

# Towards adversarial robustness with 01 loss neural networks

1<sup>st</sup> Yunzhe Xue

Department of Computer Science  
New Jersey Institute of Technology  
Newark, USA  
yx277@njit.edu

2<sup>nd</sup> Meiyan Xie

Department of Computer Science  
New Jersey Institute of Technology  
Newark, USA  
mx42@njit.edu

3<sup>rd</sup> Usman Roshan

Department of Computer Science  
New Jersey Institute of Technology  
Newark, USA  
usman@njit.edu

**Abstract**—Motivated by the general robustness properties of the 01 loss we propose a single hidden layer 01 loss neural network trained with stochastic coordinate descent as a defense against adversarial attacks in machine learning. One measure of a model’s robustness is the minimum distortion required to make the input adversarial. This can be approximated with the Boundary Attack (Brendel *et. al.* 2018) and HopSkipJump (Chen *et. al.* 2019) methods. We compare the minimum distortion of the 01 loss network to the binarized neural network and the standard sigmoid activation network with cross-entropy loss all trained with and without Gaussian noise on the CIFAR10 benchmark binary classification between classes 0 and 1. Both with and without noise training we find our 01 loss network to have the largest adversarial distortion of the three models by non-trivial margins. To further validate these results we subject all models to substitute model black box attacks under different distortion thresholds and find that the 01 loss network is the hardest to attack across all distortions. At a distortion of 0.125 both sigmoid activated cross-entropy loss and binarized networks have almost 0% accuracy on adversarial examples whereas the 01 loss network is at 40%. Even though both 01 loss and the binarized network use sign activations their training algorithms are different which in turn give different solutions for robustness. Finally we compare our network to simple convolutional models under substitute model black box attacks and find their accuracies to be comparable. Our work shows that the 01 loss network has the potential to defend against black box adversarial attacks better than convex loss and binarized networks.

**Index Terms**—adversarial attacks, adversarial distortion, 01 loss, stochastic coordinate descent, convolutional neural networks, deep learning

## I. INTRODUCTION

State of the art machine learning algorithms can achieve high accuracies in classification tasks but misclassify minor perturbations in the data known as adversarial attacks [1]–[5]. Adversarial examples have been shown to transfer across models which makes it possible to perform transfer-based (substitute model) black box attacks [6]. To counter adversarial attacks many defense methods been proposed with adversarial training being the most popular [7]. This is known to improve robustness to adversarial examples but also tends to lower accuracy on clean test data that has no perturbations [8], [9]. Many previously proposed defenses have also shown to be vulnerable [4], [10], [11] thus leaving adversarial robustness an open problem in machine learning.

The 01 loss is known to be more robust to outliers than convex loss models [12]–[14]. In addition to being robust to outliers the 01 loss is also robust to noise in the training data [15], [16]. Under this loss minimizing the empirical risk amounts to minimizing the empirical adversarial risk [17], [18] with certain assumptions of noise. Convex losses also fail to minimize the adversarial 01 loss on linear models [19].

Motivated by the above robustness properties of 01 loss we propose a 01 loss dual layer neural network as a defense against adversarial attacks. Computationally 01 loss presents a considerable challenge because it is NP-hard to solve [20]. Previous attempts [13], [21]–[24] lack on-par test accuracy with convex solvers and are slow and impractical for large image benchmarks. However, a recent stochastic coordinate descent method for linear 01 loss models [12] has shown to attain comparable accuracies to state of the art linear solvers like the support vector machine. Thus we extend the coordinate descent as an optimizer to train our network along the lines of related work [25].

We compare the adversarial robustness of our model to an equivalent one that uses sigmoid activation and cross-entropy loss. This is the standard activation and loss that are widely used in neural networks today. We also compare our model to the binarized neural network [26]–[28], that also uses sign activations like our model but it has two differences. First its weights are also constrained to be binary +1 and -1, or 1 and 0 [26]–[28]. Second, it is trained with gradient descent by approximating the sign activation whereas we take a direct coordinate descent approach.

