An ADMM-Based Universal Framework for Adversarial Attacks on Deep Neural Networks

Pu Zhao

Department of ECE, Northeastern University
zhao.pu@husky.neu.edu

Yanzhi Wang

Department of ECE, Northeastern University
yanzhi.wang@northeastern.edu

Sijia Liu

IBM Research AI
sijia.liu@ibm.com

Xue Lin

Department of ECE, Northeastern University
xue.lin@northeastern.edu

ABSTRACT

Deep neural networks (DNNs) are known vulnerable to adversarial attacks. That is, adversarial examples, obtained by adding delicately crafted distortions onto original legal inputs, can mislead a DNN to classify them as any target labels. In a successful adversarial attack, the targeted mis-classification should be achieved with the minimal distortion added. In the literature, the added distortions are usually measured by L_0 , L_1 , L_2 , and L_∞ norms, namely, L_0 , L_1 , L_2 , and L_∞ attacks, respectively. However, there lacks a versatile framework for all types of adversarial attacks.

This work for the first time unifies the methods of generating adversarial examples by leveraging ADMM (Alternating Direction Method of Multipliers), an operator splitting optimization approach, such that L_0 , L_1 , L_2 , and L_∞ attacks can be effectively implemented by this general framework with little modifications. Comparing with the state-of-the-art attacks in each category, our ADMM-based attacks are so far the strongest, achieving both the 100% attack success rate and the minimal distortion.

CCS CONCEPTS

• **Theory of computation** → **Mathematical optimization**; • **Computing methodologies** → **Computer vision problems**; **Neural networks**; • **Security and privacy** → *Software and application security*;

KEYWORDS

Deep Neural Networks; Adversarial Attacks; ADMM (Alternating Direction Method of Multipliers)

ACM Reference Format:

Pu Zhao, Sijia Liu, Yanzhi Wang, and Xue Lin. 2018. An ADMM-Based Universal Framework for Adversarial Attacks on Deep Neural Networks. In 2018 ACM Multimedia Conference (MM'18), October 22-26, 2018, Seoul, Republic of Korea. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3240508.3240639>

1 INTRODUCTION

Deep learning has been demonstrating exceptional performance on several categories of machine learning problems and has been

applied in many settings [7, 13, 14, 18, 21, 27, 31]. However, people recently find that deep neural networks (DNNs) could be vulnerable to adversarial attacks [4, 19, 22], which arouses concerns of applying deep learning in security-critical tasks. Adversarial attacks are implemented through generating adversarial examples, which are crafted by adding delicate distortions onto legal inputs. Fig. 1 shows adversarial examples for targeted adversarial attacks that can fool DNNs.

The security properties of deep learning have been investigated from two aspects: (i) enhancing the robustness of DNNs under adversarial attacks and (ii) crafting adversarial examples to test the vulnerability of DNNs. For the former aspect, research works have been conducted by either filtering out added distortions [2, 9, 12, 34] or revising DNN models [8, 10, 25] to defend against adversarial attacks. For the later aspect, adversarial examples have been generated heuristically [11, 28], iteratively [15, 19, 24, 33], or by solving optimization problems [1, 5, 6, 30]. These two aspects mutually benefit each other towards hardening DNNs under adversarial attacks. And our work deals with the problem from the later aspect.

For targeted adversarial attacks, the crafted adversarial examples should be able to mislead the DNN to classify them as any target labels, as done in Fig. 1. Also, in a successful adversarial attack, the targeted mis-classification should be achieved with the minimal distortion added to the original legal input. Here comes the question of how to measure the added distortions. Currently, in the literature, L_0 , L_1 , L_2 , and L_∞ norms are used to measure the added distortions, and they are respectively named L_0 , L_1 , L_2 , and L_∞ adversarial attacks. Even though no measure can be perfect for human perceptual similarity, these measures or attack types may be employed for different application specifications. This work bridges the literature gap by unifying all the types of attacks with a single intact framework.

In order to benchmark DNN defense techniques and to push for a limit of the DNN security level, we should develop the strongest adversarial attacks. For this purpose, we adopt the white-box attack assumption in that the attackers have complete information about the DNN architectures and all the parameters. This is also a realistic assumption, because even if we only have black-box access to the DNN model, we can train a substitute model and transfer the attacks generated using the substitute model. And for the same purpose, we adopt the optimization-based approach to generate adversarial examples. The objectives of the optimization problem should be (i) misleading the DNN to classify the adversarial example as a target label and (ii) minimizing the L_p norm of the added distortion.

By leveraging ADMM (Alternating Direction Method of Multipliers) [3], an operator splitting optimization approach, we provide a universal framework for L_0 , L_1 , L_2 , and L_∞ adversarial attacks. ADMM decomposes an original optimization problem into two correlated subproblems, each of which can be solved more

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions.acm.org.

MM '18, October 22–26, 2018, Seoul, Republic of Korea

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5665-7/18/10...\$15.00

<https://doi.org/10.1145/3240508.3240639>

Figure 1: Adversarial examples generated by our ADMM L_0 , L_1 , L_2 , and L_∞ attacks for MNIST (left) and CIFAR-10 (right) datasets. The leftmost column contains the original legal inputs. The next four columns are the corresponding adversarial examples crafted using our ADMM L_0 , L_1 , L_2 , and L_∞ attacks, respectively. If the original inputs are correctly classified as label l , then the adversarial examples mislead the DNN to classify them as target label $l + 2$.

efficiently or analytically, and then coordinates solutions to the subproblems to construct a solution to the original problem. This decomposition-alternating procedure of ADMM blends the benefits of dual decomposition and augmented Lagrangian for solving problems with non-convex and combinatorial constraints. Therefore, ADMM introduces no additional sub-optimality besides the original gradient-based backpropagation method commonly used in DNNs and provides a faster linear convergence rate than state-of-the-art iterative attacks [15, 19, 24, 33]. We also compare with the optimization-based approaches, i.e., Carlini & Wagner (C&W) attack [5] and Elastic-net (EAD) attack [6], which are the currently strongest attacks in the literature.

The major contributions of this work and its differences from C&W and EAD attacks are summarized as follows:

- With our ADMM-based universal framework, all the L_0 , L_1 , L_2 , and L_∞ adversarial attacks can be implemented with little modifications, while C&W only performs L_0 , L_2 , and L_∞ attacks and EAD only performs L_1 and L_2 attacks.
- C&W L_0 attack needs to run their L_2 attack iteratively to find the pixels with the least effect and fix them, thereby identifying a minimal subset of pixels for modification to generate an adversarial example.
- C&W L_∞ attack through naively optimization with gradient descent may produce very poor initial results. They solve the issue by introducing a limit on the L_∞ norm and reducing the limit iteratively.
- EAD attack minimizes a weighted sum of L_1 and L_2 norms. However, a universal attack generation model is missing.

- Our extensive experiments show that we are so far the best attacks. Besides the 100% attack success rate, our ADMM-based attacks outperform C&W and EAD in each type of attacks in terms of minimal distortion.

Besides comparing with C&W, EAD and other attacks, we also test our attacks against defenses such as defensive distillation [25] and adversarial training [32], demonstrating the success of our attacks. In addition, we validate the transferability of our attacks onto different DNN models. The codes of our attacks to reproduce the results are available online¹.

