



**HAL**  
open science

# FUNN: Flexible Unsupervised Neural Network

David Vigouroux, Sylvain Picard

► **To cite this version:**

David Vigouroux, Sylvain Picard. FUNN: Flexible Unsupervised Neural Network. 2018. hal-01908392

**HAL Id: hal-01908392**

**<https://hal.science/hal-01908392>**

Preprint submitted on 31 Oct 2018

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# FUNN: Flexible Unsupervised Neural Network

David Vigouroux, Sylvain Picard

IRT Saint Exupery, Toulouse, France

## Abstract

Deep neural networks have demonstrated high accuracy in image classification tasks. However, they were shown to be weak against adversarial examples: a small perturbation in the image which changes the classification output dramatically. In recent years, several defenses have been proposed to solve this issue in supervised classification tasks. We propose a method to obtain robust features in unsupervised learning tasks against adversarial attacks. Our method differs from existing solutions by directly learning the robust features without the need to project the adversarial examples in the original examples distribution space. A first auto-encoder A1 is in charge of perturbing the input image to fool another auto-encoder A2 which is in charge of regenerating the original image. A1 tries to find the less perturbed image under the constraint that the error in the output of A2 should be at least equal to a threshold. Thanks to this training, the encoder of A2 will be robust against adversarial attacks and could be used in different tasks like classification. Using state-of-art network architectures, we demonstrate the robustness of the features obtained thanks to this method in classification tasks.

## 1 Introduction

Neural networks and particularly convolutional neural networks have shown impressive results in many different tasks in computer vision such as object detection, image recognition and segmentation. A down side of these networks is their lack of robustness as described in (4), (6), (21). Indeed, imperceptible perturbations of the input (image, sound...) can confuse the network and completely change the prediction. This kind of input is described as an "Adversarial Example". This misclassification could threaten the security and the safety of critical systems. By using the concept of adversarial examples it is possible to design algorithms

in order to generate inputs that will modify the prediction with a high rate of success and a very little perturbation. Those algorithms are known as "Adversarial Attacks". We commonly distinguish two types of attacks. White box attacks: when the network architecture as well as its weights are known. Black box attacks: when they are not known by the attacker. Our work is focused on untargeted white-box attacks computed using different attack methods.

The field of adversarial generation is an active research field, many methods have been developed in order to generate adversarial examples as described in (22). In this work we use five of the most used attack methods in literature: The Fast Gradient Sign Method (FGSM) (4), Iterative FGSM (7), Single Pixel attack and LocalSearch (13) and Deepfool (11). Those methods are gradient-based or score-based approaches that try to find the minimal perturbation that will modify the model prediction. Several studies have conjectured that this lack of robustness comes from the fact that the adversarial examples are out of the datasets distribution, near the border of decision. Different methods have been developed in order to be robust against adversarial attacks. Some strategies try to project the adversarial examples in the original distribution as in (9), (17), (18). Other approaches propose to add adversarial examples into the training dataset (21), (4), (11), similar to data augmentation. Finally new methods like (12) transform the input data such as it is less sensible to perturbations. The problem with those methods is that they are attack-specific and then not efficient against new or simply different attacking methods. Recently new methods proposed to learn the attacking and defending concept using neural networks. Some methods use Generative Adversarial Networks (GANs) to generate adversarial examples as in (17), (19), (8) or Auto-Encoders as in (3), (20), (1). This paper is an attempt to increase an autoencoder(A2) robustness by adversarial training, using another autoencoder(A1) used to generate perturbed images during training. This new method computes robust features without the need to add any additional networks. Only the weights of the network are modified to obtain a robust network without loss of performance.

