

Singapore Management University

## Institutional Knowledge at Singapore Management University

---

Research Collection School Of Computing and  
Information Systems

School of Computing and Information Systems

---

10-2018

### A learning and masking approach to secure learning

Linh NGUYEN

Sky WANG

Arunesh SINHA

Singapore Management University, aruneshs@smu.edu.sg

Follow this and additional works at: [https://ink.library.smu.edu.sg/sis\\_research](https://ink.library.smu.edu.sg/sis_research)



Part of the [Databases and Information Systems Commons](#), and the [Software Engineering Commons](#)

---

#### Citation

NGUYEN, Linh; WANG, Sky; and SINHA, Arunesh. A learning and masking approach to secure learning. (2018). *Proceedings of the 9th Conference on Decision and Game Theory for Security: GameSec 2018, Seattle, USA, October 29-31*. 11199, 453-464.

Available at: [https://ink.library.smu.edu.sg/sis\\_research/4793](https://ink.library.smu.edu.sg/sis_research/4793)

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email [cherylds@smu.edu.sg](mailto:cherylds@smu.edu.sg).

# A Learning and Masking Approach to Secure Learning

Linh Nguyen, Sky Wang, Arunesh Sinha  
University of Michigan, Ann Arbor  
{lvynguyen, skywang, arunesh}@umich.edu

## ABSTRACT

Deep Neural Networks (DNNs) have been shown to be vulnerable against adversarial examples, which are data points cleverly constructed to fool the classifier. Such attacks can be devastating in practice, especially as DNNs are being applied to ever increasing critical tasks like image recognition in autonomous driving. In this paper, we introduce a new perspective on the problem. We do so by first defining robustness of a classifier to adversarial exploitation. Next, we show that the problem of adversarial example generation can be posed as learning problem. We also categorize attacks in literature into high and low perturbation attacks; well-known attacks like FGSM [11] and our attack produce higher perturbation adversarial examples while the more potent but computationally inefficient Carlini-Wagner [5] (CW) attack is low perturbation. Next, we show that the dual approach of the attack learning problem can be used as a defensive technique that is effective against high perturbation attacks. Finally, we show that a classifier masking method achieved by adding noise to the a neural network’s logit output protects against low distortion attacks such as the CW attack. We also show that both our learning and masking defense can work simultaneously to protect against multiple attacks. We demonstrate the efficacy of our techniques by experimenting with the MNIST and CIFAR-10 datasets.

## KEYWORDS

adversarial examples; robust learning

## 1 INTRODUCTION

Recent advances in deep learning have led to its wide adoption in various challenging tasks such as image classification. However, the current state of the art has been shown to be vulnerable to *adversarial examples*, small perturbations of the original inputs, often indistinguishable to a human, but carefully crafted to misguide the learning models into producing incorrect outputs. Recent results have shown that generating these adversarial examples are inexpensive [11]. Moreover, as safety critical applications such as autonomous driving increasingly rely on these tasks, it is imperative that the learning models be reliable and secure against such adversarial examples.

Prior work has yielded a lot of attack methods that generate adversarial samples, and defense techniques that improve the accuracy on these samples (see related work for details). However, defenses are often specific to certain attacks and cannot adaptively defend against any future attack and some general defense techniques have been shown to be ineffective against more powerful novel attacks. More generally, attacks and defenses have followed the cat-and-mouse game that is typical of many security settings.

Further, traditional machine learning theory assumes a fixed stochastic environment hence accuracy in the traditional sense is not enough to measure performance in presence of an adversarial agent.

In this paper, with the goal of generality, we pursue a principled approach to attacks and defense. Starting from a theoretical robustness definition, we present a attack and a defense that learns to generate adversarial examples against any given classifier and learns to defend against any attack respectively. Based on formal intuition, we categorize known attacks into high and low perturbation attacks. Our *learning* attack is a high perturbation attack and analogously our *learning* defense technique defends against high perturbation attack. For low perturbation attacks, we provide a *masking* approach that defends against such attacks. Our two defense techniques can be combined to defend against multiple types of attacks. While our guiding principle is general, this paper focuses on the specific domain of adversarial examples in image classification.

Our *first contribution* is a definition of *robustness of classifiers* in presence of an adversarial agent. Towards the definition, we define the exploitable space by the adversary which includes data points already mis-classified (errors) by any given classifier and any data points that can be perturbed by the adversary to force mis-classifications. Robustness is defined as the probability of data points occurring in the exploitable space. We believe our definition captures the essence of the multi-agent defender-adversary interaction, and is natural as our robustness is a strictly stronger concept than accuracy. We also analyze why accuracy fails to measure robustness. The formal set-up also provides an intuition for all the techniques proposed in this paper.

Our *second contribution* is an *attack learning neural network* (ALN). ALN is motivated by the fact that adversarial examples for a given classifier  $C$  are subsets of the input space that the  $C$  mis-classifies. Thus, given a data distribution with data points  $x$  and a classifier  $C$  trained on such data, we train a feed forward neural network  $A$  with the goal of generating output points  $A(x)$  in the mis-classified space. Towards this end, we re-purpose an autoencoder to work as our ALN  $A$  with a special choice of loss function that aims to make (1) the classifier  $C$  mis-classify  $A(x)$  and (2) minimize the difference between  $x$  and  $A(x)$ .

Our *third contribution* are two defense techniques: *defense learning neural network* (DLN) and *noise augmented classifier* (NAC). Following the motivation and design of ALN, we motivate DLN  $D$  as a neural network that, given any classifier  $C$  attacked by an attack technique  $A$ , takes in an adversarial example  $A(x)$  and aims to generate benign example  $D(A(x))$  that *does not* lie in the mis-classified space of  $C$ . The DLN is prepended to the classifier  $C$  acting as a sanitizer for  $C$ . Again, similar to the ALN, we re-purpose an autoencoder with a special loss function suited for the goal of the DLN. For non-adversarial inputs the DLN is encouraged to reproduce the input as well as make the classifier predict correctly. We show that

DLN allows for attack and defense to be set up as a repeated competition leading to more robust classifiers. Next, while DLN works efficiently for attacks that produces adversarial examples with high perturbation, such as fast gradient sign method [11] (FGSM), it is not practical for low perturbation attacks (discussed in details in Section 3.4) such as Carlini-Wagner [5] (CW). For low perturbation attacks, we present NAC which masks the classifier boundary by adding a very small noise at the logits output of the neural network classifier. The small noise added affects classification in rare cases, thereby ensuring original accuracy is maintained, but also fools low perturbation attacks as the attack is misled by the *incorrect* logits. DLN and NAC can work together to defend simultaneously against both high and low perturbation attacks.