2 RELATED WORK

We introduce the representative attacks and defenses in this section.

2.1 Adversarial Attacks

L-BFGS Attack [30] is the first optimization-based attack and is an L_2 attack that uses L_2 norm to measure the distortion in the optimization objective function.

JSMA Attack [24] is an L_0 attack and uses a greedy algorithm that picks the most influential pixels by calculating Jacobian-based Saliency Map and modifies the pixels iteratively. The computational complexity is prohibitive even for applying to ImageNet dataset.

FGSM [11] and IFGSM [19] Attacks are L_∞ attacks and utilize the gradient of the loss function to determine the direction to modify the pixels. They are designed to be fast, rather than optimal. They can be used for adversarial training by directly changing the loss function instead of explicitly injecting adversarial examples into the training data. The fast gradient method (FGM) and the iterative fast gradient method (IFGM) are improvements of FGSM and IFGSM, respectively, that can be fitted as L_1 , L_2 , and L_∞ attacks.

C&W Attacks [5] are a series of L_0 , L_2 , and L_∞ attacks that achieve 100% attack success rate with much lower distortions comparing with the above-mentioned attacks. In particular, the C&W L_2 attack is superior to L-BFGS attack (which is also an L_2 attack) because it uses a better objective function.

EAD Attack [6] formulates the process of crafting adversarial examples as an elastic-net regularized optimization problem. Elastic-net regularization is a linear mixture of L_1 and L_2 norms used in the penalty function. EAD attack is able to craft L_1 -oriented adversarial examples and includes the C&W L_2 attack as a special case.

2.2 Representative Defenses

Defensive Distillation [25] introduces *temperature* into the softmax layer and uses a higher temperature for training and a lower temperature for testing. The training phase first trains a teacher model that can produce soft labels for the training dataset and then trains a distilled model using the training dataset with soft labels. The distilled model is preserved for testing.

Adversarial Training [32] injects adversarial examples with correct labels into the training dataset and then retrain the neural network, thus increasing robustness of DNNs.

3 AN ADMM-BASED UNIVERSAL FRAMEWORK FOR ADVERSARIAL ATTACKS

ADMM was first introduced in the mid-1970s with roots in the 1950s, and the algorithm and theory have been established by the mid-1990s. It was proposed and made popular recently by S. Boyd et al. for statistics and machine learning problems with a very large

¹https://github.com/LinLabNEU/ADMM_adversarial_attack

number of features or training examples [3]. ADMM method takes the form of a decomposition-alternating procedure, in which the solutions to small local subproblems are coordinated to find a solution to a large global problem. It can be viewed as an attempt to blend the benefits of dual decomposition and augmented Lagrangian methods for constrained optimization.

ADMM was developed in part to bring robustness to the dual ascent method, and in particular, to yield convergence without assumptions like strict convexity or finiteness of the objective. ADMM is also capable of dealing with combinatorial constraints due to its decomposition property. It can be used in many practical applications, where the convexity of the objective can not be guaranteed. Besides, it converges fast in many cases since the two arguments are updated in an alternating or sequential fashion, which accounts for the term *alternating direction*.

3.1 Notations and Definitions

In this paper, we mainly evaluate the adversarial attacks with image classification tasks. A two dimensional vector $\mathbf{x} \in \mathbb{R}^{hw}$ represents a gray-scale image with height h and width w . For a colored RGB image with three channels, a three dimensional tensor $\mathbf{x} \in \mathbb{R}^{3hw}$ is utilized to denote it. Each element x_i represents the value of the i -th pixel and is scaled to the range of $[0, 1]$. A neural network has the model $F(\mathbf{x}) = \mathbf{y}$, where F generates an output $\mathbf{y}$ given an input $\mathbf{x}$. Model F is fixed since we perform attacks on given neural network.

The output layer performs softmax operation and the neural network is an m -class classifier. Let the logits $Z(\mathbf{x})$ denote the input to the softmax layer, which represents the output of all layers except for the softmax layer. We have $F(\mathbf{x}) = \text{softmax}(Z(\mathbf{x})) = \mathbf{y}$. The element y_i of the output vector $\mathbf{y}$ represents the probability that input $\mathbf{x}$ belongs to the i -th class. The output vector $\mathbf{y}$ is treated as a probability distribution, and its elements satisfy $0 \leq y_i \leq 1$ and $y_1 + y_2 + \dots + y_m = 1$. The neural network classifies input $\mathbf{x}$ according to the maximum probability, i.e., $C(\mathbf{x}) = \arg \max_i y_i$.

The adversarial attack can be either targeted or untargeted. Given an original legal input $\mathbf{x}_0$ with its correct label t^* , the untargeted adversarial attack is to find an input $\mathbf{x}$ satisfying $C(\mathbf{x}) \neq t^*$ while $\mathbf{x}$ and $\mathbf{x}_0$ are close according to some measure of the distortion. The untargeted adversarial attack does not specify any target label to mislead the classifier. In the targeted adversarial attack, with a given target label $t \neq t^*$, an adversarial example is an input $\mathbf{x}$ such that $C(\mathbf{x}) = t$ while $\mathbf{x}$ and $\mathbf{x}_0$ are close according to some measure of the distortion. In this work, we consider targeted adversarial attacks since they are believed stronger than untargeted attacks.

3.2 General ADMM Framework for Adversarial Attacks

The *initial* problem of constructing adversarial examples is defined as: **Given** an original legal input image $\mathbf{x}_0$ and a target label t , **find** an adversarial example $\mathbf{x}$, **such that** $\mathcal{D}(\mathbf{x} - \mathbf{x}_0)$ is minimized, $C(\mathbf{x}) = t$, and $\mathbf{x} \in [0, 1]^n$. $\mathbf{x} - \mathbf{x}_0$ is the distortion added onto the input $\mathbf{x}_0$. $C(\cdot)$ is the classification function of the neural network and the adversarial example $\mathbf{x} \in [0, 1]^n$ is classified as the target label t .

$\mathcal{D}(\mathbf{x} - \mathbf{x}_0)$ is a measure of the distortion $\mathbf{x} - \mathbf{x}_0$. We need to measure the distortion between the original legal input $\mathbf{x}_0$ and the adversarial example $\mathbf{x}$. L_p norms are the most commonly used measures in the literature. The L_p norm of the distortion between

$\mathbf{x}$ and $\mathbf{x}_0$ is defined as:

$$\|\mathbf{x} - \mathbf{x}_0\|_p = \left(\sum_{i=1}^n |\mathbf{x}_i - \mathbf{x}_{0i}|^p \right)^{\frac{1}{p}} \quad (1)$$

We see the use of L_0 , L_1 , L_2 , and L_∞ norms in different attacks.

- L_0 norm: measures the number of mismatched elements between $\mathbf{x}$ and $\mathbf{x}_0$.
- L_1 norm: measures the sum of the absolute values of the differences between $\mathbf{x}$ and $\mathbf{x}_0$.
- L_2 norm: measures the standard Euclidean distance between $\mathbf{x}$ and $\mathbf{x}_0$.
- L_∞ norm: measures the maximum difference between $\mathbf{x}_i$ and $\mathbf{x}_{0i}$ for all i 's.