## 2 Related Work

### 2.1 Attack strategies

Many attack strategies have been proposed, which can be classified in two categories: black-box and white-box attacks. White-box attacks have access to all the weights and gradients of the classifier while Black-box strategies have access only to the predictions of the network. In this work we will focus on white-box attacks. We use two kinds of white-box attacks, gradient-based and score-based. Gradient-based methods are:

**Fast Gradient Sign Method (FGSM)** Given an image  $x$  and its corresponding label  $y$ , the FGSM attack sets the perturbation  $\delta$  to:

$$\delta = \epsilon \cdot \text{sign}(\nabla_x J(x, y)) \quad (1)$$

FGSM (4) uses the sign of the gradient at each pixel to determine the direction with which to change the corresponding pixel value.

**DeepFool** DeepFool (11) perform an iterative attack with a linear approximation in order to find the closest distance from the original input to the decision boundary of adversarial examples. If  $f$  is a binary differentiable classifier, they used an iterative method to approximate the perturbation by considering  $f$  is linearized around  $x_i$  at each iteration. The minimal perturbation is computed as:

$$\begin{aligned} \underset{\eta_i}{\text{argmin}} \|\eta_i\|_2 \\ \text{s.t. } f(x_i) + \nabla f(x_i)^T \eta_i = 0 \end{aligned} \quad (2)$$

This result can also be extended to the multi-class classifier by finding the closest hyperplanes.

DeepFool provided less perturbation and reduced intensity compared to FGSM and BIM.

Score-based methods are:

**Single Pixel** Change the value of a random pixel to the min or max value of the image. If the perturbed image is not fooling the classifier, the pixel value is reseted to it original value and the algorithm choose another random pixel. The method is iterative and stop if an adversarial image is generated or if the max-step limit is reached.

**Local Search** Use a local search procedure to find a pixel (or a group of pixels) that is critical for the classifier robustness and then modify it value to generate adversarial images. Doing this there might exist a pixel (or a group of pixels) in the adversarial image whose coordinate value could lie outside the valid range of image values. The LocSearchAdv algorithm finds pixel locations to perturb using and then applies a defined transformation function to these selected pixels to generated an adversarial image.

## 2.2 Defense strategies

Mutiple defense methods have been used to increase robustness of deep neural networks against adversarial attacks. This section describes the most used strategies.

**Adversarial training** Adversarial training is an intuitive methods consisting of augmenting the training dataset with adversarial example. Thi method is efficient to increase robustness against the adversarial examples used in training. The main issue with this methods is that it is attack-agnostic and badly transfers to new adversarials examples not used in training. Additionally, it tends to make the model more robust to white-box attacks than to black-box attacks due to gradient masking.

**Defensive distillation** Defensive distillation (15) is a method in which a classifier is trained in two rounds using a variant of the distillation (5) method. This induce learning a smoother network and reducing the amplitude of gradients around input points, increase the robustness

against adversarials. This method have been proven to be inefficient against new, black-box attacks (2).

**Adversarial Detectors** Another idea of defense is to detect adversarial examples (10). For each attack generating method considered, a deep neural network classifier (detector) is trained to tell whether an input is normal or adversarial. The detector was directly trained on both normal and adversarial examples. It showed good performance when the training and testing attack examples were generated from the same process and the perturbation was large enough, but it did not generalize well across different attack parameters and attack generation processes.

**MagNet** MagNet (9) is a method based on adversarial detectors strategy. It trains a reformer network (which is an auto-encoder or a collection of auto-encoders) to move adversarial examples closer to the manifold of legitimate, or natural, examples. It was proven to be an effective strategy against gray-box attacks where the attacker is aware of the network architecture and defenses but do not know it parameters.