Measuring adversarial robustness is not trivial and several best practices have been recommended [29]. We incorporate several of them in our study. In particular we study (1) the robustness of models to random Gaussian noise, (2) the minimum distortion required to make a datapoint adversarial, and (3) the accuracy of substitute model black box attacks. We focus mainly on binary classification between classes 0 and 1 on the CIFAR10 image benchmark [30] where we make the following findings.

- All models are more robust to Gaussian noise than adversarial attacks, but our 01 loss network augmented with Gaussian noise during training has a higher accuracy on large distortions

- The minimum distortion to make an image adversarial is higher for our 01 loss network compared to the standard sigmoid activated cross-entropy loss network and binarized networks
- Substitute model black box attacks are far less effective on our 01 loss network compared to the standard sigmoid activated cross-entropy loss network and binarized networks
- Compared to simple convolutional neural networks like LeNet [31] our model (without convolutions) has higher accuracies on adversarial examples from substitute model black box attacks when the distortions are high.

## II. METHODS

### A. Background

The problem of determining the hyperplane with minimum number of misclassifications in a binary classification problem is known to be NP-hard [20]. In mainstream machine learning literature this is called minimizing the 01 loss [32] given in Objective 1,

$$\frac{1}{2n} \arg \min_{w, w_0} \sum_i (1 - \text{sign}(y_i(w^T x_i + w_0))) \quad (1)$$

where  $w \in R^d$ ,  $w_0 \in R$  is our hyperplane, and  $x_i \in R^d$ ,  $y_i \in \{+1, -1\}$ .  $\forall i = 0 \dots n - 1$  are our training data. Popular linear classifiers such as the linear support vector machine, perceptron, and logistic regression [33] can be considered as convex approximations to this problem that yield fast gradient descent solutions [14]. However, they are also more sensitive to outliers than the 01 loss [12]–[14] and more prone to mislabeled data than 01 loss [15]–[17].

### B. A dual layer 01 loss neural network

We extend the 01 loss to a simple two layer neural network with  $k$  hidden nodes and sign activation that we call the MLP01 loss. This objective for binary classification can be given as

$$\frac{1}{2n} \arg \min_{W, W_0, w, w_0} \sum_i (1 - \text{sign}(y_i(w^T (\text{sign}(W^T x_i + W_0)) + w_0))) \quad (2)$$

where  $W \in R^{d \times k}$ ,  $W_0 \in R^k$  are the hidden layer parameters,  $w \in R^k$ ,  $w_0 \in R$  are the final layer node parameters,  $x_i \in R^d$ ,  $y_i \in \{+1, -1\}$ .  $\forall i = 0 \dots n - 1$  are our training data, and  $\text{sign}(v \in R^k) = (\text{sign}(v_0), \text{sign}(v_1), \dots, \text{sign}(v_{k-1}))$ . While this is a straightforward model to define optimizing it is a different story altogether. Optimizing even a single node is NP-hard which makes optimizing this network much harder.

### C. Stochastic coordinate descent for 01 loss

In Algorithm 1 we sketch our coordinate descent for our 01 loss network that is based upon earlier work [12]. We initialize all parameters to random values from the Normal distribution with mean 0 and variance 1. We then randomly select a subset of the training data (known as a batch) and

---

### Algorithm 1 Stochastic coordinate descent for two layer 01 loss network

---

#### Procedure:

1. Initialize all network weights  $W$ ,  $w$  to random values from the Normal distribution  $N(0, 1)$ .
2. Set network thresholds  $W_0$  to the median projection value on their corresponding weight vectors and  $w_0$  to the projection value that minimizes our network objective.

#### while $i < \text{epochs}$ do

1. Randomly sample a batch of data equally from each class
2. Perform coordinate descent separately first on the final node  $w$  and then a randomly selected hidden node  $u$  (a random column from the hidden layer weight matrix  $W$ )
3. In the coordinate descent we randomly pick a set of features and perform a single update to each one. For each update we determine the threshold that optimizes the 01 loss on the sampled data: we sort the projections  $w^T x_i$  and pick the optimal middle value between each consecutive pair in the projected values.
4. After making the best update we evaluate the 01 loss on the full dataset and accept the change if it improves the loss.