In this work, with a general ADMM-based framework, we implement L_0 , L_1 , L_2 , and L_∞ attacks, respectively. When generating adversarial examples in the four attacks, $\mathcal{D}(\mathbf{x} - \mathbf{x}_0)$ in the objective function becomes L_0 , L_1 , L_2 , and L_∞ norms, respectively. For the simplicity of expression, the form of $\mathcal{D}(\mathbf{x} - \mathbf{x}_0)$ is used to denote the measure of $\mathbf{x} - \mathbf{x}_0$. When introducing the detailed four attacks based on the ADMM framework, we utilize the form of L_p norm to represent the distortion measure.

ADMM provides a systematic way to deal with non-convex and combinatorial constraints by breaking the initial problem into two subproblems. To do this, the initial problem is first transformed into the following problem, introducing an auxiliary variable $\mathbf{z}$:

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{z}} \quad & \mathcal{D}(\mathbf{x} - \mathbf{x}_0) + g(\mathbf{z}) \\ \text{s.t.} \quad & \mathbf{x} = \mathbf{z} \\ & \mathbf{z} \in [0, 1]^n \end{aligned} \quad (2)$$

where $g(\mathbf{x})$ has the form:

$$g(\mathbf{x}) = \begin{cases} 0 & \text{if } \max_{i \neq t} (Z(\mathbf{x})_i) - Z(\mathbf{x})_t \leq 0 \\ +\infty & \text{otherwise} \end{cases} \quad (3)$$

Here $Z(\mathbf{x})$ is the logits before the softmax layer. $Z(\mathbf{x})_i$ means the i -th element of $Z(\mathbf{x})$. The function $g(\mathbf{x})$ ensures that the input is classified with target label t . The augmented Lagrangian function of problem (2) is as follows:

$$L_\rho(\mathbf{x}, \mathbf{z}, \mathbf{u}) = \mathcal{D}(\mathbf{x} - \mathbf{x}_0) + g(\mathbf{z}) + \mathbf{u}^T(\mathbf{x} - \mathbf{z}) + \frac{\rho}{2} \|\mathbf{x} - \mathbf{z}\|_2^2 \quad (4)$$

where $\mathbf{u}$ is the dual variable or Lagrange multiplier and $\rho > 0$ is called the penalty parameter. Using the scaled form of ADMM by defining $\mathbf{u} = \rho \mathbf{s}$, we have:

$$L_\rho(\mathbf{x}, \mathbf{z}, \mathbf{s}) = \mathcal{D}(\mathbf{x} - \mathbf{x}_0) + g(\mathbf{z}) + \frac{\rho}{2} \|\mathbf{x} - \mathbf{z} + \mathbf{s}\|_2^2 - \frac{\rho}{2} \|\mathbf{s}\|_2^2 \quad (5)$$

ADMM solves problem (2) through iterations. In the k -th iteration, the following steps are performed:

$$\mathbf{x}^{k+1} = \arg \min_{\mathbf{x}} L_\rho(\mathbf{x}, \mathbf{z}^k, \mathbf{s}^k) \quad (6)$$

$$\mathbf{z}^{k+1} = \arg \min_{\mathbf{z}} L_\rho(\mathbf{x}^{k+1}, \mathbf{z}, \mathbf{s}^k) \quad (7)$$

$$\mathbf{s}^{k+1} = \mathbf{s}^k + \mathbf{x}^{k+1} - \mathbf{z}^{k+1} \quad (8)$$

In Eqn. (6), we find $\mathbf{x}^{k+1}$ which minimizes L_ρ with fixed $\mathbf{z}^k$ and $\mathbf{s}^k$. Similarly, in Eqn. (7), $\mathbf{x}^{k+1}$ and $\mathbf{s}^k$ are fixed and we find $\mathbf{z}^{k+1}$ minimizing L_ρ . $\mathbf{s}^{k+1}$ is then updated accordingly. Note that the two variables $\mathbf{x}$ and $\mathbf{z}$ are updated in an alternating or sequential fashion, from which the term *alternating direction* comes. It converges when:

$$\|\mathbf{x}^{k+1} - \mathbf{z}^{k+1}\|_2^2 \leq \epsilon, \quad \|\mathbf{z}^{k+1} - \mathbf{z}^k\|_2^2 \leq \epsilon \quad (9)$$

Equivalently, in each iteration, we solve two optimization subproblems corresponding to Eqns. (6) and (7), respectively:

$$\min_{\mathbf{x}} \quad \mathcal{D}(\mathbf{x} - \mathbf{x}_0) + \frac{\rho}{2} \|\mathbf{x} - \mathbf{z} + \mathbf{s}\|_2^2 \quad (10)$$

and

$$\min_z \quad g(\mathbf{z}) + \frac{\rho}{2} \|\mathbf{x} - \mathbf{z} + \mathbf{s}\|_2^2 \quad (11)$$

The non-differentiable $g(\mathbf{x})$ makes it difficult to solve the second subproblem (11). Therefore, a new differentiable $g(\mathbf{x})$ inspired by [5] is utilized as follows:

$$g(\mathbf{x}) = c \cdot \max \left(\left(\max_{i \neq t} (Z(\mathbf{x})_i) - Z(\mathbf{x})_t \right), -\kappa \right) \quad (12)$$

Then, stochastic gradient decent methods can be used to solve this subproblem. The Adam optimizer [16] is applied due to its fast and robust convergence behavior. In the new $g(\mathbf{x})$ of Eqn. (12), κ is a confidence parameter denoting the strength of adversarial example transferability. The larger κ , the stronger transferability of the adversarial example. It can be kept as 0 if we do not evaluate the transferability.

3.3 Box Constraint

The constraint on $\mathbf{z}$ i.e., $\mathbf{z} \in [0, 1]^n$ is known as a “box constraint” in the optimization literature. We use a new variable $\mathbf{w}$ and instead of optimizing over $\mathbf{z}$ defined above, we optimize over $\mathbf{w}$, based on:

$$\mathbf{z} = \frac{1}{2} (\tanh(\mathbf{w}) + 1) \quad (13)$$

Here the $\tanh(\cdot)$ is performed elementwise. Since $-1 \leq \tanh(w_i) \leq 1$, the method will automatically satisfy the box constraint and allows us to use optimization algorithms that do not natively support box constraints.

3.4 Selection of Target Label

There are different ways to choose the target labels:

- *Average Case*: select at random the target label uniformly among all the labels that are not the correct label.
- *Best Case*: perform attacks using all incorrect labels, and report the target label that is the least difficult to attack.
- *Worst Case*: perform attacks using all incorrect labels, and report the target label that is the most difficult to attack.

We evaluate the performs of the proposed ADMM attacks in the three cases mentioned above.

3.5 Discussion on Constants

There are two constants c and ρ in the two subproblems (10) and (11). Different policies are adopted for choosing appropriate c and ρ in L_0 , L_1 , L_2 and L_∞ attacks. In L_2 attack, since ρ acts in both problems (10) and (11), we fix ρ and change c to improve the solutions. We find that the best choice of $c > 0$ is the smallest one that can help achieve $g(\mathbf{x}) = 0$ in the subproblem (11). Thus, a modified binary search is used to find a satisfying c . For the ADMM L_0 attack, as ρ has stronger and more direct influence on the solutions, c is fixed and adaptive search of ρ is utilized. More details are provided in Section 4.2. For the ADMM L_1 and L_∞ attacks, as we find fixed c and ρ can achieve good performance, c and ρ are kept unchanged and adaptive search method is not used.

