification tasks, which use 3 different classifiers. By exploring the **FAIL** dimensions, we observe new transferability properties in existing targeted evasion attacks and highlight characteristics that could provide resiliency against targeted poisoning. This exploration generalizes the prior work on attack transferability and provides new results on the transferability of poison samples.

Acknowledgments We thank Ciprian Baetu, Jonathan Katz, Daniel Marcu, Tom Goldstein, Michael Maynard, Ali Shafahi, W. Ronny Huang, our shepherd, Patrick McDaniel and the anonymous reviewers for their feedback. We also thank the Drebin authors for giving us access to their data set and VirusTotal for access to their service. This research was partially supported by the Department of Defense and the Maryland Procurement Office (contract H98230-14-C-0127).

References

- [1] ALEXEY MALANOV 12 POSTS MALWARE EXPERT, ANTI-MALWARE TECHNOLOGIES DEVELOPMENT, K. L. The multilayered security model in kaspersky lab products, Mar 2017.
- [2] ARP, D., SPREITZENBARTH, M., HUBNER, M., GASCON, H., AND RIECK, K. Drebin: Effective and explainable detection of android malware in your pocket. In *NDSS* (2014).
- [3] BARRENO, M., NELSON, B., JOSEPH, A. D., AND TYGAR, J. D. The security of machine learning. *Machine Learning 81* (2010), 121–148.
- [4] BIGGIO, B., CORONA, I., MAIORCA, D., NELSON, B., ŠRNDIĆ, N., LASKOV, P., GIACINTO, G., AND ROLI, F. Evasion attacks against machine learning at test time. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases* (2013), Springer, pp. 387–402.
- [5] BIGGIO, B., NELSON, B., AND LASKOV, P. Poisoning attacks against support vector machines. *arXiv preprint arXiv:1206.6389* (2012).
- [6] BOJARSKI, M., YERES, P., CHOROMANSKA, A., CHOROMANSKI, K., FIRNER, B., JACKEL, L., AND MULLER, U. Explaining how a deep neural network trained with end-to-end learning steers a car. *arXiv preprint arXiv:1704.07911* (2017).
- [7] BRÜCKNER, M., AND SCHEFFER, T. Stackelberg games for adversarial prediction problems. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining* (2011), ACM, pp. 547–555.
- [8] CARLINI, N., AND WAGNER, D. Adversarial examples are not easily detected: Bypassing ten detection methods. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security* (2017), ACM, pp. 3–14.
- [9] CARLINI, N., AND WAGNER, D. Towards evaluating the robustness of neural networks. In *Security and Privacy (SP), 2017 IEEE Symposium on* (2017), IEEE, pp. 39–57.
- [10] CHAU, D. H. P., NACHENBERG, C., WILHELM, J., WRIGHT, A., AND FALOUTSOS, C. Polonium : Tera-scale graph mining for malware detection. In *SIAM International Conference on Data Mining (SDM)* (Mesa, AZ, April 2011).
- [11] CHEN, X., LIU, C., LI, B., LU, K., AND SONG, D. Targeted Backdoor Attacks on Deep Learning Systems Using Data Poisoning. *ArXiv e-prints* (Dec. 2017).
- [12] COLVIN, R. Stranger danger - introducing smartscreen application reputation. <http://blogs.msdn.com/b/ie/archive/2010/10/13/stranger-danger-introducing-smartscreen-application-reputation.aspx>, Oct 2010.
- [13] CRETU, G. F., STAVROU, A., LOCASO, M. E., STOLFO, S. J., AND KEROMYTIS, A. D. Casting out demons: Sanitizing training data for anomaly sensors. In *Security and Privacy, 2008. SP 2008. IEEE Symposium on* (2008), IEEE, pp. 81–95.