### 3 Proposed Method

The previously described defense methods provide intuition on what can make a network robust to adversarial examples. Our strategy differs from them by doing sort of adversarial training with data generated thanks to a separated network during training. The generating network is used only in training, and not in inference unlike in the Adversarial Detectors strategies. This paper is focused only on unsupervised learning strategies, this is the reason why we decided to use Auto-Encoders. An auto-encoder is a deep neural network composed by an encoder responsible of deep features extraction from input image and a decoder responsible of image re-generation from deep features extracted by the encoder. During training the network

is optimised in order to generate images similar to input image. We propose a defense strategy to increase robustness against white box attacks using  $L_2$  norm and demonstrate its efficiency against  $L_\infty$ . The strategy used is to train two Auto-Encoders, an “attacker”: a generator of adversarial examples from a given image and a “defender”: capable to extract the original image from adversarial examples given by the attacker. The attacker and the defender are trained simultaneously and adapt themselves all along of the training until convergence. At the end of training, we expect the deep features extracted by the defender’s encoder to be robust against the images generated by the attacker and that they could be used in further tasks like classification. We expect this training method to generate deep features representing a larger distribution of input images and so increase the robustness of the defending auto-encoder(A2).



Figure 1: Concept Architecture

The goal of the attacker A1 is to find the minimal perturbation which fools the defender A2 with a loss value at least equal to a threshold  $\beta$ . The defender should avoid to be deceived.

$$\min_{\theta_{defense}} E_{x \sim \chi}(L_\alpha(x, Defense(Attack(x, \theta_{attack}), \theta_{defense}))) \quad (3)$$

$$\min_{\theta_{attack}} E_{x \sim \chi}(L_\alpha(Attack(x, \theta_{attack}), x)) \quad (4)$$

Under the constraint:

$$E_{x \sim \chi}(L_\alpha(x, Defense(Attack(x, \theta_{attack}), \theta_{defense}))) \geq \beta \quad (5)$$

Where:

$\chi$  is the ensemble of distribution examples

$L_\alpha$  is the loss function :  $\chi, \chi \rightarrow R$

$Attack(x, \theta_{attack}) : \chi \rightarrow \chi$  with parameters  $\theta_{attack}$

$Defense(x, \theta_{defense}) : \chi \rightarrow \chi$  with parameters  $\theta_{defense}$

$\theta_{attack}$  is the weight of the attacker.

$\theta_{defense}$  is the weight of the defender.

NB: The optimization problem under constraint is a not zero-sum game.

### 3.1 Motivation

Another representation of the problem which could seem more intuitive is to minimize/maximize, zero-sum game, the final error for the defender/attacker while the adversarial examples are not too distant to the original image (constraint).

$$\max_{\theta_{attack}} \min_{\theta_{defense}} E_{x \sim \chi}(L_\alpha(x, Defense(Attack(x, \theta_{attack}), \theta_{defense}))) \quad (6)$$

Under the constraint:

$$E_{x \sim \chi}(L_\alpha(Attack(x, \theta_{attack}), x)) \leq \beta \quad (7)$$

The first representation (Equations 3, 4, 5) is preferred to the second (Equations 6, 7) for two reasons:

- The attacks against neural networks try to minimize the perturbation of the input, while the second representation does not.
- The problem optimization is harder to solve because the perturbations generated by state of art attacks like DeepFool (11) are not on the constraint border (when  $\beta$  is large enough to accept all types of perturbations).



During the training, the optimization process will follow these strategies (See figure 2).

- In the first representation: at the beginning, the optimizer will reach the border of the constraint and then it will follow it until reaching the optimum.
- In the second representation, the optimizer will explore the space of solution while continuously closing the gap with the border, reaching it only at the end of the training.




Figure 2: Two representations of the optimization problem.

This is the second reason why the first representation is preferred to the second: the constraint border is easier to follow. To solve this under constraint optimization problem, we choose to add the constraint in the training loss:

$$\min_{\theta_{attack}} E_{x \sim \chi} (L_{\alpha}(Attack(x, \theta_{attack}), x)) + \gamma * \max(E_{x \sim \chi} (L_{\alpha}(x, Defense(Attack(x, \theta_{attack}), \theta_{defense}))) - \beta, 0.0) \quad (8)$$

Where  $\gamma$  is big enough to satisfy the constraint. This representation differs from the classic Lagrangian because we want the constraint gradient part to be free when the constraint is satisfied. This approximation of the under constraint optimization problem is justified by the fact that the inequality loss should be equal or close to  $\beta$ .