We tested our approach on two datasets: MNIST and CIFAR-10. Our ALN based attack was able to attack all classifiers we considered and achieve performance comparable to other high perturbation attacks. Our defense approach made the resultant classifier robust to the FGSM and CW. Detailed experiments are presented in Section 4 and 5. Missing proofs are in an online appendix<sup>1</sup> (see footnote).

## 2 ATTACK MODEL

Given the adversarial setting, it is imperative to define the capabilities of the adversary, which we do in this section. First, we use *inference phase* of a classifier to mean the stage when the classifier is actually deployed as an application (after all training and testing is done). The attacker attacks *only* in the inference phase and can channel his attack *only* through the inputs. In particular, the attacker cannot change the classifier weights or inject any noise in the hidden layers. The attacker has access to the classifier weights, so that it can compute gradients if required. The attacker’s goal is to produce adversarial data points that get mis-classified by the classifier. These adversarial examples should be legitimate (that is not a garbage noisy image) and the true class and the predicted class of the data point could be additional constraints for the adversary.

## 3 APPROACH

This section formally describes our approach to the adversarial example generation and defense problem using the notion of robustness we define. We start by defining basic notations. Let the function  $C : X \rightarrow Y$  denote a classifier that takes input data points with feature values in  $X$  and outputs a label among the possible  $k$  labels  $Y = \{1, \dots, k\}$ . Further, for neural networks based classifiers we can define  $C_p : X \rightarrow \Delta Y$  as the function that takes in data and produces a probability distribution over labels. Thus,  $C = \max\{C_p(x)\}$ , where  $\max$  provides the maximum component of the vector  $C_p(x)$ . Let  $\overline{sim}(x, x')$  denote the dissimilarity between  $x$  and  $x'$ . Let  $H(p, q)$  denote the cross entropy  $-\sum_i p_i \log(q_i)$ . In particular, let  $H(p)$  denotes the entropy given by  $H(p, p)$ . For this paper, we assume  $X$  is the set of legitimate images (and not garbage images or ambiguous images). Legitimate images are different for different domains, e.g., they are digits for digit classification. Given a label  $y$ , let  $Cat(y)$  denote the categorical probability distribution with the component for  $y$  set to 1 and all else 0. Let  $\overline{opsim}(y, y')$  denote the dissimilarity between output distributions  $y, y' \in \Delta Y$ .

## 3.1 Robustness

We first introduce some concepts from PAC learning [1], in order to present the formal results in this section. It is assumed that data points arise from a fixed but unknown distribution  $\mathcal{P}$  over  $X$ . We denote the probability mass over a set  $Z \subset X$  as  $\mathcal{P}(Z)$ . A loss function  $l(y_x, C(x))$  captures the loss of predicting  $C(x)$  when the true label for  $x$  is  $y_x$ . As we are focused on classification, we restrict ourselves to the ideal 0/1 loss, that is, 1 for incorrect classification and 0 otherwise. A classifier  $C$  is chosen that minimizes the empirical loss over the  $n$  training data points  $\sum_{i=1}^n l(y_{x_i}, x_i)$ . Given enough data, PAC learning theory guarantees that  $C$  also minimizes the expected loss  $\int_X l(y_x, C(x))\mathcal{P}(x)$ . Given, 0/1 loss this quantity is just  $\mathcal{P}(M_C(X))$ , where  $M_C(X) \subset X$  denote the region where the classifier  $C$  misclassifies. Accuracy for a classifier is then just  $1 - \mathcal{P}(M_C(X))$ . In this paper we will assume that the amount of data is always enough to obtain low expected loss. Observe that a classifier can achieve high accuracy (low expected loss) even though its predictions in the low probability regions may be wrong.

All classifier families have a capacity that limits the complexity of separators (hypothesis space) that they can model. A higher capacity classifier family can model more non-smooth separators<sup>2</sup>. Previous work [11] has conjectured that adversarial examples abound due to the low capacity of the classifier family used. See Figure 1A for an illustration.

**Adversarial exploitable space:** Define  $E_{C, \epsilon}(X) = M_C(X) \cup \{x \mid \overline{sim}(x, M_C(X)) \leq \epsilon\}$ , where  $\overline{sim}$  is a dissimilarity measure that depends on the domain and  $\overline{sim}(x, M_C(X))$  denotes the lowest dissimilarity of  $x$  with any data point in  $M_C(X)$ . For image classification  $\overline{sim}$  can just be the  $l_2$  (Euclidean) distance:  $\sqrt{\sum_i (x_i - x'_i)^2}$  where  $i$  indexes the pixels.  $E_{C, \epsilon}(X)$  is the adversarial exploitable space, as this space includes all points that are either mis-classified or can be mis-classified by a minor  $\epsilon$ -perturbation. Note that we assume that any already present mis-classifications of the classifier is exploitable by the adversary without the need of any perturbation. For example, if a stop sign image in a dataset is mis-classified then an adversary can simply use this image as is to fool an autonomously driven vehicle.

**Robustness:** Robustness is simply defined as  $1 - \mathcal{P}(E_{C, \epsilon}(X))$ . First, it is easy to see that robustness is a strictly stronger concept than accuracy, that is, a classifier with high robustness has higher accuracy. We believe this property makes our definition more natural than other current definitions. Further, another readily inferable property from the definition of  $E_{C, \epsilon}$  that we utilize later is that a classifier  $C'$  with  $M_{C'}(X) \subset M_C(X)$  is more robust than classifier  $C$  in the same setting. We call a classifier  $C'$  perfect if the robustness is 100%.