#### end while

---

perform the coordinate descent analog of a single step gradient update in stochastic gradient descent [34].

When the gradient is known we step in its negative direction by a factor of the learning rate:  $w = w - \eta \nabla(f)$  where  $f$  is the objective. In our case since the gradient does not exist we randomly select  $k$  features (set to 128 in our experiments), modify the corresponding entries in  $w$  by the learning rate (set to 0.17) one at a time, and accept the modification that gives the largest decrease in the objective. Key to our search is a heuristic to determine the optimal threshold each time we modify an entry of  $w$ . In this heuristic we perform a linear search on a subset of the projection  $w^T x_i$  and select  $w_0$  that minimizes the objective.

In Figure 1 we show the effect of the batch size (as a percentage of each class to ensure fair sampling) on a linear 01 loss search on CIFAR10 between classes 0 and 1. We see that a batch size of 75% reaches a train accuracy of 80% faster than the other batch sizes. Thus we use this batch size in all our experiments going forward.

We also see that for this batch size the search flattens after 15 iterations (or epochs as given in the figure). We run 1000 iterations to ensure a deep search with an intent to maximize test accuracy. The problem with our search described above is that it will return different solutions depending upon the initial starting point. To make it more stable we run it a 100 times from different random seeds and use the majority vote for prediction. Full details of our training algorithms are in the Supplementary Material.




Fig. 1. Train and test accuracy of our stochastic coordinate descent on CIFAR10 class 0 vs 1 with different batch sizes (denoted as nrows).

#### D. Implementation, experimental platform, and image data

1) *Implementation*: We implement our 01 loss network (MLP01) in Python and Pytorch [35], the sigmoid activated cross-entropy loss network (MLP) in scikit-learn [36], and binarized neural network (BNN) with the Larq software suite <https://github.com/larq/larq>. We train MLP with stochastic gradient descent that has a batch size of 200, momentum of 0.9, and learning rate of 0.01. For BNN we use the approximate sign activation [27] that has been shown to give higher test accuracies than other variants and the original straight through estimator [37].

2) *Computational platform*: We ran all experiments on Intel Xeon 6142 2.6GHz CPUs and NVIDIA Titan RTX GPU machines (for parallelizing multiple votes). Our MLP01 source code, supplementary programs, data, and Supplementary Material file are available from <https://github.com/zero-one-loss/mlp01>.

3) *Data*: We experiment on the popular image benchmark CIFAR10 [30] that has  $32 \times 32$  color images with 50000 training and 10000 test. We extract data from classes 0 and 1 and experiment on binary classification between them. This gives us a total of 10000 training and 1000 test examples. We normalize each image by dividing each pixel value by 255.

### III. RESULTS

We refer to our 01 loss neural network as MLP01, the sigmoid activated cross-entropy loss network as MLP, and the binarized network as BNN. We use one hidden layer of 20 nodes in all networks. For each model we run it a 100 times with different random number generator seeds and return the majority vote as the prediction.

In addition to training each model on the training data, we also study three versions trained with augmented Gaussian noise. In the augmentation we take each datapoint  $x$  from the training set and add Gaussian noise to it:  $x' = x + N(0, \sigma)$  where  $N(0, \sigma)$  is a vector of the same dimension as  $x$  and each entry is selected from the Normal distribution with mean 0 and standard deviation  $\sigma$ .

We refer to the accuracy on the test data as clean data test accuracy. An incorrectly classified adversarial example is considered a successful attack whereas a correctly classified adversarial is a failed one. Thus when we refer to accuracy of adversarial examples it is the same as  $100 - \text{attacksuccessrate}$ . The lower the accuracy the more effective the attack.

