

# ‘Security Theater’: On the Vulnerability of Classifiers to Exploratory Attacks

Tegjyot Singh Sethi<sup>1</sup>(✉), Mehmed Kantardzic<sup>1</sup>, and Joung Woo Ryu<sup>2</sup>

<sup>1</sup> Data Mining Lab, University of Louisville, Louisville, USA  
{tegjyotsingh.sethi,mehmedkantardzic}@louisville.edu

<sup>2</sup> Onycom Inc., Seoul, Republic of Korea  
ryu0914@gmail.com

**Abstract.** The increasing scale and sophistication of cyber-attacks has led to the adoption of machine learning based classification techniques, at the core of cybersecurity systems. These techniques promise scale and accuracy, which traditional rule/signature based methods cannot. However, classifiers operating in adversarial domains are vulnerable to evasion attacks by an adversary, who is capable of learning the behavior of the system by employing intelligently crafted probes. Classification accuracy in such domains provides a false sense of security, as detection can easily be evaded by carefully perturbing the input samples. In this paper, a generic data driven framework is presented, to analyze the vulnerability of classification systems to black box probing based attacks. The framework uses an *exploration-exploitation* based strategy, to understand an adversary’s point of view of the attack-defense cycle. The adversary assumes a black box model of the defender’s classifier and can launch indiscriminate attacks on it, without information of the defender’s model type, training data or the domain of application. Experimental evaluation on 10 real world datasets demonstrates that even models having high perceived accuracy (>90%), by a defender, can be effectively circumvented with a high evasion rate (>95%, on average). The detailed attack algorithms, adversarial model and empirical evaluation, serve as a background for developing secure machine learning based systems.


**Keywords:** Adversary · Reverse engineering · Classification · Cybersecurity

## 1 Introduction

The Big Data revolution has fueled the development of scalable and practical machine learning systems, which has in turn led to their widespread adaptation and popularity. The domain of cybersecurity has also recognized the need for a data driven solution [1, 2, 4, 6, 27], owing to the increased scale and sophistication of attacks in recent times<sup>1</sup>. Although the use of machine learning techniques has

---

<sup>1</sup> <https://www.statista.com/chart/2540/data-breaches/>.



**Fig. 1.** Classifier systems in adversarial environments. An adversary making probes to the model  $C(x)$  can use active learning to arrive at their own understanding of the model, as  $C'(x)$ . Probes are made through the same channels as by other users accessing the ML service, and as such can be indistinguishable.

found early success in many cybersecurity applications (such as spam filtering [11], CAPTCHA systems [10], and intrusion detection [6]), its own vulnerabilities have mostly been overlooked. Machine learning systems were designed under the assumption of stationarity, i.e. the training and the testing dataset should be identically and independently distributed [9]. This assumption is often violated in cybersecurity domains, as the systems operate in a dynamic and adversarial environment [1, 4].

In an adversarial environment, the accuracy of classification has little significance if the deployed classifier can be easily evaded by an intelligent adversary [11]. Classifiers operating in such environments are susceptible to exploratory attacks by an adversary [1], who uses the same channel as the input data to probe the system in order to gain information about it. As seen in Fig. 1, a model trained and deployed by a data miner (the *defender*) can provide services to end users, but it is also vulnerable to attacks, which use carefully crafted input samples to evade the classification. In doing so, the classifier system can be viewed as a black box ( $C$ ), providing tacit *Accept/Reject* feedback [18]. This feedback can be harnessed by an adversary, who is also equipped with the knowledge of machine learning, to reverse engineer ( $C'$ ) the behavior of the black box and avoid detection on future samples. The symmetry of the attack-defense cycle and the new gamut of vulnerabilities introduced by using classification at the core of cybersecurity systems, warrants a data driven analysis of the problem and its effects.

The security of machine learning has garnered recent interest in literature [1, 3, 17, 18, 22, 23]. A taxonomy of attacks against machine learning systems was proposed in [4], with Causative and Exploratory attacks being the broad classification of attacks, based on the portion of the data mining process they affect. Causative attacks affect the training data and are aimed at misleading the learned model. These attacks can poison the trained model, but can be prevented by careful curation of the training data [12] and by using data encryption techniques to safeguard the original training data. Exploratory attacks are more commonplace and dangerous, as they affect the testing phase data, which is often unlabeled and difficult to detect [2]. Using the same channels as a client

user these attacks can masquerade as regular data samples, posing a risk to any deployed machine learning model. Mimicry attacks [22], spoofing [2] and reverse engineering [3] are all forms of exploratory attacks. Exploratory attacks cause the training and testing data distributions to drift, leading to non-stationarity and subsequent degradation in the predictive power of a classifier [11]. These attacks can affect any deployed system, even if the system is available only as a black box service. Recent advancements in black box *Machine-Learning-as-a-Service* providers (such as Amazon AWS<sup>2</sup>, Google Cloud Platform<sup>3</sup> and BigML<sup>4</sup>), promise a new era of flexibility and ubiquity in the usage of machine learning. However, preliminary analysis in [23] has shown that these services are vulnerable to exploratory attacks, by means of querying the system through their APIs. Membership inference attacks presented in [21], shows that these services are also vulnerable to data leakage, i.e. inferring whether a particular sample belongs to a model’s training data, leading to privacy concerns in using machine learning models. The scope of these attacks can be made independent of the model trained and deployed as the black box. Deep neural networks were shown to be vulnerable in [17], who later extended their work in [18] to show how various classifiers, treated as transferable black boxes, are all equally vulnerable to such attacks.