- [14] ERNST YOUNG LIMITED. The future of underwriting. <http://www.ey.com/us/en/industries/financial-services/insurance/ey-the-future-of-underwriting>, 2015.
- [15] FAIR ISAAC CORPORATION. FICO enterprise security score gives long-term view of cyber risk exposure, November 2016. <http://www.fico.com/en/newsroom/fico-enterprise-security-score-gives-long-term-view-of-cyber-risk-exposure-10-27-2016>.
- [16] GILAD-BACHRACH, R., DOWLIN, N., LAINE, K., LAUTER, K., NAEHRIG, M., AND WERNING, J. Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. In *International Conference on Machine Learning* (2016), pp. 201–210.
- [17] GOODFELLOW, I. J., SHLENS, J., AND SZEGEDY, C. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014).
- [18] GOOGLE RESEARCH BLOG. Assisting pathologists in detecting cancer with deep learning. <https://research.googleblog.com/2017/03/assisting-pathologists-in-detecting.html>, Mar 2017.
- [19] GROSSE, K., PAPERNOT, N., MANOHARAN, P., BACKES, M., AND MCDANIEL, P. Adversarial perturbations against deep neural networks for malware classification. *arXiv preprint arXiv:1606.04435* (2016).
- [20] GU, T., DOLAN-GAVITT, B., AND GARG, S. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733* (2017).
- [21] HEARN, M. Abuse at scale. In *RIPE 64* (Ljubljana, Slovenia, Apr 2012). <https://ripe64.ripe.net/archives/video/25/>.
- [22] HUANG, L., JOSEPH, A. D., NELSON, B., RUBINSTEIN, B. I., AND TYGAR, J. Adversarial machine learning. In *Proceedings of the 4th ACM workshop on Security and artificial intelligence* (2011), ACM, pp. 43–58.
- [23] KOH, P. W., AND LIANG, P. Understanding black-box predictions via influence functions. *arXiv preprint arXiv:1703.04730* (2017).
- [24] KRIZHEVSKY, A., AND HINTON, G. Learning multiple layers of features from tiny images. *CiteSeer* (2009).
- [25] LASKOV, P., ET AL. Practical evasion of a learning-based classifier: A case study. In *Security and Privacy (SP), 2014 IEEE Symposium on* (2014), IEEE, pp. 197–211.
- [26] LECUN, Y. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/> (1998).
- [27] LIU, W., AND CHAWLA, S. A game theoretical model for adversarial learning. In *Data Mining Workshops, 2009. ICDMW'09. IEEE International Conference on* (2009), IEEE, pp. 25–30.
- [28] LIU, Y., CHEN, X., LIU, C., AND SONG, D. Delving into transferable adversarial examples and black-box attacks. *arXiv preprint arXiv:1611.02770* (2016).
- [29] LIU, Y., MA, S., AAFER, Y., LEE, W.-C., ZHAI, J., WANG, W., AND ZHANG, X. Trojaning attack on neural networks. Tech. Rep. 17-002, Purdue University, 2017.
- [30] LIU, Y., SARABI, A., ZHANG, J., NAGHIZADEH, P., KARIR, M., BAILEY, M., AND LIU, M. Cloudy with a chance of breach: Forecasting cyber security incidents. In *24th USENIX Security Symposium (USENIX Security 15)* (2015), pp. 1009–1024.
- [31] MIT TECHNOLOGY REVIEW. How to upgrade judges with machine learning. <https://www.technologyreview.com/s/603763/how-to-upgrade-judges-with-machine-learning/>, Mar 2017.
- [32] MOZAFFARI-KERMANI, M., SUR-KOLAY, S., RAGHUNATHAN, A., AND JHA, N. K. Systematic poisoning attacks on and defenses for machine learning in healthcare. *IEEE Journal of biomedical and health informatics* 19, 6 (2015), 1893–1905.
- [33] MUÑOZ-GONZÁLEZ, L., BIGGIO, B., DEMONTIS, A., PAUDICE, A., WONGRASSAMEE, V., LUPU, E. C., AND ROLI, F. Towards poisoning of deep learning algorithms with back-gradient optimization. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security* (2017), ACM, pp. 27–38.
- [34] NELSON, B., BARRENO, M., CHI, F. J., JOSEPH, A. D., RUBINSTEIN, B. I. P., SAINI, U., SUTTON, C., TYGAR, J. D., AND XIA, K. Exploiting machine learning to subvert your spam filter. In *Proceedings of the 1st Usenix Workshop on Large-Scale Exploits and Emergent Threats* (Berkeley, CA, USA, 2008), LEET'08, USENIX Association, pp. 7:1–7:9.
- [35] PAPERNOT, N., MCDANIEL, P., JHA, S., FREDRIKSON, M., CELIK, Z. B., AND SWAMI, A. The limitations of deep learning in adversarial settings. In *2016 IEEE European Symposium on Security and Privacy (EuroS&P)* (2016), IEEE, pp. 372–387.
- [36] PAPERNOT, N., MCDANIEL, P. D., AND GOODFELLOW, I. J. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *CoRR abs/1605.07277* (2016).
- [37] PAPERNOT, N., MCDANIEL, P. D., GOODFELLOW, I. J., JHA, S., CELIK, Z. B., AND SWAMI, A. Practical black-box attacks against deep learning systems using adversarial examples. In *ACM Asia Conference on Computer and Communications Security* (Abu Dhabi, UAE, 2017).
- [38] PAPERNOT, N., MCDANIEL, P. D., WU, X., JHA, S., AND SWAMI, A. Distillation as a defense to adversarial perturbations against deep neural networks. In *IEEE Symposium on Security and Privacy* (2016), IEEE Computer Society, pp. 582–597.
- [39] RAJAB, M. A., BALLARD, L., LUTZ, N., MAVROMMATIS, P., AND PROVOS, N. CAMP: Content-agnostic malware protection. In *Network and Distributed System Security (NDSS) Symposium* (San Diego, CA, Feb 2013).
- [40] SABOTTKE, C., SUCIU, O., AND DUMITRA, T. Vulnerability disclosure in the age of social media: exploiting twitter for predicting real-world exploits. In *24th USENIX Security Symposium (USENIX Security 15)* (2015), pp. 1041–1056.
- [41] SAINI, U. Machine learning in the presence of an adversary: Attacking and defending the spambayes spam filter. Tech. rep., DTIC Document, 2008.
- [42] STEINHARDT, J., KOH, P. W. W., AND LIANG, P. S. Certified defenses for data poisoning attacks. In *Advances in Neural Information Processing Systems* (2017), pp. 3520–3532.
- [43] SUCIU, O., MĂRGINEAN, R., KAYA, Y., DAUMÉ III, H., AND DUMITRAȘ, T. When does machine learning fail? generalized transferability for evasion and poisoning attacks. *arXiv preprint arXiv:1803.06975* (2018).
- [44] SZEGEDY, C., ZAREMBA, W., SUTSKEVER, I., BRUNA, J., ERHAN, D., GOODFELLOW, I., AND FERGUS, R. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199* (2013).
- [45] TAMERSOY, A., ROUNDY, K., AND CHAU, D. H. Guilt by association: large scale malware detection by mining file-relation graphs. In *KDD* (2014).
- [46] TRAMÈR, F., ZHANG, F., JUELS, A., REITER, M., AND RISTENPART, T. Stealing machine learning models via prediction APIs. In *25th USENIX Security Symposium (USENIX Security 16)* (Austin, TX, Aug. 2016), USENIX Association.