4 INSTANTIATIONS OF L_0 , L_1 , L_2 AND L_∞ ATTACKS BASED ON ADMM FRAMEWORK

The ADMM framework for adversarial attacks now need to solve two subproblems (10) and (11). The difference between L_0 , L_1 , L_2 and L_∞ attacks lies in the subproblem (10), while the processes to find the solutions of the subproblem (11) based on stochastic gradient descent method are the very similar for the four attacks.

4.1 L_2 Attack

For L_2 attack, the subproblem (10) has the form:

$$\min_{\mathbf{x}} \quad \|\mathbf{x} - \mathbf{x}_0\|_2^2 + \frac{\rho}{2} \|\mathbf{x} - \mathbf{z} + \mathbf{s}\|_2^2 \quad (14)$$

the solution to which can be directly derived in an analytical format:

$$\mathbf{x} = \frac{\rho}{2 + \rho} (\mathbf{z} - \mathbf{s}) + \frac{2}{2 + \rho} \mathbf{x}_0 \quad (15)$$

Then the complete solution to the L_2 attack problem using the ADMM framework is as follows: for the k -th iteration,

$$\mathbf{x}^{k+1} = \frac{\rho}{2 + \rho} \left(\left(\frac{1}{2} (\tanh(\mathbf{w}^k) + 1) \right) - \mathbf{s}^k \right) + \frac{2}{2 + \rho} \mathbf{x}_0 \quad (16)$$

$$\begin{aligned} \mathbf{w}^{k+1} = \arg \min_{\mathbf{w}} \quad & \left(g \left(\frac{1}{2} (\tanh(\mathbf{w}) + 1) \right) \right. \\ & \left. + \frac{\rho}{2} \left\| \mathbf{x}^{k+1} - \left(\frac{1}{2} (\tanh(\mathbf{w}) + 1) \right) + \mathbf{s}^k \right\|_2^2 \right) \end{aligned} \quad (17)$$

$$\mathbf{s}^{k+1} = \mathbf{s}^k + \mathbf{x}^{k+1} - \left(\frac{1}{2} (\tanh(\mathbf{w}^{k+1}) + 1) \right) \quad (18)$$

Eqn. (16) corresponds to the analytical solution to the subproblem (10) i.e., problem (14) with Eqn. (13) replacing $\mathbf{z}$ in Eqn. (15). Eqn. (17) corresponds to the subproblem (11) with Eqn. (13) replacing $\mathbf{z}$ and g taking the form of Eqn. (12). The solution to Eqn. (17) is derived through the Adam optimizer with stochastic gradient descent.

4.2 L_0 Attack

For L_0 attack, the subproblem (10) has the form:

$$\min_{\mathbf{x}} \quad \|\mathbf{x} - \mathbf{x}_0\|_0 + \frac{\rho}{2} \|\mathbf{x} - \mathbf{z} + \mathbf{s}\|_2^2 \quad (19)$$

Its equivalent optimization problem is as follows:

$$\min_{\boldsymbol{\delta}} \quad \|\boldsymbol{\delta}\|_0 + \frac{\rho}{2} \|\boldsymbol{\delta} - \mathbf{z} + \mathbf{s} + \mathbf{x}_0\|_2^2 \quad (20)$$

The solution to problem (19) can be obtained through $\mathbf{x}^* = \mathbf{x}_0 + \boldsymbol{\delta}^*$ where $\boldsymbol{\delta}^*$ is the solution to problem (20). The solution to problem (20) can be derived in this way: let $\boldsymbol{\delta}$ be equal to $\mathbf{z} - \mathbf{s} - \mathbf{x}_0$ first, then for each element in $\boldsymbol{\delta}$, if its square is smaller than $\frac{2}{\rho}$, make it zero. A proof for the solution is given in the following.

LEMMA 4.1. *Suppose that two matrices $\mathbf{A}$, $\mathbf{B}$ are of the same size, and that there are at least k zero elements in $\mathbf{A}$. Then the optimal value of the following problem is the sum of the square of the k smallest elements in $\mathbf{B}$.*

$$\min_{\mathbf{A}} \quad \|\mathbf{A} - \mathbf{B}\|_2^2 \quad (21)$$

The proof for the lemma is straightforward and we omit it for the sake of brevity. We use $h(\mathbf{x}, k)$ to denote the sum of the k smallest $\mathbf{x}_i^2$ ($\mathbf{x}_i$ is an element in $\mathbf{x}$).

THEOREM 4.2. *Set $\boldsymbol{\delta} = \mathbf{z} - \mathbf{s} - \mathbf{x}_0$ and then make those elements in $\boldsymbol{\delta}$ zeros if their square are smaller than $\frac{2}{\rho}$. Such $\boldsymbol{\delta}$ would yield the minimum objective value of problem (20).*

PROOF. Suppose that δ_1 is constructed according to the above rule in Theorem 1, and δ_1 has k_1 elements equal to 0. We need to prove that δ_1 is the optimal solution with the minimum objective value. Suppose we have another arbitrary solution δ_2 with k_2 elements equal to 0. Both δ_1 and δ_2 have n elements. The objective value of solution δ_1 is:

$$\|\delta_1\|_0 + \frac{\rho}{2} h(z - s - \mathbf{x}_0, k_1) = n - k_1 + \frac{\rho}{2} h(z - s - \mathbf{x}_0, k_1) \quad (22)$$

The objective value of solution δ_2 is:

$$\begin{aligned} \|\delta_2\|_0 + \frac{\rho}{2} \|\delta_2 - z + s + \mathbf{x}_0\|_2^2 &\geq \|\delta_2\|_0 + \frac{\rho}{2} h(z - s - \mathbf{x}_0, k_2) \\ &= n - k_2 + \frac{\rho}{2} h(z - s - \mathbf{x}_0, k_2) \end{aligned} \quad (23)$$

The inequality in Eqn. (23) holds due to Lemma 4.1.

If $k_2 > k_1$, then according to the definition of δ_1 , we have

$$h(z - s - \mathbf{x}_0, k_2) - h(z - s - \mathbf{x}_0, k_1) > \frac{2}{\rho} (k_2 - k_1) \quad (24)$$

$$(n - k_2 + \frac{\rho}{2} h(z - s - \mathbf{x}_0, k_2)) - (n - k_1 + \frac{\rho}{2} h(z - s - \mathbf{x}_0, k_1)) > 0 \quad (25)$$

If $k_1 > k_2$, then according to the definition of δ_1 , we have

$$h(z - s - \mathbf{x}_0, k_1) - h(z - s - \mathbf{x}_0, k_2) < \frac{2}{\rho} (k_1 - k_2) \quad (26)$$

$$(n - k_2 + \frac{\rho}{2} h(z - s - \mathbf{x}_0, k_2)) - (n - k_1 + \frac{\rho}{2} h(z - s - \mathbf{x}_0, k_1)) > 0 \quad (27)$$

Thus, we can see that our solution δ_1 can achieve the minimum objective value and it is the optimal solution. $\square$

When solving the subproblem (19) according to Theorem 4.2, we enforce a hidden constraint on the distortion $\delta = \mathbf{x} - \mathbf{x}_0$, that the square of each non-zero element in δ must be larger than $\frac{2}{\rho}$. Therefore, a smaller ρ would push ADMM method to find δ with larger non-zero elements, thus reducing the number of non-zero elements and decreasing L_0 norm. Empirically, we find the constant ρ represents a trade-off between attack success rate and L_0 norm of the distortion, i.e., a larger ρ can help find solutions with higher attack success rate at the cost of larger L_0 norm of the distortion.