There are a number of subtle aspects of the definition that we elaborate upon below:

- A 100% robust classifier can still have  $M_{C'}(X) \neq \phi$ . This is because robustness is still defined w.r.t.  $\mathcal{P}$ , for example, large compact regions of zero probability with small sub-region of erroneous prediction far away from the boundary can still make robustness 100%. However,  $M_{C'}(X) = \phi$  provides

<sup>1</sup>[https://drive.google.com/open?id=1CBaHsU6IL9jQ4UN\\_2yYUteUsbLP2aKMKk](https://drive.google.com/open?id=1CBaHsU6IL9jQ4UN_2yYUteUsbLP2aKMKk)

<sup>2</sup>While capacity is defined for any function class [1] (includes deep neural networks), the value is known only for simple classifiers like single layered neural networks.

100% robustness for any  $\mathcal{P}$ . Thus, robustness based on just  $M_{C'}(X)$  and not  $\mathcal{P}$  is a stronger but much more restrictive concept of robustness than ours.

- A perfect classifier (100% robust) is practically impossible due to large data requirement especially as the capacity of the classifier family grows. As shown in Figure 1 low capacity classifiers cannot model complex separators, thus, large capacity is required to achieve robustness. On the other hand, classifiers families with large capacity but not enough data tend to overfit the data [1]. Thus, there is a delicate balance between the capacity of the classifier family used and amount of data available. The relation between amount of data and capacity is not very well understood for Dep Neural Networks. In any case, perfect robustness provides a goal that robust classifiers should aim to achieve. In this paper, for the purpose of defense, we seek to increase the robustness of classifiers.
- Robustness in practice may apparently seem to be computable by calculating the accuracy for the test set and the adversarially perturbed test set for any given dataset, which we also do and has been done in all prior work. However, this relies on the fact that the attack is all powerful, i.e., it can attack *all* perturb-able points. It is easy to construct abstract examples with probability measure zero mis-classification set (single mis-classified point in a continuous Euclidean space) that is computationally intractable for practical attacks to discover. A detailed analysis of computing robustness is beyond the scope of this paper and is left for future work.
- The definition can be easily extended to weigh some kinds of mis-classification more, if required. For example, predicting a malware as benign is more harmful than the opposite erroneous prediction. For our focus area of image classification in this paper, researchers have generally considered all mis-classification equally important. Also the  $\overline{sim}$  function in the definition is reasonably well agreed upon in literature on adversarial learning as the  $l_2$  distance; however, we show later in experiments that  $l_2$  distance does not capture similarity well enough. Instantiating the definition for other domains such as malware classification requires exploring  $\overline{sim}$  further such as how to capture that two malwares are functionally similar.

Lastly, compared to past work [9, 27], our robustness definition has a clear relation to accuracy and not orthogonal to it. Also, our definition uses the ideal 0/1 loss function rather than an approximate loss function  $l$  (often used in training due to smoothness) as used in other definitions [7? ]. We posit that the 0/1 loss measures robustness more precisely, as these other approaches have specified the adversary goal as aiming to perturb in order to produce the maximum loss within an  $\epsilon$  ball  $B(x, \epsilon)$  of any given point  $x$ , with the defender expected loss defined as  $\int_X \max_{z \in B(x, \epsilon)} l(y_x, C(z)) \mathcal{P}(x)$ . However, this means that even if the class is same throughout the  $\epsilon$  ball, with a varying  $l$  the adversary still conducts a supposed ‘‘attack’’ and increases loss for the defender without flipping labels. For example, the well-known hinge loss varies rapidly within one of the classes and such supposed attacks could lead to an overestimation

of the loss for defender and hence underestimate robustness. Further, use of an approximation in the definition allows an adversary to bypass the definition by exploiting the approximation by  $l$  when the true loss is 0/1. It is an interesting question for future on what kind of approximations can help in computing robustness within reasonable error bounds.

Finally, we analyze if accuracy is ever suitable to capture robustness. First, we make a few mild technical assumptions that there exists a density  $p(x)$  for the data distribution  $\mathcal{P}$  over  $X$ ,  $X$  is a metric space with metric  $d$  and  $vol(X) = 1$ . We have the following result:

**THEOREM 3.1.** *1 - a accuracy implies at least 1 - (a + v + Kε/T) robustness for any output C if*

- For all  $x \in X$ ,  $\overline{sim}(x, x') \geq Td(x, x')$  for some  $T > 0$ .
- $M_C(X)$  lies in a low density region, that is, for all  $x \in M_C(X)$  we have  $p(x) \leq v$  for some small  $v$ .
- $p(x)$  is  $K$ -Lipschitz, that is,  $|p(x) - p(x')| \leq Kd(x, x')$  for all  $x, x' \in X$ .

The first two conditions in the above result are quite natural. In simple words, the first two conditions says dissimilarity increases with distance (high  $T$ ) and the regions that the output classifier predicts badly has low amount of data in the data-set (low  $v$ ).

However, the final condition may not be satisfied in many natural settings. This condition states that the data distribution must not change abruptly (low  $K$ ). This is required as the natural behavior of most classifiers is to predict bad in a low data density region and if this region is near a high data density region, the adversary can successfully modify the data points in the high density region causing loss of robustness. But in high dimensional spaces, data distribution is quite likely to be not distributed smoothly with many pockets or sub-spaces of zero density as pointed out in a recent experimental work [? ]. Thus, data distribution is an important contributing factor that determines robustness.


### 3.2 ALN

Our goal is to train a neural network ALN to produce samples in the misclassification region of a given neural network based classifier. The ALN acts on a data point  $x$  producing  $x'$ . Thus, we choose the following loss function for the ALN that takes into account the output for the given classifier:

$$\alpha \overline{sim}(x, x') - \overline{opsim}(Cat(y_x), C_p(x')),$$

The input dissimilarity term in the loss function aims to produce data points  $x'$  that are similar to the original input  $x$  while the output dissimilarity term aims to maximize the difference between the true label of  $x$  and prediction of  $C$  on  $x'$ . The  $\alpha$  is a weight that is tuned through a simple search. Observe that this loss function is general and can be used with any classifier (by inferring  $C_p$  from  $C$  in case of specific non neural network based classifiers). For the image classification problem we use the  $l_2$  distance for  $\overline{sim}$ . For  $\overline{opsim}$  we tried a number of functions, but the best performance was for the  $l_1$  loss  $\|Cat(y_x) - C_p(x')\|_1$ .