This paper analyzes the vulnerability of classification systems to exploratory attacks, from a data driven perspective. The effects of an adversary, capable of accessing the system only as a black box; without any information about the learning process of the defender’s classifier, is presented. To the best of our knowledge, this is the first work which aims at understanding attacks as a exploration-exploitation problem and presents data generation attack algorithms, using the classifier as a black box oracle. The following are the contributions of the proposed work:

- The vulnerabilities of classifiers operating in adversarial environments, to probing based attacks, is demonstrated. These attacks show that classification systems should not be naively used in cybersecurity applications, as they can be easily evaded.
- A domain independent and data-driven framework is presented, which can be used to simulate attacks on classifiers. Under this general framework, two specific attack algorithms are provided: the Anchor Points (AP) attacks and the Reverse Engineering (RE) attacks.
- Experimental analysis on 10 real world datasets demonstrate that only information about the feature space is sufficient to launch an attack against classifiers, while being agnostic of the type of classifier, the training dataset and the domain of application. This analysis serves as a background for developing secure machine learning frameworks.

The rest of the paper is organized as follows. Section 2 presents related work in the area of exploratory attacks on classifiers. Data driven attacks on binary

<sup>2</sup> <https://aws.amazon.com/machine-learning/>.

<sup>3</sup> <https://cloud.google.com/prediction/>.

<sup>4</sup> <https://bigml.com/>.

classifiers are presented in Sect. 3. Two specific strategies and algorithms are presented in Sects. 3.1 and 3.2, for generating simple probing attacks to complex reverse engineering attacks. Section 4 presents experimental evaluation on 10 real world datasets, 7 from classification domains and 3 from cybersecurity domains. Avenues for further development are presented in Sect. 5.

## 2 Related Work on Exploratory Attacks on Machine Learning Based Classifiers

Once a model is trained and deployed in a cybersecurity application, it is vulnerable to exploratory attacks. These attacks are non-intrusive and aim at gaining information about the system, which is then exploited to create evasive samples, to avoid detection. These attacks are universal and are difficult to eliminate by traditional encryption/security techniques, because they use the same access channels as regular client users and see the same black box view of the system. Work in [16], shows that linear and convex inducing classifier are all vulnerable to probing based exploratory attacks. Exploratory attacks are classified as either: *Targeted* or *Indiscriminate*, based on the specificity of the attacks [4]. Targeted attacks aim at modifying a specific set of malicious input samples, minimally, to disguise them as legitimate. Indiscriminate attacks are more general in their goals, as they aim to produce any sample which will avoid detection by the defender’s model. Most work on exploratory attacks are concentrated on the targeted case, considering it as a constrained form of indiscriminate attacks, with the goal of starting with a malicious sample and making minimal modifications to them to avoid detection [5, 14].

Particular strategies developed for performing exploratory attacks vary based on the amount of information available to the adversary, with a broad classification presented in [3] as: (a) Evasion attacks and (b) Reverse Engineering attacks. Evasion attacks are used when limited information about the system is available, such as a few legitimate samples only. An example of evasion is seen in case of the ‘Good Words’ attacks on spam classifiers, where the word *SALE* is modified to *SA1E*, to avoid being flagged [15]. In [26], a general purpose domain independent evasion technique was developed. Genetic programming was used to generate variants of a set of malicious samples, as per a monotonically increasing fitness function describing success of evasion. This technique is attractive due to its generality, but its practicality is limited by the lack of a graded fitness function and limited probing budgets. The gradient descent attacks of [5] provides an efficient heuristic approach to utilizing the information about the classifier model, to generate optimal modification for targeted evasion of a set of data samples. However, the attack strategy relies on knowing the exact model of the defender and cannot be effectively used when only a black box interface to the defender’s classifier is presented.

Reverse engineering attacks on classifiers, provides avenues for large scale evasion, as it conveys important internal information about the importance of features to the classification task. Reverse engineering was employed in [14],

where a signed witness test was used to identify if a particular feature has a positive or a negative impact on the prediction decision. Reverse engineering of a decision tree model was presented in [25]. In [26], genetic programming was used as a model independent reverse engineering tool, assuming that the training data is known. The idea of reverse engineering was linked to that of active learning in [3]. Here, the robustness of Support Vector Machines (SVM) classifiers, to reverse engineering, was tested using active learning techniques of random sampling, uncertainty sampling and selective sampling.