Then the complete solution to the L_0 attack problem using the ADMM framework can be derived similar to the L_2 attack. More specifically, in each iteration, Theorem 4.2 is applied to obtain the optimal δ and $\mathbf{x}$. Then we solve Eqn. (17) with Adam optimizer and update parameters through (18).

4.3 L_1 Attack

For L_1 attack, the subproblem (10) has the form:

$$\min_{\mathbf{x}} \|\mathbf{x} - \mathbf{x}_0\|_1 + \frac{\rho}{2} \|\mathbf{x} - z + s\|_2^2 \quad (28)$$

Problem (28) has the closed-form solution. If we change the variable $\delta = \mathbf{x} - \mathbf{x}_0$, then problem becomes

$$\min_{\delta} \|\delta\|_1 + \frac{\rho}{2} \|\delta + \mathbf{x}_0 - z + s\|_2^2. \quad (29)$$

The solution of problem (29) is given by the soft thresholding operator evaluated at the point $(z - s - \mathbf{x}_0)$ with a parameter $1/\rho$ [26],

$$\delta^* = (z - s - \mathbf{x}_0 - 1/\rho)_+ - (-(z - s - \mathbf{x}_0) - 1/\rho)_+, \quad (30)$$

where $(\cdot)_+$ is taken in elementwise, and $(x)_+ = x$ if $x \geq 0$, and 0 otherwise. Therefore, the solution to problem (28) is given by

$$\mathbf{x}^* = \mathbf{x}_0 + \delta^*. \quad (31)$$

Table 1: Number of Iterations in Adversarial Attacks

	C&W L_2	C&W L_∞	C&W L_0	ADMM
# of iterations	$N \times M$	$J \times N \times M$	$L \times N \times M$	$K \times N \times M$

N is the number of binary search steps and M is the number of iterations in each binary search step. J and L denote the iteration number of L_2 attacks in L_∞ and L_0 attacks, respectively. K is the number of ADMM iterations. Note that $N = 1$ in ADMM L_1 and L_∞ attacks when binary search is not utilized.

The complete solution to the L_1 attack problem using the ADMM framework is similar to the L_2 attack. In each iteration, we obtain the closed-form solution of the first subproblem (28) and then Adam optimizer is utilized to solve the second subproblem (17). Next we update the parameters through Eqn. (18).

4.4 L_∞ Attack

For L_∞ attack, the subproblem (10) has the form:

$$\min_{\mathbf{x}} \|\mathbf{x} - \mathbf{x}_0\|_\infty + \frac{\rho}{2} \|\mathbf{x} - z + s\|_2^2 \quad (32)$$

This problem does not have a closed form solution. One possible method is to derive the KKT conditions of problem (32) [26]. Here we use stochastic gradient decent methods to solve it. In the experiments, we find that the Adam optimizer [16] could achieve fast and robust convergence results. So Adam optimizer is utilized to solve Eqn. (32). Since Eqn. (32) is relatively simpler compared with Eqn. (17), the complexity for solving Eqn. (32) is negligible.

The complete solution to the L_∞ attack problem using the ADMM framework can be derived similar to the L_2 attack. In the k -th iteration, we first use Adam optimizer to get the optimal $\mathbf{x}^{k+1}$ in Eq. (32). Then we solve Eq. (17) and update parameters through Eq. (18) as the L_2 attack.

4.5 Computation Complexity

We evaluate the computation cost in terms of iteration complexity in Table 1. For comparison, we also present the complexity of C&W attacks. The C&W L_∞ and L_0 attacks utilize L_2 attack method as the inner loop, so their total iterations are J and L times of the iterations in L_2 attack, respectively, where J and L vary for different parameter settings. In the ADMM attack, the number of iterations should be $K \times N \times M$. Note that in ADMM L_1 and L_∞ attacks, binary search is not applied, which yields $N = 1$ and $K \times M$ iterations. It has been proved in [15] that ADMM has at least the linear convergence rate, and it empirically converges in a few tens of iterations.

Besides, this paper focuses on the offline generation of adversarial examples. Therefore, the computation complexity is not the most important performance metric only if we successfully craft adversarial examples. Therefore, the attack success rate and the distortion are the most important performance measures.

Table 2: Parameter setting of ADMM attacks

datasets	attacks	ρ	c	K	# of BS	LR of BS
MNIST	ADMM L_2	20	0.001	10	1000	0.02
	ADMM L_0	3	20	10	1000	0.01
	ADMM L_1	10	2	80	2000	0.1
	ADMM L_∞	0.1	0.1	100	2000	0.001
CIFAR-10	ADMM L_2	20	0.001	10	1000	0.002
	ADMM L_0	40	200	10	1000	0.01
	ADMM L_1	300	40	80	2000	0.001
	ADMM L_∞	0.1	0.1	100	2000	0.002
ImageNet	ADMM L_2	20	0.001	10	1000	0.002
	ADMM L_1	2000	200	80	2000	0.001
	ADMM L_∞	0.1	0.1	100	2000	0.002

BS is short for "binary search". If binary search is applied, the ρ or c value in the table is the initial value. # of BS denotes the number of iterations in each binary search. LR is short for learning rate.

5 PERFORMANCE EVALUATION

The proposed ADMM attacks are compared with state-of-the-art attacks, including C&W attacks [5], EAD attack, FGM and IFGM attacks, on three image classification datasets, MNIST [20], CIFAR-10 [17] and ImageNet [7]. We also test our attacks against two defenses, defensive distillation [25] and adversarial training [32], and evaluate the transferability of ADMM attacks.

5.1 Experiment Setup and Parameter Setting

Our experiment setup is based on C&W attack setup for fair comparisons. Two networks are trained for MNIST and CIFAR-10 datasets, respectively. For the ImageNet dataset, a pre-trained network is utilized. The network architecture for MNIST and CIFAR-10 has four convolutional layers, two max pooling layers, two fully connected layers and a softmax layer. It can achieve 99.5% accuracy on MNIST and 80% accuracy on CIFAR-10. For ImageNet, a pre-trained Inception v3 network [29] is applied so there is no need to train our own model. The Google Inception model can achieve 96% top-5 accuracy with image inputs of size $299 \times 299 \times 3$. All experiments are conducted on machines with an Intel I7-7700K CPU, 32 GB RAM and an NVIDIA GTX 1080 TI GPU.

The implementations of FGM and IFGM are based on the CleverHans package [23]. The key distortion parameter ϵ is determined through a fine-grained grid search. For each image, the smallest ϵ in the grid leading to a successful attack is reported. For IFGM, we perform 10 FGM iterations. The distortion parameter ϵ' in each FGM iteration is set to be $\epsilon/10$, which is effective shown in [32].