Note that an alternate loss function is possible that does not use the actual label  $y_x$  of  $x$ , rather using  $C_p(x)$ . This would also work assuming that the classifier is good; for poor classifiers a lot of the data points are as it is mis-classified and hence adversarial example generation is interesting only for good classifiers. Further,



**Figure 1: Intuition behind ALN and DLN. (A) shows a linear classifier (low capacity) is not able to accurately model a non-linear boundary. (B) shows the ALN as the distribution mapping function  $F$ . (C) shows that DLN does the reverse mapping of ALN.**

this choice would allow using unlabeled data for conducting such an attack, making attack easier for an attacker. However, in our experiments we use the more powerful attack using the labels.

Next, we provide a formal intuition of what ALN actually achieves. Any adversarial example generation can be seen as a distribution transformer  $F$  such that acting on the data distribution  $\mathcal{P}$  the resultant distribution  $F(\mathcal{P})$  has support mostly limited to  $M_C(X)$ . The support may not completely limited to  $M_C(X)$  as the attacks are never 100% effective. Also, attacks in literature aim to find points in  $M_C(X)$  that are close to given images in the original dataset. ALN is essentially a neural network representation of such a function  $F$  against a given classifier  $C$ . See Figure 1B for an illustration. We return to this interpretation in the next sub-section to provide a formal intuition about the DLN defense.

Lastly, we show in our experiments that ALN produces adversarial examples whose perturbations are roughly of the same order as the prior attack FGSM. We categorize these as high perturbation attacks. On the other the attack CW produces adversarial perturbation with very small perturbations, we call such attacks low perturbation attacks. As mentioned earlier, we provide two separate defenses for these two types of attacks. Both these defense can be used simultaneously to defend against both types of attacks.

### 3.3 DLN

Our first defense approach is to insert a neural network DLN  $D$  between the input and classifier so that  $D$  sanitizes the input enabling the classifier to correctly classify the input. Each data point for training DLN has three parts:  $x'$  is the image to sanitize,  $x$  is the expected output and  $y_x$  is the correct label of  $x$ .  $x$  are images from the provided dataset, and there are two possibilities for  $x'$ : (1)  $x' = x$  so that DLN attempts to satisfy  $C(D(x)) = y_x$ , even if  $C(x) \neq y_x$ , and (2)  $x' = A(x)$  so that DLN undoes the attack and make the classifier  $C$  correctly classify  $x'$ .

We formulate a loss function for DLN that, similar to ALN, has two terms:  $\overline{sim}(x, D(x'))$  that aims to produce output  $D(x')$  close to  $x$  and  $\overline{opsim}(Cat(y_x), C_p(D(x')))$  that aims to make the classifier output on  $D(x')$  be the same as  $y_x$ . Thus, the loss function is

$$\overline{sim}(x, D(x')) + \overline{opsim}(Cat(y_x), C_p(D(x'))).$$

In this paper we only use  $\alpha = 1$ . Note that the attack  $A$  is used as a black box here to generate training data and is not a part of the loss

function. After training the DLN, our new classifier is  $C'$  which is  $C$  prepended by the DLN. The working of DLN can be interpreted as an inverse map  $F^{-1}$  for the mapping  $F$  induced by the attack  $A$ . See Figure 1C for an illustration. For the image classification problem we use the  $l_2$  distance for  $\overline{sim}$ . For  $\overline{opsim}$  we tried a number of functions, but the best performance was for the cross-entropy loss  $H(Cat(y_x), C_p(D(x')))$ .

An important point to note is that the original classifier  $C$  is unchanged. What this ensures is the mis-classification space  $M_C(X)$  does not change and allows us to prove an important result about  $C'$  under certain assumptions. For the sake of this result, we assume that attacks  $A$  generate adversarial examples in a sub region  $M_{C,A}(X) \subset M_C(X)$ . We also assume a good DLN  $D$ , that is,  $C(D(x))$  is correct for a non-empty subset  $Z \subset M_{C,A}(X)$  and  $C(D(x))$  continues to be correct for all  $x \notin M_C(X)$ . Then, we prove

LEMMA 3.2. Assuming  $M_{C,A}(X) \subset M_C(X)$ , DLN is good as defined above, and  $M_{C,A}(X) \neq \phi$ , then  $M_{C'}(X) \subset M_C(X)$

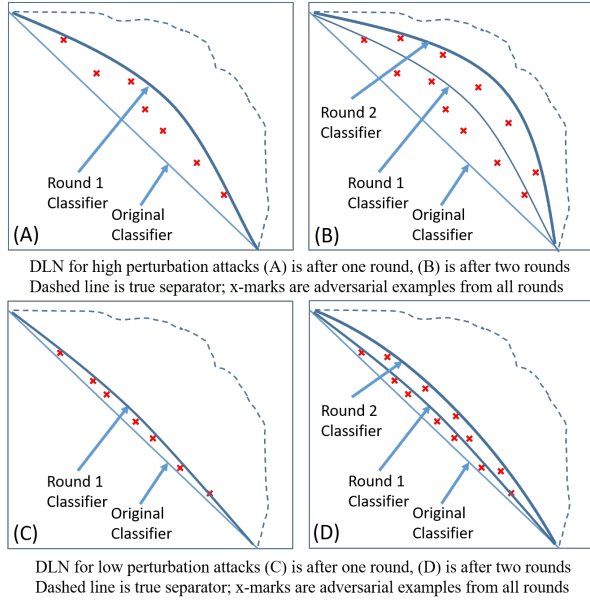
PROOF. Since DLN does not decrease the performance of  $C$  on points outside  $M_C(X)$ ,  $C'$ 's prediction on inputs outside  $M_C(X)$  is correct, hence  $M_{C'}(X) \subseteq M_C(X)$ . Any data point not mis-classified by a classifier does not belong to its mis-classification space. Good sanitization by DLN makes  $C'$  predict correctly on  $Z \subset M_{C,A}(X)$ , which makes  $M_{C,A}(X) \cap M_{C'}(X) \subset M_{C,A}(X)$ . Thus, we can claim the result in the lemma statement.  $\square$

While the above proof is under ideal assumptions, it provides an intuition to how the defense works. Namely, the reduction in the adversarial exploitable space makes the new classifier more robust (see robustness properties earlier). This also motivates the generalization of this technique to multiple attacks presented in the next sub-section.