In the above described targeted-exploratory attacks, it is assumed that an adversary would give up if an attack is expensive (far from the original samples). These attacks do not consider the case of a determined adversary intending to launch an indiscriminate attack. These type of attacks have been largely ignored, with the only mention we found was in [28], where it is termed - the free range attack, as an adversary is free to move about in the data space. Analyzing performance of models under such attack scenarios is essential to understanding its vulnerabilities in a more general and real world situation, where all types of attacks are possible. Also, while most recent methodologies develop attacks as an experimental tool to test their safety mechanisms, there is very few works [18, 21, 23], which have attempted to study the attack generation process itself. Our proposed work analyzes the vulnerability of classifiers to *Indiscriminate-Exploratory* attacks, where only a black box model of the defender is available. A data driven framework is proposed, thereby highlighting the symmetry of the problem of attack and defense, to motivate a data driven solution.

### 3 Data Driven Attacks on Classifiers

Data driven attacks on classification systems are exploratory in nature. An adversary proceeds by making probes to the classifier, by means of generated input samples, which it presents to the system. The feedback, a simple accept/reject in most cases, can then be used to infer the nature of the trained model. The classifier is seen only as a black box, which can provide binary feedback on input samples, in the same way as it provides classification on regular benign input data (Fig. 1). As an example, a spam classification system will not provide any information other than its intended behavior of marking input emails as spam, based on its analysis. In this setting, the model of an adversary can be formalized as follows:

- **Knowledge:** Adversary is aware of the number, range and type of features used by the system. This can be approximated from publicly available case studies or by educated guessing. No information about classifier type and training data is known.
- **Goals:** Adversary intends to produce false negatives for the classification.
- **Resources:** The attacker has access to the system as a client user. It can submit probe samples and receive feedback, up to a budget  $B$ , without being detected.

Based on the above model of an adversary, the problem of generating exploratory attacks can be formalized. A binary classifier  $C$ , trained on a set of training data  $D$ , is deployed to classify input samples as *Legitimate* or *Malicious*. An adversary aims to generate samples  $D'_{Attack}$ , such that  $C(D'_{Attack})$  has a high false negative rate. The adversary has at its disposal a budget  $B_{Explore}$  of probing data  $D'_{Explore}$ , which it can use to learn about  $C$  and understand it as  $C'$ . This setting represents a natural scenario where attackers start with limited reconnaissance and then launch a dedicated campaign, based on the learned vulnerabilities. An adversary operating in this scenario can utilize an **Explore-Exploit** strategy, popular in search based optimization techniques [24], to best utilize the budget  $B_{Explore}$  to produce  $D'_{Explore}$ . Two specific instantiations of this general idea are presented here as the Anchor Points (AP) attack and the Reverse Engineering (RE) attack.

### 3.1 The Anchor Points(AP) Attack

Anchor Points attacks start with an adversary gaining information about a set of samples classified as *Legitimate* by the classifier  $C$ , which it then uses as *Anchors* to launch new attack instances. These attacks are characteristic of an adversary who has a limited probing budget  $B_{Explore}$  and who wishes to quickly exploit a new found vulnerability, as is common in the case of zero day exploits [7].

From a data driven perspective, the attacks begin by obtaining a set of seed *Legitimate* samples  $D'_{Seed}$ . As an example, the adversary could start with a set of legitimate emails obtained from its' own inbox, for a spam evasion task. The obtained seed samples are then used to trigger the exploration phase as described in Algorithm 1. The exploration phase intends to obtain a set of *Anchor Points*, which will serve as ground truth for representing the space of samples classified as *Legitimate*. This exploration is performed using a radius based incremental neighborhood search around the seed samples, guided by the feedback from the black box classifier  $C$ . Diversity of search is ensured by dynamically updating the neighborhood radius  $R_i$  in every iteration, as given by (Line 5). This equation causes the radius of exploration to increase in cases where the number of legitimate samples obtained are high, thereby balancing diversity of search with its accuracy. Samples are explored by searching for a random sample within the radius of exploration, as given by (Line 6). The final exploration dataset of *Anchor Points* -  $D'_{Explore}$ , is comprised of all the explored samples  $x_i$  such that  $C(x_i)$  is *Legitimate*. This is depicted in Fig. 2, as the set of positive points which are obtained at the end of the exploration cycle.

The explored anchor points set  $D'_{Explore}$ , serves as the basis to launch a dedicated attack campaign. Algorithm 2 incorporates information learned in the exploration phase, to generate the attack samples while also imparting diversity to the attack set  $D'_{Attack}$ . Diversity is imparted by adding random perturbation to the samples (Line 4) and then generating a sample based on their convex combination (Line 6), as inspired by the Synthetic Minority Oversampling Technique(SMOTE) for imbalanced datasets [8]. Random perturbation is controlled by the input parameter  $R_{Exploit}$ , which is kept close to  $R_{min}$  (Algorithm 1), as

**Algorithm 1.** AP- Exploration Phase

---

**Input** : Seed Data  $D'_{Seed}$ , Defender black box  $C$ . *Parameters*: Exploration budget  $B_{Explore}$ , Exploration neighborhood-  $[R_{min}, R_{max}]$