#### A. Sensitivity to Gaussian noise

We start with accuracy of models trained without and with noise on CIFAR10 class 0 vs. 1. In addition to evaluating the accuracy of each model on clean test data, we add noise to each test datapoint as  $x' = x + N(0, \sigma)$  where  $N(0, \sigma)$  is a vector of the same dimension as  $x$  and each entry is selected from the Normal distribution with mean 0 and standard deviation  $\sigma$ . We consider  $\sigma$  ranging from .004 to 1. The lower bound is the minimum distance of  $\frac{1}{255}$  between two pixels and 1 is the maximum distortion. If a pixel is negative or above 1 after adding noise we clip it to 0 and 1 respectively.

In Figure 2 we see that noise does not affect the accuracy of models trained without and with noise upto distortion threshold of 0.125. After that all models begin to dip in accuracy with MLP01 model trained without noise showing the steepest descent. At the same time MLP01 trained with noisy augmentation of  $\sigma = .2$  (denoted as mlp01\_ep2 in Figure 2) is also most robust to high levels of noise.




Fig. 2. Accuracy of test data without and with random noise of different random Gaussian distortions. In addition to models trained on clean training data we consider two versions trained with augmented Gaussian noise with distortion thresholds of 0.1 and 0.2 (denoted with ep1 and ep2 respectively).

#### B. Minimum adversarial distortion

Determining the minimum distortion to an image such that it will fool a classifier is itself an NP-hard problem for ReLU activated neural networks [38], [39] and tree ensemble classifiers [40]. Even approximating the minimum distortion for ReLU activated neural networks is NP-hard [41]. Recent heuristics such as Boundary attack [5] and HopSkipJump attack [42] attempt to find an adversarial example with the minimum distance to the clean correctly classified version.

We use both methods to evaluate the minimum distortion of all three models including their noise trained versions. Both

methods can take long to finish with default parameters even for a single example. For example for a single image BNN takes 3 days on an exclusive CPU core. Thus we use all default parameters except for maxiter that we set to 100 so that the program finishes within our computing limitations. To confirm that this does not severely affect the relative distortions we ran both Boundary and HopSkipJump attacks with maxiter set to 10, 100, and 500 (which is the default) on a single image. We found the relative distortions between MLP and MLP01 to be the same across the three values.

Since both attack methods start with a random initialization we run each of them 10 times for a single example on each model and report the minimum value found. In Table I we report these values for a single random test datapoint from CIFAR10 classes 0 and 1 that is correctly classified by all models. We see that the minimum distortion of MLP01 is much higher than both MLP and BNN by both attack methods and under both  $L_2$  and  $L_\infty$  norms. We also see that HopSkipJump attack is more effective than Boundary attack and finds a smaller distortion.

TABLE I  
MINIMUM ADVERSARIAL DISTORTION OF A SINGLE RANDOM TEST IMAGE

	$L_2$ distance			$L_\infty$ distance		
	BNN	MLP	MLP01	BNN	MLP	MLP01
Boundary	2.72	1.33	10.58	0.17	0.08	0.58
HopSkipJump	0.82	0.44	2.21	0.04	0.03	0.16

In Table II we report the HopSkipJump distortions for four more randomly selected images from CIFAR10 classes 0 and 1 that are correctly classified. For the first image both have comparable distortion but for the other three MLP01 is higher.

TABLE II  
MINIMUM ADVERSARIAL DISTORTION GIVEN BY HOPSKIPJUMP OF FOUR RANDOM CORRECTLY CLASSIFIED TEST IMAGES

	$L_2$ distance		$L_\infty$ distance	
	MLP	MLP01	MLP	MLP01
Image1	.64	.52	.043	.041
Image2	.75	2.42	.06	.15
Image3	.88	1.12	.06	.09
Image4	1.15	4.86	.09	.28
Average	.86	2.23	.063	0.14

Training our models with noise has an interesting effect on the distortions. In Table III we see the minimum distortions of models trained with Gaussian noise with distortions of 0.1 and 0.2 (as described earlier). We report the distortions for the same image as in Table I. As we increase the noise threshold the MLP01 model’s minimum distortion also rises whereas the other two models are stable or fluctuate.