## 3.2 Model and training procedure

The attacker is implemented as an auto-encoder where a uniform random vector and the features of the encoder are concatenated and then passed through the generator. Adding a uniform random vector will allow the attacker to generate noise without modifying the "real" features. The defender is implemented as a classic auto-encoder. In a first phase, to initialize the weights of the networks, the two auto-encoders are trained to generate realistic images (the constraint is ignored, see Equation [3] and [4]).

In the second phase, the weights of the Attacker's encoder and the Defender's generator are fixed. This choice is justified by:

- Attacker's encoder: The goal of the attacker is to generate adversarial examples from given non-perturbed examples. Because image information is compressed in the features, the attacker could be summarized by a generator which has the features generated by the encoder as inputs, already learned, and a random noise.
- Defender's generator: The goal of the defender is to be robust against adversarial attacks. This means that the features of adversarial examples should be identical to the features of non-perturbed examples. This is done by fixing the generator parameters after the initialization phase, in this manner the features of perturbed input will be generated in the same way as normal inputs.

During this second phase we optimize the equations [3] and [8].

## 3.3 Testing procedure

As an auto-encoder is an architecture made for unsupervised learning, no classifier is used during training. In order to be able to use adversarial attacks and evaluate the network robustness a

classifier is trained with the defender’s features generated by the encoder with data coming from the dataset, not the adversarial data generated by the attacker. The robustness of the classifier is evaluated against the followings attacks: DeepFool  $L_2$  and  $L_\infty$ , Single Pixel, LocalSearch, FGSM and BIM. Robustness is evaluated by two criteria, the attack success rate and the noise level of adversarial examples. Noise level is defined by:

$$Noise_i = \sqrt{\sum_i (X_i - Adv_i)^2} \tag{9}$$

Where  $X_i$  is the original image and  $Adv_i$  is the generated adversarial example.

## 4 MNIST Experiments

### 4.1 Configuration

This method has been evaluated on the MNIST dataset with the following networks:

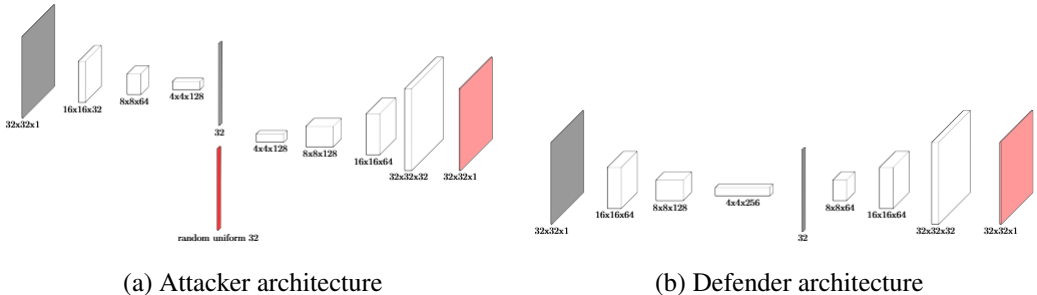


Figure 3: Network architectures

The neural network use convolutions with kernel size 3x3, strides 2x2 and ReLU activation function. The feature layers are not passed through an activation function. The last layers have  $(1 + \tanh)/2$  as activation function. The batch size is equal to 128.

The Adam optimizer is used with the following learning rates:

- Defender : Initialization phase =  $5 * 10^{-4}$ , Optimization phase =  $10^{-5}$
- Attacker: Initialization phase =  $10^{-3}$ , Optimization phase =  $10^{-5}$

The constants were fixed at  $\gamma = 5.0$ ,  $\beta = 0.01$ , the  $L_2$  loss is used. The initialization has been done during 9 epochs and the optimization phase during 31 epochs. Then a classifier is trained during epochs 20 with Adam (learning rate  $10^{-3}$ , loss: softmax-cross-entropy). The classifier uses the Defender’s encoder (whose weights are fixed) and a hidden dense layer of 64 units with ReLU activation. NB: The networks have not been optimized for this task and no search was done for the optimal parameters  $\gamma$  and  $\beta$ . However, several learning rates have been tested to achieve the presented results: a bad choice of learning rate could imply a non-robust network.