**Output**: Exploration data set  $D'_{Explore}$

- 1  $D'_{Explore} \leftarrow D'_{Seed}$
- 2 count\_legitimate=0
- 3 **for**  $i = 1 .. B_{Explore}$  **do**
- 4      $x_i \leftarrow$  Select random sample from  $D'_{Explore}$
- 5      $R_i = (R_{max} - R_{min}) * (count\_legitimate/i) + R_{min}$
- 6      $\hat{x}_i \leftarrow Perturb(x_i, R_i)$  ▷ Perturbed sample
- 7     **if**  $C.predict(\hat{x}_i)$  is Legitimate **then**
- 8          $D'_{Explore} \cup \hat{x}_i$
- 9         count\_legitimate ++

10 **Procedure** Perturb(sample,  $R_{Neigh}$ )

11 return sample+=random(mean=0, std= $R_{Neigh}$ )

---

**Algorithm 2.** AP- Exploitation Phase

---

**Input** : Exploration data set  $D'_{Explore}$ , Number of attacks  $N_{Attack}$ , Radius of Exploitation  $R_{Exploit}$

**Output**: Attacks set  $D'_{Attack}$

- 1  $D'_{Attack} \leftarrow []$
- 2 **for**  $i = 1 .. N_{Attack}$  **do**
- 3      $x_A, x_B \leftarrow$  Select random samples from  $D'_{Explore}$
- 4      $\hat{x}_A, \hat{x}_B \leftarrow Perturb(x_A, R_{Exploit}), Perturb(x_B, R_{Exploit})$
- 5      $\lambda = random(0, 1)$  ▷ Random number in  $[0,1]$
- 6      $attack\_sample_i \leftarrow \hat{x}_A * \lambda + (1 - \lambda) * \hat{x}_B$
- 7      $D'_{Attack} \cup attack\_sample_i$


8 **Procedure** Perturb(sample,  $R_{Exploit}$ )

9 return sample+=random(mean=0, std= $R_{Exploit}$ )

---

no explicit feedback from the black box  $C$ , is available in this phase. The final set of attack samples  $D'_{Attack}$  is submitted as attack on the classifier  $C$ , shown as red samples in Fig. 2. An adversary aims to cause a high false negative rate for  $C$ , while at the same time have high diversity in its attacks.  $D'_{Attack}$  is kept much larger than  $B_{Explore}$  to justify adversarial budget expenditure.

The performance of the Anchor Points attack depends on the diversity and accuracy of samples collected in the exploration phase. Larger coverage ensures flexibility in the attack phase. By the nature of these attacks, they could be thwarted by blacklists capable of approximate matching [20]. Nevertheless, these techniques are suited for adhoc swift exploits, to cause impact before a defender has time to respond.



**Fig. 2.** Illustration of AP attacks on 2D synthetic data. (*Left - Right*): The defender’s model from training data. The Exploration phase depicting the seed (blue) and the anchor points samples. The Exploitation phase samples generated based on the anchor points, and submitted as attack payload. (Color figure online)

### 3.2 The Reverse Engineering(RE) Attack

These attacks aim to directly reverse engineer the classification boundary, so as to better understand the classification landscape, which can then be leveraged to launch large scale evasion attacks on the black box  $C$ . Reverse engineering could be a goal in itself, as it provides information about features importance to the classification task, or it could be a first step to launching an evasion or availability attack [3]. A reverse engineering attack, if done effectively, can avoid detection and can make retraining more difficult on the part of the defender. However, unlike the AP attack, the reverse engineering process is affected by the type of classifier model used by  $C$ , and is also dependent on the availability of sufficient exploration budget  $B_{Explore}$ , for the reverse engineering learning task. Nevertheless, an adversary motivated to evade the classification system is not concerned with fitting the decision boundary of  $C$  exactly. A linear approximation to the non linear defender’s boundary is sufficient to launch a reduced accuracy attack, which can be compensated for by launching a massive attack campaign, utilizing the information provided by the reverse engineered model.

Effective reverse engineering depends on making best use of the  $B_{Explore}$ . Random sampling can lead to wasted probes, with no new information added. The query synthesis strategy of [24] generates samples close to the decision boundary and spreads these samples across the boundary, for better learning of the decision landscape. The approach in [24] was used for selecting samples for active learning. We modify the approach for the task of reverse engineering in Algorithm 3, where a surrogate classifier  $C'$  is learned as a result of the exploration phase. The algorithm begins by accepting a seed datasets, which is comprised of atleast one *Legitimate* and *Malicious* sample. The algorithm then employs the Gram-Schmidt process [24], to generate orthonormal samples near the midpoint of two randomly selected points of the opposite classes, as shown in Fig. 3. The magnitude of the orthonormal mid-perpendicular vector is set to  $\lambda_i$ , selected as random value in  $[0, \lambda]$ , to incorporate variability in the exploration



**Algorithm 3.** RE Exploration - Using Gram-Schmidt process.

---

**Input** : Seed Data  $D'_{Seed}$ , Defender black box model  $C$ . *Parameters*:  
Exploration budget  $B_{Explore}$ , Magnitude of dispersion  $\lambda$