### 3.4 Repeated DLN Against Multiple Attacks

The above DLN can be naturally extended to multiple attacks, say  $A_1, \dots, A_n$ . The only change required is to feed in all possible adversarial examples  $A_1(x)$ 's,  $\dots, A_n(x)$ 's. It is straightforward to see that under assumptions of Lemma 3.2 for all the attacks, the resultant classifier  $C'$  has an adversarial example space  $M_{C'}(X)$  that removes subsets of  $M_{C,A_i}(X)$  for all  $A_i \in \mathcal{A}$  from  $M_C(X)$ . This provides, at least theoretically under ideal assumptions, a monotonic robustness improvement property with increasing number of attacks for the DLN based approach. In fact, if all the attacks combined act as a generator for all the points in  $M_C(X)$ , then given enough data and perfect sanitization the resultant classifier  $C'$  tends towards achieving  $M_{C'}(X) = \phi$  which essentially would make  $C'$  a perfect classifier. Perfect classifiers have no adversarial examples.

However, attacks rarely explore all of the mis-classified space, which is why new attacks have defeated prior defense techniques. Even for our approach, attacks successfully work against the DLN that has been trained only once (accuracy numbers are in Experiments). However, DLN allows for easy retraining (without retraining the classifier) as follows: repeatedly attack and re-learn a DLN in rounds, that is, conduct an attack on the classifier obtained in every round and train a DLN in a round using the attacked training data from all the previous rounds and the original training data. More formally, at round  $i$  our training data consists of  $i$  copies of



**Figure 2: Intuition behind working of repeated DLN against high and low perturbation attacks. (A),(B) shows a high perturbation attack causes a faster improvement in resultant classifier. Further, beyond some rounds the attack does not work as it can only find adversarial examples with high perturbation. (C),(D) shows a low perturbation attack causes a slow improvement in resultant classifier.**

original training data and  $i$  instances of attacked training data from previous rounds. Observe that we add copies of the original training data in each round, this is because the adversarial data swamps out the original training data and accuracy suffers in regions where the original training data is distributed. See Figure 2 for an illustration of how repeated DLN works.

The following result provides formal intuition for this approach:

**LEMMA 3.3.** *Assume the following conditions hold for every round  $i$ :  $M_{C_{i-1},A}(X) \subset M_{C_{i-1}}(X)$  and the DLN  $D_i$  has good memory, which means that given there exists a largest set  $Z_i \subset M_{C_{i-1},A}(X)$  which the DLN  $D_i$  correctly sanitizes so that  $C(D_i(x))$  is correct for all  $x \in Z_i$  then  $Z_{i-1} \subset Z_i$ . That is DLN  $D_i$  can correctly sanitize data points that the previous round DLN did plus possibly more data points. Further,  $C(D_i(x))$  continues to be correct for all  $x \notin M_C(X)$ . Then the classifier  $C_n$  after  $n$  rounds satisfies  $M_{C_n}(X) \subset M_{C_{n-1}}(X)$ .*

**PROOF.** Arguing similarly to Lemma 3.2 we can show that  $M_{C_n}(X) \subseteq M_C(X)$  due to the correct classification outside of  $M_C(X)$ . Further, it is easily inferable that  $M_{C_n}(X) = M_C(X) \setminus Z_n$  given  $Z_n$  is a subset of  $M_C(X)$  and given the largest such set condition on  $Z_i$ . Then, the good memory property leads to the required result.  $\square$

The attack-defense competition technique is somewhat akin to GANs [10]. However, there is a big difference, since in every round the dataset used to train the DLN grows. Practically, this requires DLN to have a large capacity in order to be effective; also depending on the capacity and the size of dataset over or under

fitting problems could arise, which needs to be taken care of in practice. Also, the training become more expensive over rounds with increasing data size. In particular, low perturbation attacks are not defeated with few rounds. We do observe improvement with the low perturbation CW attack over rounds, but the improvement is very small, as represented visually in Figure 2. The main reason is that low perturbation attacks only exposes a very small volume of misclassified space, thus, it would require a huge number of rounds for repeated DLN to reduce the mis-classified space to such a small volume that cannot be attacked. This motivates our next approach of noise augmented classifier.

### 3.5 NAC

Figure 2 also provides a hint on how to overcome low perturbation attacks. In order to achieve low perturbation, such attacks rely a lot on the exact classifier boundary. Thus, masking the classifier boundary can fool low perturbation attacks. We achieve this by adding a small noise to the logits of the neural network calling the resultant classifier a noise augmented classifier (NAC). This noise should be small enough that it does not affect the classification of original data points by much, but is able to mis-lead the low perturbation attack. Also, following our intuition NAC should not provide any defense against high perturbation attacks, which we indeed observe in our experiments. However, observe that DLN and NAC can be used simultaneously, thus, providing benefits of both defense which we show in our experiments.

Further, a natural idea to bypass the defense provided by NAC is to take the average of multiple logit outputs for the same given input image (to cancel the randomness) and then use the average logits as the logits required for the CW attack. We show experimentally that this idea does not work effectively even after averaging over many logit outputs.

## 4 EXPERIMENTS FOR ATTACKS

All our experiments, including for DLN and NAC, were conducted using the Keras framework on a NVIDIA K40 GPU. The learning problem can be solved using any gradient-based optimizer. In our case, we used Adam with learning rate 0.0002. We use two well-known datasets: MNIST digits and CIFAR-10 colored images.

We consider two classifiers one for MNIST and one for CIFAR-10: we call them  $C_M$  and  $C_C$ . These classifiers are variants of well-known architectures that achieved state-of-the-art performances on their respective datasets. As stated earlier, we consider three attacks: ALN, FGSM and CW. CW has been referred to in the literature [28] as one of the best attacks till date (at the time of writing of this paper), while FGSM runs extremely fast. For the autoencoder we use a fourteen hidden layer convolutional architecture. Our code is publicly available, but the github link is elided in this submitted version for the review process. For our experiments we pre-process all the images so that the pixels values lie between  $[-0.5, 0.5]$ , so all components (attacks, autoencoders, classifiers) work in space  $[-0.5, 0.5]$ . We use FGSM with values of 0.03 and 0.01 for its parameter  $\epsilon$  on MNIST and CIFAR, respectively.