## 4.2 Results

**Adversarial attacks tools** Tools have been developed to facilitate adversarial attacks implementation such as Cleverhans (14) or Foolbox (16). Those tools are open-source Frameworks that propose state-of-art algorithms used to generate adversarial examples on any model.

In order to compare our method against a traditional network, we trained a classic autoencoder, with an architecture similar to Defender architecture in Figure 3b. We then train a classifier as previously described on each network. After 20 epochs, the classifier achieves an accuracy around 97% on MNIST test set. Then, the network with classifier is attacked by several methods on the whole test set, using Foolbox (16)

Attack Name	Classic Auto Encoder	Our Method
DeepFool (iter=50 - $L_2$ )	94,41 % / 1,25	<b>0.54 % / 0.0075</b>
DeepFool (iter=50 - $L_\infty$ )	100 % / 1,81	<b>4.13 % / 1.13</b>
FGSM ( $\epsilon=0.005$ )	0.41 % / 0.07	<b>1.61 % / 0.04</b>
Iterative ( $\epsilon=0.005$ )	1.04 % / 0.08	<b>1.25 % / 0.04</b>
Single Pixel	9.17 % / 1.0	<b>4.72 % / 1.0</b>
LocalSearch	44.41 % / 5.22	<b>27.51 % / 9.71</b>

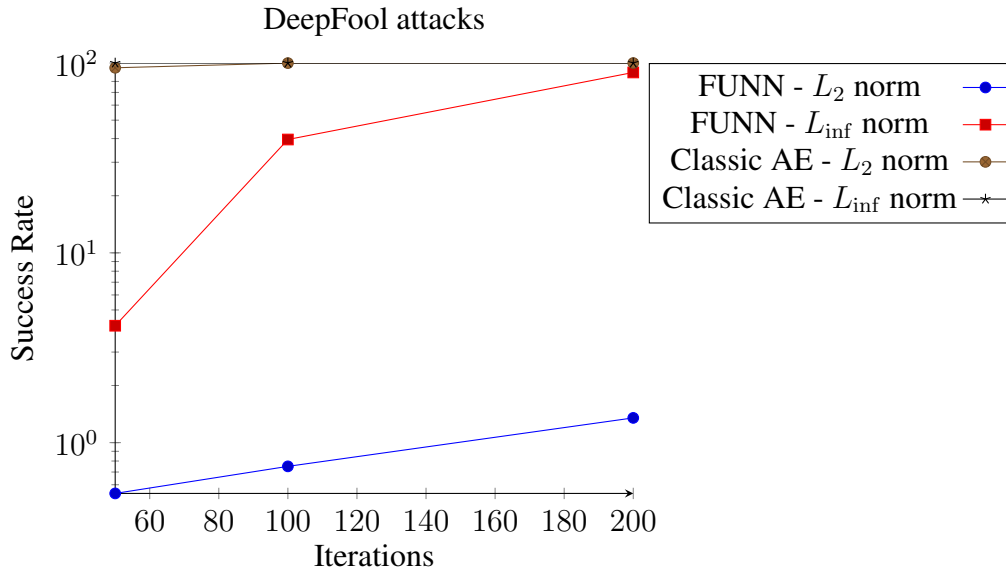
Table 1: Attack results : Attack Success Rate / Noise Level

\* The evaluation is done only on the 1000 first image of the dataset, due to computation time.

\*\*The noise level is the mean of the  $L_2$  norm between the original image and the adversarial image for all successful attacks.