**Output**: Exploration data Set  $D'_{Explore}$ , Surrogate classifier  $C'$

- 1  $D'_{Explore\_L}, D'_{Explore\_M} =$  Legitimate, Malicious samples of  $D'_{Seed}$
- 2 **for**  $i = 1 .. B_{Explore}$  **do**
- 3      $x_L, x_M \leftarrow$  Select random samples from  $D'_{Explore\_L}, D'_{Explore\_M}$
- 4      $x_0 = x_L - x_M$
- 5     Generate random vector  $x_R$
- 6      $x_R = x_R - \frac{\langle x_R, x_0 \rangle}{\langle x_0, x_0 \rangle} * x_0$       $\triangleright$  Gram-Schmidt process -  $x_R$  orthogonal to  $x_0$
- 7      $\lambda_i = \text{random}(0, \lambda)$
- 8      $x_R = \frac{\lambda_i}{\text{norm}(x_R)} * x_R$       $\triangleright$  Set magnitude of orthogonal midperpendicular
- 9      $x_S = (x_R + (x_L + x_M))/2$       $\triangleright$  Set  $x_R$  to midpoint
- 10    **if**  $C.\text{predict}(x_S)$  is legitimate **then**
- 11      $\lfloor D'_{Explore\_L} \cup x_S$
- 12    **else**
- 13      $\lfloor D'_{Explore\_M} \cup x_S$
- 14  $D'_{Explore} = D'_{Explore\_L} \cup D'_{Explore\_M}$
- 15 Train  $C'$  using  $D'_{Explore}$       $\triangleright$  Training can be based on linear classifier of choice

---


phase (Line 8). The resulting exploration samples are then classified as Legitimate/Malicious by  $C$ , which can be probed upto  $B_{Explore}$ . The final combined dataset  $D'_{Explore}$  (Line 14) is then used to train a linear classifier of choice, to form the surrogate reverse engineered model  $C'$  (Line 15).

The reverse engineered model  $C'$  can be used to crosscheck the randomly generated samples in the exploitation phase, to ensure that attacks have high accuracy. A practical and effective exploitation strategy is to use the  $D'_{Explore}$  as the seed set for Algorithm 1, with the exception that we use  $C'$  to probe instead of the original  $C$ . Since the  $C'$  is a locally trained model, probing it does not impact  $B_{Explore}$ . Thereby allowing an adversary to make, theoretically, infinite queries to  $C'$ , at effectively zero cost. The anchor points obtained can then be used to perform exploitation using Algorithm 2. The exploitation can use a large  $R_{Exploit}$ , as the results can be verified against the surrogate  $C'$ , to ensure higher diversity and higher accuracy of attacks, than the AP attacks.

## 4 Empirical Evaluation

### 4.1 Experimental Setup

Experimental evaluation, presented here, shows an adversary's point of view of the classification system. The adversary is capable of generating data driven attacks on the system, by making limited probes to it and then generating a dedicated campaign of evasive samples. Since the adversary aims at evading the classification system, its efficacy is measured as the Effective Attack Rate(EAR),



**Fig. 3.** Illustration of RE attacks on 2D synthetic data. (Left - Right): The defender’s model based on training data. The Exploration phase depicting reverse engineering using the Gram-Schmidt orthonormalization process. The Exploitation attack phase samples generated after validation from the surrogate classifier.

given by Eq. 1. This equation measures the accuracy of the attacks, from an adversary’s point of view, and it indicates the false negative rate of the defender’s classifier on the attack samples.

$$EAR = \frac{|\{x : C(x) = Legitimate \wedge x \in D'_{Attack}\}|}{|D'_{Attack}|} \tag{1}$$

Here,  $C$  represents the defender’s black box classifier and  $D'_{Attack}$  is the set of attack campaign samples generated by the adversary to attack  $C$ . The goal of an adversary is to maximize EAR.

The evaluation is performed using 10 binary classification datasets, shown in Table 1 (Column 1). The first 7 datasets represent standard classification tasks and were obtained from the UCI machine learning repository [13]. The Spam-base dataset for email classification [13], the KDD99 intrusion detection dataset [13] and the CAPTCHA dataset for classifying human-bots based on behavioral data [10], represent 3 different cybersecurity applications which employ machine learning based classification at its core. All datasets were transformed to have numerical values normalized in the range of [0,1]. Instances were shuffled to remove any bias due to inherent concept drift. The class label 1 is taken as the *Malicious* class and 0 is taken as the *Legitimate* class, as convention.