We make similar observations between MLP and MLP01 on four random examples as shown in Table IV. This suggests that perhaps training MLP01 with augmented noise examples increases their robustness.

### C. Substitute model black box attacks

As further verification of the above distortions we perform substitute model black box attacks on all three models. In

TABLE III  
MINIMUM ADVERSARIAL DISTORTION OF THE SINGLE TEST IMAGE FROM TABLE I AS GIVEN BY HOPSKIPJUMP. MODELS ARE TRAINED WITH GAUSSIAN NOISE OF INCREASING DISTORTION SHOWN BY  $\epsilon$ .

	$L_2$ distance			$L_\infty$ distance		
	$\epsilon = .004$	$\epsilon = .1$	$\epsilon = .2$	$\epsilon = .004$	$\epsilon = .1$	$\epsilon = .2$
BNN	.51	.47	.53	.025	.021	.025
MLP	.39	.35	.36	.023	.021	.022
MLP01	2.52	2.61	3.32	.196	.215	.22

TABLE IV  
MINIMUM ADVERSARIAL DISTORTION OF FOUR RANDOM CORRECTLY CLASSIFIED TEST IMAGES FROM TABLE II AS GIVEN BY HOPSKIPJUMP. MODELS ARE TRAINED WITH GAUSSIAN NOISE OF INCREASING DISTORTION SHOWN BY  $\epsilon$ .

	$L_2$ distance					
	MLP			MLP01		
	$\epsilon = .004$	$\epsilon = .1$	$\epsilon = .2$	$\epsilon = .004$	$\epsilon = .1$	$\epsilon = .2$
Image1	.58	.47	.47	.69	.61	.47
Image2	.6	.52	.54	2.65	2.89	2.51
Image3	.72	.67	.66	1.92	2.72	4.03
Image4	.97	.85	.84	4.73	4.08	5.17
Average	.72	.63	.63	2.5	2.58	3.05

	$L_\infty$ distance					
	MLP			MLP01		
	$\epsilon = .004$	$\epsilon = .1$	$\epsilon = .2$	$\epsilon = .004$	$\epsilon = .1$	$\epsilon = .2$
Image1	.04	.03	.03	.05	.05	.03
Image2	.04	.03	.03	.18	.2	.18
Image3	.05	.04	.04	.14	.21	.35
Image4	.07	.05	.05	.28	.29	.35
Average	.05	.04	.04	.16	.19	.23

this method we try to approximate the target model with a substitute and then generate white box adversaries from the substitute to attack the target model. The success of this method relies upon transferability of adversarial examples between models. We use the standard adversarially augmented training algorithm of Papernot et. al. [43] to train the substitute. In the Supplementary Material we provide full details of the algorithm.

This in fact is a powerful attack method that needs only predicted labels from the target (like Boundary and HopSkipJump) but requires much fewer queries. Once the substitute is trained it can produce adversaries for any input. Recent advances in transferability have made this method more effective and broken defenses based on adversarial training [44], [45]. For the substitute model we use a three layer sigmoid activated cross-entropy loss network with 200 nodes in each hidden layer. We start with 200 random test data points from which we iteratively train the substitute model with augmented adversaries.

In Figure 3 we see the accuracy of adversarial examples at the end of the 20<sup>th</sup> epoch. We also show the accuracy of the three models on random Gaussian noise of the same distortions (from our earlier subsection above). Clearly the black box adversaries are far more effective than random noise indicating that the substitute model training was successful. In agreement with our distortions from Boundary and HopSkipJump above we see that MLP01 can correctly classify images of much higher distortion than BNN and MLP.