Observe that all these attacks work against a given classifier  $C$ , thus, we use the notation  $A(C, \cdot)$  to denote the attack  $A$  acting on an image  $x$  to produce the adversarial example  $A(C, x)$  ( $A$  can






Figure 3: Targeted attacks by ALN: target class on bottom

be any of the three attacks).  $A(C, Z)$  denotes the set of adversarial examples  $\{A(C, x) \mid x \in Z\}$ . We report accuracies on various test sets: (1) original test dataset ( $OTD$ ): these images are the original test dataset from the dataset under consideration, (2)  $A(C, OTD)$  is the adversarially perturbed dataset using attack  $A$  against classifier  $C$ , for example, this could be  $FGSM(C_M, OTD)$ . We also report distortion numbers as has been done in literature [5]. Distortion measures how much perturbation on average was added by an attack and is meant to capture how visually similar the image is to the original image. Distortion is measured as the average over all test images of the  $l_2$  distance between the original and perturbed image.

**Results:** Untargeted attacks refers to attack that aim to produce mis-classification but not with the goal of making the classifier output a specific label. Targeted attacks aim to produce an adversarial example that gets classified as a given class  $y$ . It is also possible to modify ALN to perform targeted attacks. This is achieved by modifying the loss function to use a positive *opsim* term, like the DLN loss function, but using the target class label  $y$  instead of the original class label  $y_x$  in the *opsim* term. Then, we can perform an ALN attack in two ways: ALNU uses the ALN loss function as stated and ALNT constructs a targeted attack per class label differing from the original label and chooses the one with least distortion. Figure 3 shows an example of targeted attack with different target labels. Table 1 shows this approach for MNIST with the targeted ALNT version performance better than other attacks.

Test data type	Accuracy	Distortion
$OTD$	99.45 %	–
$FGSM(C_M, OTD)$	0.72 %	14.99
$CW(C_M, OTD)$	0.03 %	1.51
$ALNU(C_M, OTD)$	1.65 %	4.43
$ALNT(C_M, OTD)$	0.0 %	4.34

Table 1: Attacks on MNIST Dataset

Table 2 shows the result of untargeted attacks using ALN, FGSM and CW on the CIFAR-10 dataset. We can see that ALN, just like FGSM, produces slightly higher adversarial accuracy for MNIST, but the distortion of FGSM is much higher. This does result in a large difference in visual quality of the adversarial examples produced—see Figure 4 for randomly chosen 25 perturbed images using ALN and Figure 5 for randomly chosen 25 perturbed images using FGSM. Also, we can attribute the higher distortion of ALN for CIFAR as compared to MNIST partially to the CIFAR being a higher dimensional space problems and the same capacity of the autoencoder we used for CIFAR and MNIST.

Test data type	Accuracy	Distortion
$OTD$	84.59 %	–
$FGSM(C_C, OTD)$	4.21 %	10.03
$CW(C_C, OTD)$	0 %	0.18
$ALNU(C_C, OTD)$	6.16 %	2.57
$ALNT(C_C, OTD)$	8.21 %	3.01

Table 2: Attacks on CIFAR-10 Dataset



Figure 4: Untargeted attack by ALN for CIFAR-10



Figure 5: Attack using FGSM for CIFAR-10

## 5 EXPERIMENTS FOR DEFENSE

For defense, we denote the new classifier using the DLN trained against any attack  $A$  or NAC modified classifier as  $C'_M$  (for MNIST) or  $C'_C$  (for CIFAR). Also, we test accuracies on test data adversarially perturbed by attacks against the new classifiers, for example, following our convention one such dataset would be denoted as  $A(C'_M, OTD)$ . For defense we focus on defending against known attacks in literature.

### 5.1 DLN Defense Against Single Attack

Table 3 shows the results when the DLN is trained to defend against specific attacks using MNIST dataset to yield a new classifier  $C'_M$ .

Along expected lines, the accuracy on *OTD* drops slightly for all the cases where DLN is trained against FGSM or CW. However, the new classifier  $C'_M$  is able to perform quite good on adversarial examples that were produced by any attack on the original classifier  $C_M$ , for example,  $C'_M$  gets an accuracy of 95% of the adversarial examples  $CW(C_M, OTD)$  that was produced by CW attacking  $C_M$ . Lastly, when attacked again the new classifier  $C'_M$  is still not resilient to attacks as shown by the low accuracies for  $FGSM(C'_M, OTD)$  and  $CW(C'_M, OTD)$ .

One number that stands out is the success of the new classifier  $C'_M$  in correctly classifying the adversarial example generated by CW for the original classifier  $C_M$ . This supports our hypothesis that CW is very sensitive to the exact classifier boundary, and a newer classifier with a slightly different boundary is able to defeat prior adversarial examples. Of course, CW is able to attack  $C'_M$  against quite successfully, which we later show can be defended by our NAC approach. For FGSM, we show in later sections that the performance of the classifier greatly improves when DLN is repeatedly trained against FGSM—revealing that the DLN approach is flexible enough to keep on improving its performance against high perturbation attacks.

DLN Trained	Test data type	Accuracy	Distortion
FGSM	<i>OTD</i>	96.77 %	-
FGSM	$FGSM(C_M, OTD)$	88.5 %	4.55
FGSM	$FGSM(C'_M, OTD)$	13.75 %	6.98
CW	<i>OTD</i>	98.6 %	-
CW	$CW(C_M, OTD)$	95.42 %	5.77
CW	$CW(C'_M, OTD)$	0.14 %	3.5

**Table 3: New DLN prepended classifier  $C'_M$  for MNIST**

### 5.2 Repeated DLN

In this section, we run DLN repeatedly as described earlier in Section 3.4. We cut off the experiments when a single round took more than 24 hours to solve. We show the results for MNIST in Table 4 showing a clearly increasing trend in accuracy on adversarial examples produced by FGSM attacking the newer classifier, revealing increasing robustness. For CIFAR, the approach becomes too computationally expensive within two rounds. Thus, while the DLN approach is promising, as stated earlier it is computationally expensive, more so with larger and complex data sets. Further, as stated earlier running DLN against low perturbation attacks like CW does not show much improvement. However, we tackle that next using the NAC defense approach.