In all experiments, the exploration probe budget  $B_{Explore}$  is taken as 1000 and the number of attack samples to be generated  $N_{Attack}$  is taken as 2000. For the Anchor Points (AP) attack, the neighborhood radius  $[R_{min}, R_{max}]$  is taken as [0.1,0.5] and the exploitation radius as  $R_{Exploit} = 0.1$ . In case of the reverse engineering (RE) attacks, a larger exploitation radius ( $R_{Exploit}=0.5$ ) is considered, due to additional validation provided by the surrogate learned classifier  $C'$ . A SVM with linear kernel and high regularization constant ( $c=10$ ) is taken for the surrogate classifier, to prevent overfitting to  $D'_{Explore}$ . The magnitude of dispersion ( $\lambda$ ) is taken as 0.25, and it was found that changing this had little affect on the final results. In all experiments, no information about the black box  $C$  is known by the adversary. All experiments in this section are performed using

**Table 1.** Results of AP and RE attacks on a linear defender model on 10 real world datasets (EAR - Effective Attack Rate).


Dataset (#Instances, #Attributes)	Defender’s initial accuracy	Explored anchor points/ $B_{Explore}$	Accuracy of RE model $C'$	EAR	
				AP	RE
Digits08 (1500, 16)	0.98	0.63	0.92	$0.96 \pm 0.01$	$0.93 \pm 0.06$
Credit (1000, 61)	0.79	0.71	0.71	$0.98 \pm 0.01$	$0.8 \pm 0.15$
Cancer (699, 10)	0.97	0.99	0.95	$0.99 \pm 0.01$	$0.99 \pm 0.01$
Qsar (1055, 41)	0.87	0.99	0.42	$0.99 \pm 0.01$	$0.99 \pm 0.01$
Sonar (208, 60)	0.88	0.98	0.61	$0.99 \pm 0.01$	$0.98 \pm 0.01$
Theorem (3060, 51)	0.72	0.67	0.57	$0.97 \pm 0.01$	$0.87 \pm 0.08$
Diabetes (768, 8)	0.78	0.5	0.71	$0.98 \pm 0.01$	$0.95 \pm 0.04$
Spambase (4600, 57)	0.91	0.5	0.59	$0.93 \pm 0.01$	$0.71 \pm 0.2$
KDD99 (494021, 41)	0.99	0.91	0.55	$0.99 \pm 0.01$	$0.93 \pm 0.04$
CAPTCHA (1885, 26)	1.0	0.92	0.91	$0.99 \pm 0.01$	$0.97 \pm 0.02$

Python 2.7 and the scikit-learn library [19]. Results are averaged over 30 runs for every experiment. Section 4.2 presents the results of a symmetric case, where both the adversary and the defender have similar model types, while Sect. 4.3 presents analysis on non symmetric model types, with 4 different classifiers for the black box.

## 4.2 Experiments with Symmetric Defender Model

Experiments with a linear kernel SVM for the defender’s model (regularization parameter,  $c = 1$ ), is presented here, to show effects of a symmetric model type between the adversary and the defender. The initial accuracy of the defender, as perceived by cross-validation on its training dataset before deployment, is shown in Table 1. The Effective Attack Rate (EAR) shows that even models which are perceived to have a high accuracy ( $>70\%$  in all 10 cases) by the defender, are effectively evaded by an adversary, with an EAR of 97.7% in case of the AP attacks and 91.2% for the RE attacks, on average. This shows the inherent vulnerability of the classification models and the misleading nature of accuracy, in an adversarial environment.

Both the attack methods of AP and RE are effective. The fundamental difference between the two approaches is that RE places its confidence in its understanding of the separating boundary, while the AP approach places its confidence only on the anchor points obtained during exploration. From Table 1, it is seen that the number of anchor points obtained is  $>50\%$  of  $B_{Explore}$  (Column 3), making it a simple attack strategy in high dimensional spaces. For the RE attack, the accuracy of the surrogate classifier  $C'$  (Column 4) is computed by evaluating it on the original dataset, as an adhoc metric of  $C'$ ’s understanding of the original data space and the extent of reverse engineering. It is observed that even in



**Fig. 4.** Improving diversity of attacks in AP leads to decreased EAR (a). In case of RE, which has high diversity, EAR can be improved by increasing  $B_{Explore}$  (b).

cases where the RE accuracy is low (0.41 for Qsar), a high EAR is seen (0.99). This is because, the goal of the adversary is not to totally reverse engineer the model  $C$ , but instead to learn it sufficiently enough to evade it. This enables the linear approximation technique of the RE approach to work as an effective attack strategy. The higher variability in EAR for the RE attacks is also a result of this approximation and its dependence on the quality of the exploration data.

While the AP attacks have a higher EAR than the RE attacks (Table 1), the RE approach provides better diversity of attacks, as it uses a larger  $R_{Exploit} = 0.5$  (for AP,  $R_{Exploit} = 0.1$ ). Diversity ensures that simple countermeasures of blacklisting will not thwart an attack [20]. The effect of increasing the  $R_{Exploit}$  for the AP attack, in an attempt to increase its diversity, is shown in Fig. 4(a). It is seen that this leads to rapid deterioration in the EAR, as the only ground truth information available is the Anchor Points obtained during exploration. Increasing distance from these points leads to uncertainty and reduced accuracy of attacks. For the RE attacks, the datasets of - Cancer, Theorem and Spambase, are seen to have a low EAR in Table 1. Effects of increasing the  $B_{Explore}$ , in an attempt to increase the EAR is shown in Fig. 4(b). Increasing number of probes leads to an increase in the understanding of the black box  $C$ , which translates to better EAR. This increase plateaus after a critical mass of samples, needed for active learning the model  $C$ , is reached. This indicates the efficacy of the RE attacks as a long term attack strategy on a classification system, where over time the additional information learned can lead to both high accuracy and high diversity.