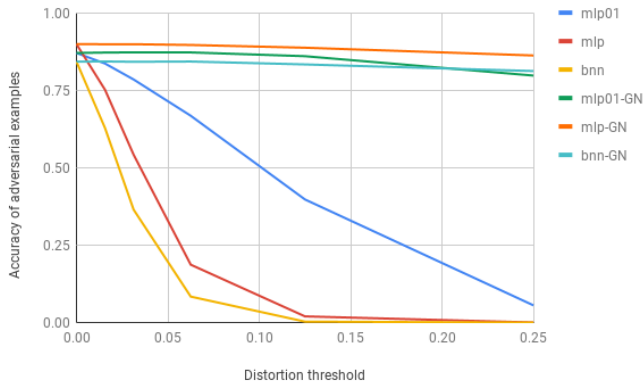


Fig. 3. Accuracy of adversarial examples and test data with Gaussian noise (denoted as -GN for each model) for various distortion thresholds. The adversarial examples are far more effective than random noise. At distortion 0.125 both MLP and BNN have near 0% accuracy whereas MLP01 has 40%.

#### D. Comparison to convolutional neural networks

As a test against state of the art classification methods we compare our 01 loss network with 500 hidden nodes to two convolutional neural networks. First is LeNet [31] which is among the first convolutional networks to be proposed and second is SimpleNet500. In this model we use the same convolutional layers as LeNet followed by one layer of 500 nodes and then the final output node.

We employ the same substitute model training algorithm as in the above subsection. However instead of a dual hidden layer model we use a convolutional network as the substitute. In each convolutional block we have a  $3 \times 3$  convolutional kernel followed by max pool and batch normalization. In the first, second, third, and fourth layer we have 32, 64, 128, and 256 kernels respectively following by a final layer for the output.

In Figure 4 we see that our model has a comparable accuracy to the convolutional models on clean test data and low distortion thresholds. However, when we cross 0.03125 then MLP01 has the highest accuracy. At threshold 0.125 it is about 11% higher than both LeNet and SimpleNet500.

#### IV. DISCUSSION

Binarized neural networks [26]–[28] have weights and activations constrained to be near +1 and -1 (or 1 and 0) whereas our model weights are real numbers. The purpose of those networks is efficiency as opposed to robustness. Indeed we see in recent work that binarized networks offer marginal improvements in robustness to substitute model black box robustness on MNIST and none in CIFAR10 (see Tables 4 and 5 in [26] and Table 8 in [46]).

We make the same observations here: BNN has similar distortions to MLP and similar accuracies on adversarial examples. Both BNN and MLP01 have sign activations yet MLP01 has higher distortions and higher adversarial accuracies. Perhaps this has to do with the optimization method.

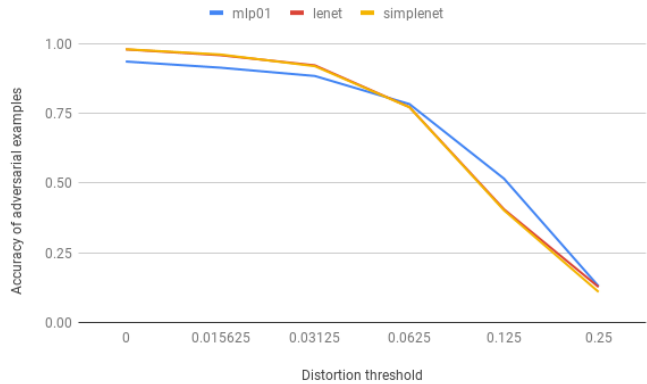


Fig. 4. Accuracy of adversarial examples for different distortion thresholds

BNNs are trained with an approximation to the sign activation that is differentiable whereas we train with direct coordinate descent.

In separate work we study transferability between our 01 loss network and the standard sigmoid activated cross-entropy loss networks [25]. There we show a lack of transferability between convex and 01 loss models in white box attacks and that both convex and 01 loss substitute model black box attacks are ineffective on our 01 loss network. However, in the work here we focus on the distortion thresholds of adversarial and Gaussian noise examples. In our related work [25] we also show robustness of our model to substitute model black box attacks in the multiclass setting, thus suggesting our work here is likely to extend to multiclass problems.