The implementations of C&W attacks² and EAD attack³ are based on the GitHub code released by the authors. The EAD attack has two decision rules when selecting the final adversarial example: the least elastic-net (EN) and L_1 distortion measurement (L_1). Usually, the L_1 decision rule can achieve lower L_1 distortion than the EN decision rule as the EN decision rule considers a mixture of L_1 and L_2 distortions. We use the L_1 decision rule for fair comparison. The key parameter β is set to 0.001.

The parameters of ADMM attacks are summarized in Table 2. If binary search is applied, the number of binary search steps is 9.

5.2 Attack Success Rate and Distortion for ADMM L_2 attack

The ADMM L_2 attack is compared with FGM, IFGM and C&W L_2 attacks. The attack success rate (ASR) represents the percentage of the constructed adversarial examples that are successfully classified as target labels. The average distortion of all successful adversarial examples is reported. For zero ASR, its distortion is not available (N.A.). We craft adversarial examples on MNIST, CIFAR-10 and ImageNet. For MNIST and CIFAR-10, 1000 correctly classified images are randomly selected from the test sets and 9 target labels are tested for each image, so we perform 9000 attacks for each dataset using each attack method. For ImageNet, 100 correctly classified images are randomly selected each with 9 random target labels.

Table 3 shows the results on MNIST, CIFAR-10 and ImageNet. As we can see, FGM fails to generate adversarial examples with high success rate since it is designed to be fast, rather than optimal. Among IFGM, C&W and ADMM L_2 attacks, ADMM achieves the lowest L_2 distortion for the best case, average case and worst case. IFGM has larger L_2 distortions compared with C&W and ADMM attacks on the three datasets, especially on ImageNet. For MNIST, the ADMM attack can reduce the L_2 distortion by about 7% compared

with C&W L_2 attack. This becomes more prominent on ImageNet that ADMM reduces L_2 distortion by 19% comparing with C&W in the worst case.

We also observe that on CIFAR-10, ADMM L_2 attack can achieve lower L_2 distortions but the reductions are not as prominent as that on MNIST or ImageNet. The reason may be that CIFAR-10 is the easiest dataset to attack since it requires the lowest L_2 distortion among the three datasets. So both ADMM L_2 attack and C&W L_2 attack can achieve quite good performance. Note that in most cases on the three datasets, ADMM L_2 attack can achieve lower L_1 , L_2 and L_∞ distortions than C&W L_2 attack, indicating a comprehensive enhancement of the ADMM L_2 attack over C&W L_2 attack.

5.3 Attack Success Rate and Distortion for ADMM L_0 attack

The performance of ADMM L_0 attack in terms of attack success rate and L_0 norm of distortion is demonstrated in this section. The ADMM L_0 attack is compared with C&W L_0 attack on MNIST and CIFAR-10. 500 images are randomly selected from the test sets of MNIST and CIFAR-10, respectively. Each image has 9 target labels and we perform 4500 attacks for each dataset using either ADMM or C&W L_0 attack.

The results of the L_0 attacks are shown in Table 4. As observed from the table, both C&W and ADMM L_0 attacks can achieve 100% attack success rate. For the best case, C&W L_0 attack and ADMM L_0 attack have relatively close performance in terms of L_0 distortion. For the worst case, ADMM L_0 attack can achieve lower L_0 distortion than C&W. ADMM L_0 attack reduces the L_0 distortion by up to 17% on MNIST. We also note that the differences between C&W and ADMM L_0 attacks are smaller on CIFAR-10 than that on MNIST.

5.4 Attack Success Rate and Distortion for ADMM L_1 attack

We compare the ADMM L_1 attack with FGM, IFGM and EAD L_1 [6] attacks. The attack success rate (ASR) and the average distortion of all successful adversarial examples are reported. We perform the adversarial L_1 attacks on MNIST, CIFAR-10 and ImageNet. For MNIST and CIFAR-10, 1000 correctly classified images are randomly selected from the test sets and 9 target labels are tested for each image, so we perform 9000 attacks for each dataset using each attack method. For ImageNet, 100 correctly classified images and 9 target labels are randomly selected.

The results of the ADMM L_1 attack are shown in Table 5. We can observe that both EAD and ADMM L_1 attacks can achieve 100% attack success rate while FGM L_1 attack has bad performance and IFGM L_1 attack can not guarantee 100% ASR on ImageNet. ADMM L_1 attack can achieve the best performance compared with FGM, IFGM, and EAD L_1 attacks. As demonstrated in Table 5, the L_1 distortion measurements of ADMM and EAD L_1 attacks are relatively close in the best case while the improvement of ADMM L_1 attack over EAD L_1 attack is much larger for the worst case. In the best case, the ADMM L_1 attack can craft adversarial examples with a L_1 norm about 14% smaller than that of the EAD L_1 attack on MNIST, CIFAR-10 and ImageNet. For the worst case, the L_1 norm of ADMM L_1 attack is about 28% lower on CIFAR-10 and 50% lower on ImageNet compared with that of EAD L_1 attack.

5.5 Attack Success Rate and Distortion for ADMM L_∞ attack

The ADMM L_∞ attack is compared with FGM and IFGM L_∞ attacks. The attack success rate (ASR) and the average distortion

²https://github.com/carlini/nn_robust_attacks

³<https://github.com/ysharma1126/EAD-Attack>

Table 3: Adversarial attack success rate (ASR) and distortion of different L_2 attacks for different datasets

Data Set	Attack Method	Best Case				Average Case				Worst Case			
		ASR	L_2	L_1	L_∞	ASR	L_2	L_1	L_∞	ASR	L_2	L_1	L_∞
MNIST	FGM(L_2)	99.4	2.245	25.84	0.574	34.6	3.284	39.15	0.747	0	N.A.	N.A.	N.A.
	IFGM(L_2)	100	1.58	18.51	0.388	99.9	2.50	32.63	0.562	99.6	3.958	55.04	0.783
	C&W(L_2)	100	1.393	13.57	0.402	100	2.002	22.31	0.54	99.9	2.598	31.43	0.689
	ADMM(L_2)	100	1.288	13.87	0.345	100	1.873	22.52	0.498	100	2.445	31.427	0.669
CIFAR-10	FGM(L_2)	99.5	0.421	14.13	0.05	42.8	1.157	39.5	0.136	0.7	3.115	107.1	0.369
	IFGM(L_2)	100	0.191	6.549	0.022	100	0.432	15.13	0.047	100	0.716	25.22	0.079
	C&W(L_2)	100	0.178	6.03	0.019	100	0.347	12.115	0.0364	99.9	0.481	16.75	0.0536
	ADMM(L_2)	100	0.173	5.8	0.0192	100	0.337	11.65	0.0365	100	0.476	16.73	0.0535
ImageNet	FGM(L_2)	12	2.29	752.9	0.087	1	6.823	2338	0.25	0	N.A.	N.A.	N.A.
	IFGM(L_2)	100	1.057	349.55	0.034	100	2.461	823.52	0.083	98	4.448	1478.8	0.165
	C&W(L_2)	100	0.48	142.4	0.016	100	0.681	215.4	0.03	100	0.866	275.4	0.042
	ADMM(L_2)	100	0.416	117.3	0.015	100	0.568	177.6	0.022	97	0.701	229.08	0.0322

Table 4: Adversarial attack success rate and distortion of ADMM and C&W L_0 attacks for MNIST and CIFAR-10

Dataset	Attack method	Best case		Average case		Worst case	
		ASR	L_0	ASR	L_0	ASR	L_0
MNIST	C&W(L_0)	100	8.1	100	17.48	100	31.48
	ADMM(L_0)	100	8	100	15.71	100	25.87
CIFAR	C&W(L_0)	100	8.6	100	19.6	100	34.4
	ADMM(L_0)	100	8.25	100	18.8	100	31.2

of all successful adversarial examples are reported. We perform the adversarial L_∞ attacks on MNIST, CIFAR-10 and ImageNet. For MNIST and CIFAR-10, 1000 correctly classified images are randomly selected from the test sets and 9 target labels are tested for each image, so we perform 9000 attacks for each dataset using each attack method. For ImageNet, 100 correctly classified images and 9 target labels are randomly selected.