### 4.3 Experiments with Non-symmetric Defender Model

The development of the AP and the RE attack strategies considers a black box defender’s model. These techniques are essentially data space search approaches, which are independent of the defender’s underlying model type and their parameters. Results of testing the behavior of these attack strategies on a non-symmetric and non-linear defender models is presented in Table 2. The following defender black box models are considered: K-Nearest Neighbor classifier (KNN) with  $K = 3$ , SVM with a radial basis function kernel (SVM-RBF,  $\gamma = 0.1$ ), C4.5 decision

**Table 2.** Effective Attack Rate (EAR) of AP and RE attacks, with non linear defender’s model (Low EAR values are italicized.)

Dataset	KNN		SVM-RBF		DT		RF	
	AP	RE	AP	RE	AP	RE	AP	RE
Digits08	0.89	0.96	0.97	0.89	0.87	0.63	0.85	<i>0.48</i>
Credit	0.96	0.78	0.94	0.53	0.79	<i>0.42</i>	0.79	<i>0.33</i>
Cancer	0.99	0.99	0.99	0.99	0.97	0.89	0.99	0.98
Qsar	1	0.99	0.99	0.99	0.96	0.76	0.99	0.99
Sonar	0.99	0.98	1	1	0.97	0.62	0.99	0.95
Theorem	0.97	0.813	0.95	0.5	0.95	0.79	0.62	0.78
Diabetes	0.99	0.935	0.99	0.9	0.83	0.63	0.88	0.61
Spambase	0.93	0.99	<i>0.48</i>	0.84	<i>0.08</i>	<i>0.11</i>	0.99	0.98
KDD99	0.99	0.93	1	0.99	0.89	0.54	0.92	<i>0.27</i>
CAPTCHA	0.99	0.92	0.99	0.92	0.97	0.83	0.93	0.89

trees (DT), and random forest of 50 decision models (RF) [19]. The parameters of the experiments are kept the same as in Sect. 4.2.

It is seen that the AP approach is minimally affected by the choice of the defender’s model. The RE attacks are affected by the choice of the model, as seen in the case of the Credit and the Theorem datasets. These datasets were seen to have a low accuracy, when trained using a linear model (Table 1), indicating non linearity in their model space. In these cases, the linear approximation in the RE approach, is not sufficient to have a high EAR. However, in a majority of the cases it is seen that a >50% attack rate is still possible with just a linear SVM model used for reverse engineering. A high average attack rate, irrespective of the underlying classifier used, indicates vulnerability of classification to purely data driven attacks. An interesting observation warranting further investigation is the comparatively low EAR on the decision tree models, which could be indicative of their attack resistance and its inverse relation to model robustness.

## 5 Conclusion and Future Work

We present a general data driven framework for demonstrating the vulnerability of classification systems to exploratory attacks at test time. Under this framework, two specific attack algorithms were developed: The Anchor Points attack (AP) and the Reverse Engineering attacks (RE). The effectiveness of these attack algorithms on 10 real world datasets, demonstrates that adversarial attacks can be launched having only the knowledge of the feature space of the data, agnostic of the defender’s classifier type, training data and the domain of application. The defender’s perceived accuracy was shown to be of little importance, if the model can be easily evaded by such probing based attacks. This is especially relevant in cybersecurity application domains, where the primary purpose of the

classifier is to provide security. In these domains, it is worth emphasizing that - higher accuracy in machine learning does not necessarily imply better security.

While the proposed work presents the adversary's point of view of the attack-defense cycle, its goal is to move towards a more secure paradigm of using classifiers in cybersecurity domains, by clearly understanding its vulnerabilities. Future work will include attack detection and effective relearning, in environments with adversarial activity.