Interestingly the adversarial accuracy of our network is on-par with simple convolutional models that have the powerful advantage of convolutions. As future work 01 loss convolutions may be a promising avenue to obtain models with high clean test accuracy and high adversarial accuracy as well.

#### V. CONCLUSION

We show that our 01 loss neural network can correctly classify images with a higher distortion than both the sigmoid activated cross-entropy loss network and binarized neural networks.

#### VI. ACKNOWLEDGEMENTS

We thank Angus Galloway and Alexandros Dimakis for helpful feedback and discussions.

#### REFERENCES

- [1] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [2] Nicolas Papernot, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Z Berkay Celik, and Ananthram Swami. The limitations of deep learning in adversarial settings. In *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 372–387. IEEE, 2016.
- [3] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, 2016.
- [4] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017.

- [5] Wieland Brendel, Jonas Rauber, and Matthias Bethge. Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. *arXiv preprint arXiv:1712.04248*, 2017.
- [6] Nicolas Papernot, Patrick McDaniel, and Ian Goodfellow. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *arXiv preprint arXiv:1605.07277*, 2016.
- [7] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [8] Aditi Raghunathan, Sang Michael Xie, Fanny Yang, John C Duchi, and Percy Liang. Adversarial training can hurt generalization. *arXiv preprint arXiv:1906.06032*, 2019.
- [9] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P Xing, Laurent El Ghauoui, and Michael I Jordan. Theoretically principled trade-off between robustness and accuracy. *arXiv preprint arXiv:1901.08573*, 2019.
- [10] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. *arXiv preprint arXiv:1802.00420*, 2018.
- [11] Amin Ghiasi, Ali Shafahi, and Tom Goldstein. Breaking certified defenses: Semantic adversarial examples with spoofed robustness certificates. *arXiv preprint arXiv:2003.08937*, 2020.
- [12] Meiyang Xie, Yunzhe Xue, and Usman Roshan. Stochastic coordinate descent for 0/1 loss and its sensitivity to adversarial attacks. In *Proceedings of 18th IEEE International Conference on Machine Learning and Applications - ICMLA 2019*, page to appear, 2019.
- [13] Tan Nguyen and Scott Sanner. Algorithms for direct 0–1 loss optimization in binary classification. In *Proceedings of The 30th International Conference on Machine Learning*, pages 1085–1093, 2013.
- [14] Peter L. Bartlett, Michael I. Jordan, and Jon D. McAuliffe. Large margin classifiers: Convex loss, low noise, and convergence rates. In S. Thrun, L.K. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*, pages 1173–1180. MIT Press, 2004.
- [15] Naresh Manwani and PS Sastry. Noise tolerance under risk minimization. *IEEE transactions on cybernetics*, 43(3):1146–1151, 2013.
- [16] Aritra Ghosh, Naresh Manwani, and PS Sastry. Making risk minimization tolerant to label noise. *Neurocomputing*, 160:93–107, 2015.
- [17] Yueming Lyu and Ivor W Tsang. Curriculum loss: Robust learning and generalization against label corruption. *arXiv preprint arXiv:1905.10045*, 2019.
- [18] Weihua Hu, Gang Niu, Issei Sato, and Masashi Sugiyama. Does distributionally robust supervised learning give robust classifiers? *arXiv preprint arXiv:1611.02041*, 2016.
- [19] Han Bao, Clayton Scott, and Masashi Sugiyama. Calibrated surrogate losses for adversarially robust classification. *arXiv preprint arXiv:2005.13748*, 2020.
- [20] Shai Ben-David, Nadav Eiron, and Philip M Long. On the difficulty of approximately maximizing agreements. *Journal of Computer and System Sciences*, 66(3):496–514, 2003.
- [21] Shaodan Zhai, Tian Xia, Ming Tan, and Shaojun Wang. Direct 0-1 loss minimization and margin maximization with boosting. In C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 872–880. Curran Associates, Inc., 2013.