- [47] VERIZON. Data breach investigations reports (dbir), February 2012. <http://www.verizonenterprise.com/DBIR/>.
- [48] VIRUSTOTAL. <http://www.virustotal.com>.
- [49] XU, W., EVANS, D., AND QI, Y. Feature squeezing: Detecting adversarial examples in deep neural networks. *arXiv preprint arXiv:1704.01155* (2017).
- [50] XU, W., QI, Y., AND EVANS, D. Automatically evading classifiers. In *Proceedings of the 2016 Network and Distributed Systems Symposium* (2016).
- [51] YANG, C., WU, Q., LI, H., AND CHEN, Y. Generative poisoning attack method against neural networks. *arXiv preprint arXiv:1703.01340* (2017).
- [52] YOSINSKI, J., CLUNE, J., BENGIO, Y., AND LIPSON, H. How transferable are features in deep neural networks? In *Advances in neural information processing systems* (2014), pp. 3320–3328.
- [53] ZHANG, C., BENGIO, S., HARDT, M., RECHT, B., AND VINYALS, O. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530* (2016).

Appendix

A The StingRay Attack

Algorithm 1 shows the pseudocode of StingRay’s two general-purpose procedures. STINGRAY builds a set I with at least N_{min} and at most N_{max} attack instances. In the sample crafting loop, this procedure invokes GETBASEINSTANCE to select appropriate base instances for the target. Each iteration of the loop crafts one poison instance by invoking CRAFTINSTANCE, which modifies the set of allowable features (according to FAIL’s L dimension) of the base instance. This procedure is specific to each application. The other application-specific elements are the distance function D and the method for injecting the poison in the training set: the crafted instances may either replace or complement the base instances, depending on the application domain. Next, we describe the steps that overcome the main challenges of targeted poisoning.

Application-specific instance modification. CRAFTINSTANCE crafts a poisoning instance by modifying the set of allowable features of the base instance. The procedure selects a random sample among these features, under the constraint of the target resemblance budget. It then alters these features to resemble those of the target. Each crafted sample introduces only a small perturbation that may not be sufficient to induce the target misclassification; however, because different samples modify different features, they collectively teach the classifier that the features of \mathbf{t} correspond to label y_d . We discuss the implementation details of this procedure for the four applications in Section 4.2.

Crafting individually inconspicuous samples. To ensure that the attack instances do not stand out from the

rest of the training set, GETBASEINSTANCE randomly selects a base instance from S^l , labeled with the desired target class y_d , that lies within τ_D distance from the target. By choosing base instances that are as close to the target as possible, the adversary reduces the risk that the crafted samples will become outliers in the training set. The adversary can further reduce this risk by trading target resemblance (modifying fewer features in the crafted samples) for the need to craft more poison samples (increasing N_{min}). The adversary then checks the negative impact of the crafted instance on the training set sample S^l . The crafted instance \mathbf{x}_c is discarded if it changes the prediction on \mathbf{t} above the attacker set threshold τ_{NI} or added to the attack set otherwise. To validate that these techniques result in individually inconspicuous samples, we consider whether our crafted samples would be detected by three anti-poisoning defenses, discussed in detail in Section 4.1.

Crafting collectively inconspicuous samples. After the crafting stage, GETPDR checks the perceived PDR on the available classifier. The attack is considered successful if both adversarial goals are achieved: changing the prediction of the available classifier and not decreasing the PDR below a desired threshold τ_{PDR} .

Guessing the labels of the crafted samples. By modifying only a few features in crafted sample, CRAFTINSTANCE aims to preserve the label y_d of the base instance. While the adversary is unable to dictate how the poison samples will be labeled, they might guess this label by consulting an oracle. We discuss the effectiveness of this technique in Section 4.3.