### 5.3 NAC defense

Recall that the NAC defense works by adding noise to the logits layer of the neural network classifier to produce a new classifier  $C'_M$  for MNIST and  $C'_C$  for CIFAR. We use Gaussian noise with 0 mean and variance 1. In this section, we show that the NAC defense is able to produce classifiers that are resilient to CW attack. Further, the new classifier’s accuracy on the original test data-set is nearly

Round	Acc. <i>OTD</i>	Acc. $FGSM(C_i, OTD)$	Distortion
0	99.36 %	0.72 %	14.99
1	97.70 %	13.70 %	13.63
2	97.61 %	24.86 %	14.58
3	97.95 %	43.39 %	14.73
4	97.79 %	52.88 %	14.57
5	97.77 %	56.57 %	14.52

**Table 4: Classifier trained repeatedly against FGSM for MNIST**

unchanged. This can be seen in Table 5. The second line in that table shows that the CW attack completely fails as the accuracy on the adversarial examples in 93%. However, it can also be observed that the new classifier is not resilient to attack by FGSM, as shown by the third line in that table. This follows the intuition we provided in Figure 2. For CIFAR, Table 6 shows that NAC is able to overcome CW to a large extent.

Attack	Test data type	Accuracy	Distortion
-	<i>OTD</i>	99.36 %	-
CW	$CW(C'_M, OTD)$	93.60 %	1.49
FGSM	$FGSM(C'_M, OTD)$	0.74 %	14.99

**Table 5: Accuracy of NAC Classifier  $C'_M$  for MNIST**

Attack	Test data type	Accuracy	Distortion
-	<i>OTD</i>	84.67 %	-
CW	$CW(C'_M, OTD)$	77.70 %	0.17
FGSM	$FGSM(C'_M, OTD)$	4.19 %	10.04

**Table 6: Accuracy of NAC Classifier  $C'_M$  for CIFAR-10**

As stated earlier, a natural idea to attack NAC would be to query an image  $n$  times and then average the logits before using it for the CW attack. This augmented attack does make CW more effective but not by much. Table 7 shows that the accuracy on the adversarial example generated for  $C'_M$  remains high. Moreover, more queries make it more difficult to conduct the CW attack in practice (as the adversary may be query limited), while also causing a small increase (1% with 5000 sample) in the already high runtime of CW.

$n$	Adversarial accuracy	Distortion
500	95.14 %	1.51
5000	82.07 %	1.51

**Table 7: Accuracy of NAC Classifier  $C'_M$  against improved CW for MNIST**



## 5.4 Defense Against Multiple Attacks

Finally, we show that DLN and NAC can work together. We show this by presenting the accuracy on the adversarial example generated in each round of DLN repetition when the classifier  $C_i$  after each round is augmented with NAC and attacked by FGSM and CW both. See Table 8. One observation is that NAC’s performance decreases slightly over rounds stabilizing at 79%, while the accuracy for original test set and FGSM perturbed test set stays almost exactly same as Table 4.

Round	Acc. $OTD$	Acc. FGSM( $C_i, OTD$ )	Acc. CW( $C_i, OTD$ )
0	99.36 %	0.72 %	94.2
1	97.70 %	13.72 %	93.7
2	97.73 %	24.28 %	84.7
3	97.60 %	43.20 %	83.3
4	97.64 %	53.17 %	79
5	97.73 %	56.45 %	79.3

**Table 8: Classifier trained repeatedly against FGSM for MNIST and augmented with NAC in each round**

## 6 RELATED WORK

A thorough survey of security issues in machine learning is present in surveys [25] and some of the first results appeared in [8, 19]. Here we discuss the most closely related work.

**Attacks:** Most previous attack work focuses on adversarial examples for computer vision tasks. Multiple techniques to create such adversarial examples have been developed recently. Broadly, such attacks can be categorized as either using costs gradients [3, 11, 14, 20] or the forward gradient of the neural network [24] and perturbing along most promising direction or directly solving an optimization problem (possibly using gradient ascent/descent) to find a perturbation [5, 21]. In addition, adversarial examples have been shown to transfer between different network architectures, and networks trained on disjoint subsets of data [22, 26]. Adversarial examples have also been shown to translate to the real world [15], that is, adversarial images can remain adversarial even after being printed and recaptured with a cell phone camera. Attacks on non-neural networks have also been explored in literature [3]. Our approach is distinctly different from all these approaches as we pose the problem of generating adversarial samples as a generative learning problem, and demonstrate generation of adversarial examples given access to any given classifier. Our approach also applies to any classifier that output class probabilities and not just neural networks.

**Defense:** Also, referred to as robust classification in many papers, defense techniques can be roughly categorized into techniques that do (1) adversarial (re)training, which is adding back adversarial examples to the training data and retraining the classifier, often repeatedly [17], or modifying loss function to account for attacks [14]; (2) gradient masking, which targets that gradient based attacks by trying to make the gradient less informative [23]; (3) input modification, which are techniques that modify (typically lower the

dimension) the feature space of the input data to make crafting adversarial examples difficult [28]; (4) game-theoretic formulation, which modifies the loss minimization problem as a constrained optimization with constraints provided by adversarial utility in performing perturbations [16], and (5) filtering and de-noising, which aims to detect/filter or de-noise adversarial examples (cited below).

Our defense approach differs from the first four kinds of defense as our DLN approach never modify the classifier or inputs but add a sanitizer (DLN) before the classifier. First, this increases the capacity of the resultant classifier  $C'$ , so that it can model more complex separators, which is not achieved when the classifier family stays the same. Further, our defense is agnostic to the type of attack and does not utilize properties of specific types of attacks. Interestingly, the DLN approach can be used with any classifier that output class probabilities and not just neural networks. Further, NAC is a very minor modification to the classifier that, distinct from other randomized approaches [?] that randomize over multiple classifiers, aims to mask the classifier boundary. Also, NAC can work with other defenses unlike techniques that modify inputs to try and defend against CW [28].

More closely related to our work are some defense techniques that have focused on detecting and filtering out adversarial samples [12, 18] or de-noising input [13]; here the filter or de-noiser with the classifier could be considered as a larger neural network. However, unlike these work, our goal for DLN is targeted sanitization. Moreover, recent attack work [4] have produced attack techniques to defeat many known detection techniques. Our technique provides the flexibility to be resilient against more powerful attacks by training the DLN with such an attack for high perturbation attacks or using NAC for low perturbation attacks.