Attack Name	MagNet	Our Method
DeepFool (iter=50 - $L_\infty$ )	<b>0.6 %</b>	4.13 %
FGSM ( $\epsilon=0.005$ )	3.2 %	<b>1.61 %</b>
Iterative ( $\epsilon=0.005$ )	4.8 %	<b>1.25 %</b>

Table 2: MagNet comparison



### 4.3 Analysis

**DeepFool with  $L_2$  norm** In the previous table we can see that only 0.5 % of attack succeeded when DeepFool was used with 50 iterations (number of iterations used by the original article of DeepFool) for a noise level mean around 0.0075. That was expected, as during learning, the attacker tries to find noise level greater than 0.01 (in fact, “equal to” like explained above) which can fool the defender. If the defender succeeded to defend against this level of noise, we expect to have as good a defense against any level of noise lower than this threshold. The classifier seems to be robust against  $L_2$  DeepFool attacks for 50 iterations. When the number of deepfool iterations increases, the success rate of the attack slowly increases: for 200 iterations, 1.35% of attack succeeded. It’s only when the number of iterations reaches 1000 steps that the percentage of successful attacks becomes very significant (15.25%). However, the noise increases significantly at the same time to reach 0.73 which is far from 0.01. We may think that increasing the threshold would allow to increase the robustness but in the case of MNIST it is impossible because doing so, the attacker converges to a black image. Indeed, because the  $L_2$  distance between a MNIST image and a black image is around 0.07, it is not possible to increase the threshold too much.

**DeepFool with  $L_\infty$  norm** With the default step used by DeepFool (50 steps), the classifier seems to be robust (4.13% of success rate) for the  $L_\infty$  norm. However, the classifier is not robust when the number of steps increases: at 200 steps, the success rate of the attacks reaches 89.20%. This is the reason why it is important to test DeepFool with several steps which is unfortunately not often done in the defense literature. However, even with 50 steps, the norm of the noise is strong with 1.13. This is the reason why it does not necessary mean that the network is not robust at all against this attack.

**Single Pixel and LocalSearch** Even if the number of success of the attack is relatively low (4.72%), this level of success is higher than what we may expect. In our case this attack looks more powerful than DeepFool which is designed precisely to minimize the  $L_2$  norm. We would expect a result lower than with DeepFool.

**Gradient Methods** FGSM and Gradient have a high attack success rate, close to 100% for Gradient Sign and Gradient. This result was expected because the defender is not trained with this high level of noise, shown by all these methods.

**MagNet Comparaison** Our method have similar performance to MagNet (9) without using any adversarial detectors that induces more computation time during inference.

## 5 Conclusion

Without adding extra computation complexity, like filtering the image or detecting adversarial images, we demonstrated that a classic network can be robust against  $L_2$  attacks by an adapted unsupervised training procedure. This procedure shows promising results toward robust unsupervised networks. However, even if it was expected due to the design of the defense, the trained network is not robust against all types of attacks such as gradient attacks. Further studies should be focused on a more general defense against different attacks with unsupervised training. We will also focus our future work on adapting this type of learning procedure to real world datasets with more complex network architectures.

## References and Notes

1. Shumeet Baluja and Ian Fischer. Learning to Attack: Adversarial Transformation Networks. page 9.
2. Nicholas Carlini and David Wagner. Defensive Distillation is Not Robust to Adversarial Examples. *arXiv:1607.04311 [cs]*, July 2016. arXiv: 1607.04311.
3. Joachim Folz, Sebastian Palacio, Joern Hees, Damian Borth, and Andreas Dengel. Adversarial Defense based on Structure-to-Signal Autoencoders. *arXiv:1803.07994 [cs, stat]*, March 2018. arXiv: 1803.07994.
4. Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and Harnessing Adversarial Examples. *arXiv:1412.6572 [cs, stat]*, December 2014. arXiv: 1412.6572.
5. Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the Knowledge in a Neural Network. *arXiv:1503.02531 [cs, stat]*, March 2015. arXiv: 1503.02531.
6. Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. *arXiv:1607.02533 [cs, stat]*, July 2016. arXiv: 1607.02533.
7. Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial Machine Learning at Scale. *arXiv:1611.01236 [cs, stat]*, November 2016. arXiv: 1611.01236.
8. Hyeungill Lee, Sungyeob Han, and Jungwoo Lee. Generative Adversarial Trainer: Defense to Adversarial Perturbations with GAN. *arXiv:1705.03387 [cs, stat]*, May 2017. arXiv: 1705.03387.
9. Dongyu Meng and Hao Chen. MagNet: A Two-Pronged Defense against Adversarial Examples. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and*



*Communications Security - CCS '17*, pages 135–147, Dallas, Texas, USA, 2017. ACM Press.