The results of the ADMM L_∞ attack are demonstrated in Table 6. We can observe that both IFGM and ADMM L_∞ attacks can achieve 100% attack success rate while FGM has bad performance. ADMM L_∞ attack can achieve the best performance compared with FGM and IFGM L_∞ attacks. We also note that the L_∞ norms of ADMM and IFGM L_∞ attacks are relatively close in the best case. Usually the L_∞ distortion measure of ADMM attack is smaller than that of IFGM attack by no larger than 10% for the best case. In the worst case, the improvement of ADMM L_∞ attack over IFGM L_∞ attack is much more obvious. The L_∞ distortion measure of ADMM attack is about 40% smaller than that of IFGM attack on MNIST or CIFAR-10 dataset for the worst case. On ImageNet, the L_∞ norm of ADMM attack is 64% lower than that of IFGM attack.

5.6 ADMM Attack Against Defensive Distillation and Adversarial Training

ADMM attacks can break the undefended DNNs with high success rate. It is also able to break DNNs with defensive distillation. We perform C&W L_2 attack, ADMM L_0 , L_1 , L_2 and L_∞ attack for different temperature parameters on MNIST and CIFAR-10. 500 randomly selected images are used as source to generate 4500 adversarial examples with 9 targets for each image on MNIST or CIFAR-10. We find that the attack success rates of C&W L_2 attack and ADMM four attacks for different temperature T are all 100%. Since distillation

Figure 2: L_2 distortion of adversarial training for three cases on MNIST

at temperature T causes the value of logits to be approximately T times larger while the relative values of logits remain unchanged, C&W attack and ADMM attack which work on the relative values of logits do not fail.

We further test ADMM attack against adversarial training on MNIST. C&W L_2 attack and ADMM L_2 attack are utilized to separately generate 9000 adversarial examples with 1000 randomly selected images from the training set as sources. Then we add the adversarial examples with correct labels into the training dataset and retrain the network with the enlarged training dataset. With the retained network, we perform ADMM attack on the adversarially trained networks (one with C&W adversarial examples, and one with ADMM adversarial examples), as shown in Fig. 2. ADMM L_2 attack can break all three networks (one unprotected, one retained with C&W adversarial examples, and one retained with ADMM adversarial examples) with 100% success rate. L_2 distortions on the latter two networks are higher than that on the first network, showing some defense effect of adversarial training. We also note that L_2 distortion on the third network is higher than the second network, which demonstrates higher defense efficiency of performing adversarial training with ADMM adversarial examples (partly because ADMM attack is stronger).

5.7 Attack Transferability

Here we test the transferability of ADMM adversarial attack. For each value of confidence parameter κ , we use ADMM L_2 attack and C&W L_2 attack to generate 9000 adversarial examples on MNIST, respectively. Then these examples are applied to attack the defensively distilled network with temperature $T = 100$. The ASR

Table 5: Adversarial attack success rate (ASR) and distortion of different L_1 attacks for different datasets

Data Set	Attack Method	Best Case				Average Case				Worst Case			
		ASR	L_1	L_2	L_∞	ASR	L_1	L_2	L_∞	ASR	L_1	L_2	L_∞
MNIST	FGM(L_1)	100	29.6	2.42	0.57	36.5	51.2	3.99	0.8	0	N.A.	N.A.	N.A.
	IFGM(L_1)	100	18.7	1.6	0.41	100	33.9	2.6	0.58	100	54.8	4.04	0.81
	EAD(L_1)	100	7.08	1.49	0.56	100	12.5	2.08	0.77	100	18.8	2.57	0.92
	ADMM(L_1)	100	6.0	2.07	0.97	100	10.61	2.72	0.99	100	16.6	3.41	1
CIFAR-10	FGM(L_1)	98.5	18.25	0.53	0.057	47	48.32	1.373	0.142	1	33.99	0.956	0.101
	IFGM(L_1)	100	6.28	0.184	0.21	100	13.72	0.394	0.44	100	22.84	0.65	0.74
	EAD(L_1)	100	2.44	0.31	0.084	100	6.392	0.6	0.185	100	10.21	0.865	0.31
	ADMM(L_1)	100	2.09	0.319	0.102	100	5.0	0.591	0.182	100	7.453	0.77	0.255
ImageNet	FGM(L_1)	12	229	0.73	0.028	1	67	0.165	0.08	0	N.A.	N.A.	N.A.
	IFGM(L_1)	93	311	0.966	0.033	67	498.5	1.5	0.051	47	720.2	2.2	0.08
	EAD(L_1)	100	65.4	0.632	0.047	100	165.5	1.02	0.06	100	290	1.43	0.08
	ADMM(L_1)	100	56.1	0.904	0.053	100	92.7	1.15	0.0784	100	142.1	1.473	0.102

Table 6: Adversarial attack success rate (ASR) and distortion of different L_∞ attacks for different datasets

Data Set	Attack Method	Best Case				Average Case				Worst Case			
		ASR	L_∞	L_1	L_2	ASR	L_∞	L_1	L_2	ASR	L_∞	L_1	L_2
MNIST	FGM(L_∞)	100	0.194	84.9	4.04	35	0.283	122.7	5.85	0	N.A.	N.A.	N.A.
	IFGM(L_∞)	100	0.148	50.9	2.48	100	0.233	71.2	3.44	100	0.378	96.8	4.64
	ADMM(L_∞)	100	0.135	35.9	2.068	100	0.178	48	2.73	100	0.218	60.2	3.37
CIFAR-10	FGM(L_∞)	100	0.015	42.8	0.78	53	0.48	136	2.5	1.5	0.31	712	14
	IFGM(L_∞)	100	0.0063	14.36	0.28	100	0.015	26.2	0.54	100	0.026	37.7	0.826
	ADMM(L_∞)	100	0.0061	12.8	0.25	100	0.0114	23.07	0.47	100	0.017	31.9	0.65
ImageNet	FGM(L_∞)	20	0.0873	22372	43.55	1.5	0.0005	134	0.26	0	N.A.	N.A.	N.A.
	IFGM(L_∞)	100	0.0046	542.4	1.27	100	0.0128	1039.6	2.54	100	0.0253	1790.2	4.4
	ADMM(L_∞)	100	0.0041	280.2	0.773	100	0.0059	427.7	1.10	100	0.0092	624.1	1.6