## References

1. Abramson, M.: Toward adversarial online learning and the science of deceptive machines. In: 2015 AAAI Fall Symposium Series (2015)
2. Akhtar, Z., et al.: Robustness of multi-modal biometric systems under realistic spoof attacks against all traits. In: BIOMS 2011, pp. 1–6. IEEE (2011)
3. Alabdulmohsin, I.M., et al.: Adding robustness to support vector machines against adversarial reverse engineering. In: Proceedings of the 23rd ACM International Conference on Information and Knowledge Management, pp. 231–240. ACM (2014)
4. Barreno, M., Nelson, B., Sears, R., Joseph, A.D., Tygar, J.D.: Can machine learning be secure? In: Proceedings of the 2006 ACM Symposium on Information, Computer and Communications Security, pp. 16–25. ACM (2006)
5. Biggio, B., Corona, I., Maiorca, D., Nelson, B., Šrndić, N., Laskov, P., Giacinto, G., Roli, F.: Evasion attacks against machine learning at test time. In: Blockeel, H., Kersting, K., Nijssen, S., Železný, F. (eds.) ECML PKDD 2013. LNCS (LNAI), vol. 8190, pp. 387–402. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-40994-3\\_25](https://doi.org/10.1007/978-3-642-40994-3_25)
6. Biggio, B., Fumera, G., Roli, F.: Security evaluation of pattern classifiers under attack. *IEEE Trans. Knowl. Data Eng.* **26**(4), 984–996 (2014)
7. Bilge, L., Dumitras, T.: Before we knew it: an empirical study of zero-day attacks in the real world. In: Proceedings of the 2012 ACM Conference on Computer and Communications Security, pp. 833–844. ACM (2012)
8. Chawla, N.V., et al.: SMOTE: synthetic minority over-sampling technique. *J. Art. Intell. Res.* **16**, 321–357 (2002)
9. Ditzler, G., Roveri, M., Alippi, C., Polikar, R.: Learning in nonstationary environments: a survey. *IEEE Comput. Intell. Mag.* **10**(4), 12–25 (2015)
10. D'Souza, D.F.: Avatar captcha: telling computers and humans apart via face classification and mouse dynamics. *Electronic Theses and Dissertations-1715* (2014)
11. Kantchelian, A., et al.: Approaches to adversarial drift. In: Proceedings of the 2013 ACM Workshop on Artificial Intelligence and Security, pp. 99–110. ACM (2013)
12. Li, H., Chan, P.P.K.: An improved reject on negative impact defense. In: Wang, X., Pedrycz, W., Chan, P., He, Q. (eds.) ICMLC 2014. CCIS, vol. 481, pp. 452–459. Springer, Heidelberg (2014). doi:[10.1007/978-3-662-45652-1\\_45](https://doi.org/10.1007/978-3-662-45652-1_45)
13. Lichman, M.: UCI machine learning repository (2013). <http://archive.ics.uci.edu/ml>
14. Lowd, D., Meek, C.: Adversarial learning. In: Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining, pp. 641–647. ACM (2005)
15. Lowd, D., Meek, C.: Good word attacks on statistical spam filters. In: CEAS (2005)
16. Nelson, B., Rubinstein, B.I., Huang, L., Joseph, A.D., Lau, S., Lee, S.J., Rao, S., Tran, A., Tygar, J.D.: Near-optimal evasion of convex-inducing classifiers. In: AISTATS, pp. 549–556 (2010)

17. Papernot, N., et al.: The limitations of deep learning in adversarial settings. In: IEEE European Symposium on Security and Privacy, pp. 372–387. IEEE (2016)
18. Papernot, N., McDaniel, P., Goodfellow, I.: Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. arXiv preprint [arXiv:1605.07277](https://arxiv.org/abs/1605.07277) (2016)
19. Pedregosa, F., et al.: Scikit-learn: machine learning in python. *J. Mach. Learn. Res.* 2825–2830 (2011)
20. Prakash, P., et al.: Phishnet: predictive blacklisting to detect phishing attacks. In: Proceedings of IEEE INFOCOM, pp. 1–5. IEEE (2010)
21. Shokri, R., Stronati, M., Shmatikov, V.: Membership inference attacks against machine learning models. arXiv preprint [arXiv:1610.05820](https://arxiv.org/abs/1610.05820) (2016)
22. Smutz, C., Stavrou, A.: When a tree falls: using diversity in ensemble classifiers to identify evasion in malware detectors. In: NDSS Symposium (2016)
23. Tramèr, F., Zhang, F., Juels, A., Reiter, M.K., Ristenpart, T.: Stealing machine learning models via prediction APIs. arXiv preprint [arXiv:1609.02943](https://arxiv.org/abs/1609.02943) (2016)
24. Wang, L., Hu, X., Yuan, B., Lu, J.: Active learning via query synthesis and nearest neighbour search. *Neurocomputing* **147**, 426–434 (2015)
25. Xu, L., Zhan, Z., Xu, S., Ye, K.: An evasion and counter-evasion study in malicious websites detection. In: 2014 IEEE Conference on Communications and Network Security (CNS), pp. 265–273. IEEE (2014)
26. Xu, W., Qi, Y., Evans, D.: Automatically evading classifiers. In: Proceedings of the 2016 Network and Distributed Systems Symposium (2016)
27. Zhou, Y., Kantarcioglu, M.: Modeling adversarial learning as nested stackelberg games. In: Bailey, J., Khan, L., Washio, T., Dobbie, G., Huang, J.Z., Wang, R. (eds.) PAKDD 2016. LNCS (LNAI), vol. 9652, pp. 350–362. Springer, Cham (2016). doi:[10.1007/978-3-319-31750-2\\_28](https://doi.org/10.1007/978-3-319-31750-2_28)
28. Zhou, Y., Kantarcioglu, M., Thuraisingham, B., Xi, B.: Adversarial support vector machine learning. In: Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1059–1067. ACM (2012)