10. Jan Hendrik Metzen, Tim Genewein, Volker Fischer, and Bastian Bischoff. On Detecting Adversarial Perturbations. *arXiv:1702.04267 [cs, stat]*, February 2017. arXiv: 1702.04267.
11. Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. DeepFool: a simple and accurate method to fool deep neural networks. *arXiv:1511.04599 [cs]*, November 2015. arXiv: 1511.04599.
12. Vishaal Munusamy Kabilan, Brandon Morris, and Anh Nguyen. VectorDefense: Vectorization as a Defense to Adversarial Examples. April 2018.
13. Nina Narodytska and Shiva Prasad Kasiviswanathan. Simple Black-Box Adversarial Perturbations for Deep Networks. *arXiv:1612.06299 [cs, stat]*, December 2016. arXiv: 1612.06299.
14. Nicolas Papernot, Fartash Faghri, Nicholas Carlini, Ian Goodfellow, Reuben Feinman, Alexey Kurakin, Cihang Xie, Yash Sharma, Tom Brown, Aurko Roy, Alexander Matyasko, Vahid Behzadan, Karen Hambardzumyan, Zhishuai Zhang, Yi-Lin Juang, Zhi Li, Ryan Sheatsley, Abhibhav Garg, Jonathan Uesato, Willi Gierke, Yinpeng Dong, David Berthelot, Paul Hendricks, Jonas Rauber, and Rujun Long. Technical report on the cleverhans v2.1.0 adversarial examples library. *arXiv preprint arXiv:1610.00768*, 2018.
15. Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. Distillation as a Defense to Adversarial Perturbations against Deep Neural Networks. *arXiv:1511.04508 [cs, stat]*, November 2015. arXiv: 1511.04508.

16. Jonas Rauber, Wieland Brendel, and Matthias Bethge. Foolbox: A Python toolbox to benchmark the robustness of machine learning models. *arXiv:1707.04131 [cs, stat]*, July 2017. arXiv: 1707.04131.
17. Pouya Samangouei, Maya Kabkab, and Rama Chellappa. DEFENSE-GAN: PROTECTING CLASSIFIERS AGAINST ADVERSARIAL ATTACKS USING GENERATIVE MODELS. page 17, 2018.
18. Gokula Krishnan Santhanam and Paulina Grnarova. Defending Against Adversarial Attacks by Leveraging an Entire GAN. *arXiv:1805.10652 [cs, stat]*, May 2018. arXiv: 1805.10652.
19. Shiwei Shen, Guoqing Jin, Ke Gao, and Yongdong Zhang. APE-GAN: Adversarial Perturbation Elimination with GAN. *arXiv:1707.05474 [cs]*, July 2017. arXiv: 1707.05474.
20. Vignesh Srinivasan, Arturo Marban, Klaus-Robert Müller, Wojciech Samek, and Shinichi Nakajima. Counterstrike: Defending Deep Learning Architectures Against Adversarial Samples by Langevin Dynamics with Supervised Denoising Autoencoder. *arXiv:1805.12017 [cs, stat]*, May 2018. arXiv: 1805.12017.
21. Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv:1312.6199 [cs]*, December 2013. arXiv: 1312.6199.
22. Xiaoyong Yuan, Pan He, Qile Zhu, and Xiaolin Li. Adversarial Examples: Attacks and Defenses for Deep Learning. *arXiv:1712.07107 [cs, stat]*, December 2017. arXiv: 1712.07107.