is reported in Fig. 3. As demonstrated in Fig. 3, when κ is small, ADMM L_2 attack can hardly achieve success on the defensively distilled network, which means the generated adversarial examples are not strong enough to break the defended network. Low transferability of the generated adversarial examples is observed when κ is low. As κ increases, the ASRs of the three cases increase, demonstrating increasing transferability. When $\kappa = 50$, the ASRs of three cases can achieve the maximum value. The ASR of average case is nearly 98%, meaning most of the generated adversarial examples on the undefended network can also break the defensively distilled network with $T = 100$. Also note that when $\kappa > 50$, the ASRs of average case and worst case decrease as κ increases. The reason is that it's quite difficult to generate adversarial examples even for the undefended network when κ is very large. Thus an decrease on the ASR is observed for average case and worst case, and the advantages of strong transferable adversarial examples are mitigated by the difficulty to generate such strong attacks. We also note that when $\kappa > 40$, the ASRs of ADMM L_2 attack for average case and worst case are higher than the ASRs of C&W L_2 attack, demonstrating higher transferability of the ADMM attack.

6 CONCLUSION

In this paper, we propose an ADMM-based general framework for adversarial attacks. Under the ADMM framework, L_0 , L_1 , L_2 and L_∞ attacks are proposed and implemented. We compare the ADMM attacks with state-of-the-art adversarial attacks, showing ADMM attacks are so far the strongest. The ADMM attack is also applied to break two defense methods, the defensive distillation and

Figure 3: transferability evaluation of C&W and ADMM L_2 attacks on MNIST

adversarial training. Experimental results show the effectiveness of the proposed ADMM attacks with strong transferability.

ACKNOWLEDGEMENTS

This work is supported by the National Science Foundation (CCF-1733701, CNS-1704662, and CNS-1739748), Air Force Research Laboratory FA8750-18-2-0058, and U.S. Office of Naval Research. We thank researchers at the US Naval Research Laboratory for their comments on previous drafts of this paper.

REFERENCES

- [1] Anish Athalye, Nicholas Carlini, and David Wagner. 2018. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. *arXiv preprint arXiv:1802.00420* (2018).

- [2] Arjun Nitin Bhagoji, Daniel Cullina, and Prateek Mittal. 2017. Dimensionality reduction as a defense against evasion attacks on machine learning classifiers. *arXiv preprint arXiv:1704.02654* (2017).
- [3] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning* 3, 1 (2011), 1–122.
- [4] Nicholas Carlini, Pratyush Mishra, Tavish Vaidya, Yuankai Zhang, Micah Sherr, Clay Shields, David Wagner, and Wenchao Zhou. 2016. Hidden Voice Commands. In *USENIX Security Symposium*. 513–530.
- [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] Pin-Yu Chen, Yash Sharma, Huan Zhang, Jinfeng Yi, and Cho-Jui Hsieh. 2017. EAD: elastic-net attacks to deep neural networks via adversarial examples. *arXiv preprint arXiv:1709.04114* (2017).
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, 248–255.
- [8] Guneet S Dhillon, Kamyar Azizzadenesheli, Zachary C Lipton, Jeremy Bernstein, Jean Kossaifi, Aran Khanna, and Anima Anandkumar. 2018. Stochastic Activation Pruning for Robust Adversarial Defense. *arXiv preprint arXiv:1803.01442* (2018).
- [9] Gintare Karolina Dziugaite, Zoubin Ghahramani, and Daniel M Roy. 2016. A study of the effect of jpg compression on adversarial images. *arXiv preprint arXiv:1608.00853* (2016).
- [10] Reuben Feinman, Ryan R Curtin, Saurabh Shintre, and Andrew B Gardner. 2017. Detecting adversarial samples from artifacts. *arXiv preprint arXiv:1703.00410* (2017).
- [11] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014).
- [12] Chuan Guo, Mayank Rana, Moustapha Cissé, and Laurens van der Maaten. 2017. Countering Adversarial Images using Input Transformations. *arXiv preprint arXiv:1711.00117* (2017).
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [14] Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al. 2012. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine* 29, 6 (2012), 82–97.
- [15] Mingyi Hong and Zhi-Quan Luo. 2017. On the linear convergence of the alternating direction method of multipliers. *Mathematical Programming* 162, 1 (01 Mar 2017), 165–199. <https://doi.org/10.1007/s10107-016-1034-2>
- [16] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. 2015 *ICLR* arXiv preprint arXiv:1412.6980 (2015). arXiv:1412.6980 <http://arxiv.org/abs/1412.6980>
- [17] A. Krizhevsky and G. Hinton. 2009. Learning multiple layers of features from tiny images. *Master's thesis, Department of Computer Science, University of Toronto* (2009).
- [18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105.
- [19] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. 2016. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533* (2016).
- [20] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (Nov 1998), 2278–2324. <https://doi.org/10.1109/5.726791>
- [21] Konstantinos Makantasis, Konstantinos Karantzas, Anastasios Doulamis, and Nikolaos Doulamis. 2015. Deep supervised learning for hyperspectral data classification through convolutional neural networks. In *Geoscience and Remote Sensing Symposium (IGARSS), 2015 IEEE International*. IEEE, 4959–4962.
- [22] Anh Nguyen, Jason Yosinski, and Jeff Clune. 2015. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 427–436.
- [23] Nicolas Papernot, Ian Goodfellow, Ryan Sheatsley, Reuben Feinman, and Patrick McDaniel. 2016. cleverhans v1.0.0: an adversarial machine learning library. *arXiv preprint arXiv:1610.00768* (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, Xi Wu, Somesh Jha, and Ananthram Swami. 2016. Distillation as a defense to adversarial perturbations against deep neural networks. In *Security and Privacy (SP), 2016 IEEE Symposium on*. IEEE, 582–597.
- [26] Neal Parikh, Stephen Boyd, et al. 2014. Proximal algorithms. *Foundations and Trends® in Optimization* 1, 3 (2014), 127–239.
- [27] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature* 529, 7587 (2016), 484–489.
- [28] Jiawei Su, Danilo Vasconcellos Vargas, and Sakurai Kouichi. 2017. One pixel attack for fooling deep neural networks. *arXiv preprint arXiv:1710.08864* (2017).
- [29] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. 2016. Rethinking the Inception Architecture for Computer Vision. 2016 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016), 2818–2826.
- [30] 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).
- [31] Yaniv Taigman, Ming Yang, Marc'Aurelio Ranzato, and Lior Wolf. 2014. Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1701–1708.
- [32] F. Tramèr, A. Kurakin, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel. 2018. Ensemble Adversarial Training: Attacks and Defenses. 2018 *ICLR* arXiv preprint arXiv:1705.07204 (2018).
- [33] Huahua Wang and Arindam Banerjee. 2014. Bregman Alternating Direction Method of Multipliers. In *Advances in Neural Information Processing Systems* 27, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 2816–2824. <http://papers.nips.cc/paper/5612-bregman-alternating-direction-method-of-multipliers.pdf>
- [34] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, and Alan Yuille. 2017. Mitigating adversarial effects through randomization. *arXiv preprint arXiv:1711.01991* (2017).