- [22] Yufang Tang, Xueming Li, Yan Xu, Shuchang Liu, and Shuxin Ouyang. A mixed integer programming approach to maximum margin 0–1 loss classification. In *2014 International Radar Conference*, pages 1–6. IEEE, 2014.
- [23] Shai Shalev-Shwartz, Ohad Shamir, and Karthik Sridharan. Learning linear and kernel predictors with the 0-1 loss function, 2011.
- [24] Ling Li and Hsuan-Tien Lin. Optimizing 0/1 loss for perceptrons by random coordinate descent. In *Neural Networks, 2007. IJCNN 2007. International Joint Conference on*, pages 749–754. IEEE, 2007.
- [25] Yunzhe Xue, Meiyang Xie, and Usman Roshan. On the transferability of adversarial examples between convex and 01 loss models. In *Proceedings of the IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2020.
- [26] Angus Galloway, Graham W Taylor, and Medhat Moussa. Attacking binarized neural networks. *arXiv preprint arXiv:1711.00449*, 2017.
- [27] Matthieu Courbariaux, Itay Hubara, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1. *arXiv preprint arXiv:1602.02830*, 2016.
- [28] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European conference on computer vision*, pages 525–542. Springer, 2016.
- [29] Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian Goodfellow, Aleksander Madry, and Alexey Kurakin. On evaluating adversarial robustness. *arXiv preprint arXiv:1902.06705*, 2019.
- [30] Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009.
- [31] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [32] Shalev-Shwartz Shai, Ohad Shamir, and Karthik Sridharan. Learning linear and kernel predictors with the 0-1 loss function. *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, 22(3), 2011.
- [33] Ethem Alpaydin. *Machine Learning*. MIT Press, 2004.
- [34] Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pages 177–186. Springer, 2010.
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [36] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [37] Zechun Liu, Baoyuan Wu, Wenhan Luo, Xin Yang, Wei Liu, and Kwang-Ting Cheng. Bi-real net: Enhancing the performance of 1-bit cnns with improved representational capability and advanced training algorithm. In *Proceedings of the European conference on computer vision (ECCV)*, pages 722–737, 2018.
- [38] Guy Katz, Clark Barrett, David L Dill, Kyle Julian, and Mykel J Kochenderfer. Reluplex: An efficient smt solver for verifying deep neural networks. In *International Conference on Computer Aided Verification*, pages 97–117. Springer, 2017.
- [39] Aman Sinha, Hongseok Namkoong, and John Duchi. Certifiable distributional robustness with principled adversarial training. *arXiv preprint arXiv:1710.10571*, 2, 2017.
- [40] Alex Kanchelian, J Doug Tygar, and Anthony Joseph. Evasion and hardening of tree ensemble classifiers. In *International Conference on Machine Learning*, pages 2387–2396, 2016.
- [41] Tsui-Wei Weng, Huan Zhang, Hongge Chen, Zhao Song, Cho-Jui Hsieh, Duane Boning, Inderjit S Dhillon, and Luca Daniel. Towards fast computation of certified robustness for relu networks. *arXiv preprint arXiv:1804.09699*, 2018.
- [42] Jianbo Chen, Michael I Jordan, and Martin J Wainwright. Hop-skipjumpattack: A query-efficient decision-based attack. *arXiv preprint arXiv:1904.02144*, 3, 2019.
- [43] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, pages 506–519. ACM, 2017.
- [44] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. Ensemble adversarial training: Attacks and defenses. *arXiv preprint arXiv:1705.07204*, 2017.
- [45] Dongxian Wu, Yisen Wang, Shu-Tao Xia, James Bailey, and Xingjun Ma. Skip connections matter: On the transferability of adversarial examples generated with resnets. *arXiv preprint arXiv:2002.05990*, 2020.
- [46] Priyadarshini Panda, Indranil Chakraborty, and Kaushik Roy. Discretization based solutions for secure machine learning against adversarial attacks. *IEEE Access*, 7:70157–70168, 2019.