Lastly, two concurrent unpublished drafts (available online) have *independently and simultaneously* proposed an attack [2] similar to ALN and a defense [6] apparently similar to DLN. The difference for the attack work is in using the class label vs classifier output in *opsim* term for the attack. For the defense work, we differ as we show how DLN technique extends to multiple attacks and can be repeatedly used in an attack-defense competition. Moreover, unlike these drafts, we provide another defense technique NAC that works against CW, define robustness and show that our defense techniques approximately aims to achieve our definition of robustness. Further, our formal reasoning reveals the underlying nature of attacks and defenses.

## 7 CONCLUSION AND FUTURE WORK

Our work provides a new learning perspective of the adversarial examples generation and defense problems with a formal intuition of how these approaches work, using which we were able to defend against multiple attacks including the potent CW. Further, unlike past work, our defense technique does not claim to a catchall or specific to any attack; in fact, it is flexible enough to possibly defend against any attack. Posing the attack and defense as learning problems allows for the possibility of using the rapidly developing research in machine learning itself to make the defense more effective in future, for example, by using a different specialized neural

## A Learning and Masking Approach to Secure Learning

network architecture rather than an autoencoder. A number of variations of our theory and tuning of the application framework provides rich content for future work.

## REFERENCES

- [1] Martin Anthony and Peter L. Bartlett. 2009. *Neural Network Learning: Theoretical Foundations* (1st ed.). Cambridge University Press, New York, NY, USA.
- [2] Shumeet Baluja and Ian Fischer. 2017. Adversarial Transformation Networks: Learning to Generate Adversarial Examples. *CoRR* abs/1703.09387 (2017). <http://arxiv.org/abs/1703.09387>
- [3] Battista Biggio, Iginio Corona, Davide Maiorca, Blaine Nelson, Nedim Šrđić, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. 2013. Evasion attacks against machine learning at test time. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 387–402.
- [4] Nicholas Carlini and David Wagner. 2017. Adversarial Examples Are Not Easily Detected: Bypassing Ten Detection Methods. *arXiv preprint arXiv:1705.07263* (2017).
- [5] Nicholas Carlini and David Wagner. 2017. Towards evaluating the robustness of neural networks. In *Security and Privacy (SP), 2017 IEEE Symposium on*. IEEE, 39–57.
- [6] Xinyun Chen, Bo Li, and Yevgeniy Vorobeychik. 2016. Evaluation of Defensive Methods for DNNs against Multiple Adversarial Evasion Models. (2016). <https://openreview.net/forum?id=ByToKu9ll&notelid=ByToKu9ll>.
- [7] Moustapha Cisse, Piotr Bojanowski, Edouard Grave, Yann Dauphin, and Nicolas Usunier. 2017. Parseval Networks: Improving Robustness to Adversarial Examples. *arXiv preprint arXiv:1704.08847* (2017).
- [8] Nilesh Dalvi, Pedro Domingos, Sumit Sanghai, Deepak Verma, et al. 2004. Adversarial classification. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 99–108.
- [9] Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. 2015. Fundamental limits on adversarial robustness. In *Proc. ICML, Workshop on Deep Learning*.
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.
- [11] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and Harnessing Adversarial Examples. *CoRR* abs/1412.6572 (2014). <http://arxiv.org/abs/1412.6572>
- [12] Kathrin Grosse, Praveen Manoharan, Nicolas Papernot, Michael Backes, and Patrick McDaniel. 2017. On the (Statistical) Detection of Adversarial Examples. *arXiv preprint arXiv:1702.06280* (2017).
- [13] Shixiang Gu and Luca Rigazio. 2014. Towards deep neural network architectures robust to adversarial examples. *arXiv preprint arXiv:1412.5068* (2014).
- [14] Ruitong Huang, Bing Xu, Dale Schuurmans, and Csaba Szepesvári. 2015. Learning with a strong adversary. *arXiv preprint arXiv:1511.03034* (2015).
- [15] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. 2016. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533* (2016).
- [16] Bo Li and Yevgeniy Vorobeychik. 2014. Feature cross-substitution in adversarial classification. In *Advances in neural information processing systems*. 2087–2095.
- [17] Bo Li, Yevgeniy Vorobeychik, and Xinyun Chen. 2016. A General Retraining Framework for Scalable Adversarial Classification. *arXiv preprint arXiv:1604.02606* (2016).
- [18] Xin Li and Fuxin Li. 2016. Adversarial Examples Detection in Deep Networks with Convolutional Filter Statistics. *arXiv preprint arXiv:1612.07767* (2016).
- [19] Daniel Lowd and Christopher Meek. 2005. Adversarial learning. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*. ACM, 641–647.
- [20] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. 2015. DeepFool: a simple and accurate method to fool deep neural networks. *CoRR* abs/1511.04599 (2015). <http://arxiv.org/abs/1511.04599>
- [21] Seyed Mohsen Moosavi Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. 2017. Universal adversarial perturbations. In *Proceedings of 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [22] Nicolas Papernot, Patrick McDaniel, and Ian Goodfellow. 2016. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *arXiv preprint arXiv:1605.07277* (2016).
- [23] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. 2016. Practical black-box attacks against deep learning systems using adversarial examples. *arXiv preprint arXiv:1602.02697* (2016).
- [24] Nicolas Papernot, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Z Berkay Celik, and Ananthram Swami. 2016. The limitations of deep learning in adversarial settings. In *Security and Privacy (EuroS&P), 2016 IEEE European Symposium on*. IEEE, 372–387.
- [25] Nicolas Papernot, Patrick McDaniel, Arunesh Sinha, and Michael Wellman. 2016. Towards the Science of Security and Privacy in Machine Learning. *arXiv preprint arXiv:1611.03814* (2016).
- [26] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199* (2013).
- [27] Beilun Wang, Ji Gao, and Yanjun Qi. 2016. A Theoretical Framework for Robustness of (Deep) Classifiers Under Adversarial Noise. *arXiv preprint arXiv:1612.00334* (2016).
- [28] Weilin Xu, David Evans, and Yanjun Qi. 2017. Feature Squeezing: Detecting Adversarial Examples in Deep Neural Networks. *arXiv preprint arXiv:1704.01155* (2017).