

# On the Design of Black-box Adversarial Examples by Leveraging Gradient-free Optimization and Operator Splitting Method

Pu Zhao<sup>1</sup>, Sijia Liu<sup>2</sup>, Pin-Yu Chen<sup>2</sup>, Trong Nghia Hoang<sup>2</sup>, Kaidi Xu<sup>1</sup>, Bhavya Kailkhura<sup>3</sup>, Xue Lin<sup>1</sup>

<sup>1</sup>Department of Electrical and Computer Engineering, Northeastern University

<sup>2</sup>MIT-IBM Watson AI Lab, IBM Research

<sup>3</sup>Syracuse University

zhao.pu@husky.neu.edu, {sijia.liu, pin-yu.chen, nghiaht}@ibm.com, xu.kaid@husky.neu.edu, bkailkhu@syr.edu, xue.lin@northeastern.edu

## Abstract

*Robust machine learning is currently one of the most prominent topics which could potentially help shaping a future of advanced AI platforms that not only perform well in average cases but also in worst cases or adverse situations. Despite the long-term vision, however, existing studies on black-box adversarial attacks are still restricted to very specific settings of threat models (e.g., single distortion metric and restrictive assumption on target model's feedback to queries) and/or suffer from prohibitively high query complexity. To push for further advances in this field, we introduce a general framework based on an operator splitting method, the alternating direction method of multipliers (ADMM) to devise efficient, robust black-box attacks that work with various distortion metrics and feedback settings without incurring high query complexity. Due to the black-box nature of the threat model, the proposed ADMM solution framework is integrated with zeroth-order (ZO) optimization and Bayesian optimization (BO), and thus is applicable to the gradient-free regime. This results in two new black-box adversarial attack generation methods, ZO-ADMM and BO-ADMM. Our empirical evaluations on image classification datasets show that our proposed approaches have much lower function query complexities compared to state-of-the-art attack methods, but achieve very competitive attack success rates. Codes are available at [https://github.com/LinLabNEU/Blackbox\\_ADMM](https://github.com/LinLabNEU/Blackbox_ADMM).*

## 1. Introduction

In recent years, deep neural networks (DNNs) have achieved significant breakthroughs [22] in many machine learning (ML) tasks. However, despite these success stories, there have been many recent studies showing that even state-of-the-art DNNs might still be vulnerable to adversar-

ial misclassification attacks [15, 38, 44]. The adversarial attacks find and add visually imperceptible noises to an originally correctly classified input and essentially cause it to be misclassified by the DNNs. This raises security concerns about the robustness of DNNs in extreme situations with high reliability and dependability requirement such as face recognition, autonomous driving car and malware detection [34, 13, 17]. Investigating adversarial examples has become an increasingly prevailing topic to develop potential defensive measures in trustworthy ML [24, 40, 42, 41]. It essentially lays the groundwork for building a new generation of highly robust and reliable ML models acting as the core engine of future AI technology.

However, most of preliminary studies on this topic are restricted to the white-box setting where the adversary has complete access and knowledge of the target system (e.g., DNNs) [15, 21, 6, 7, 43]. Despite the theoretical interest, white-box attack methods are not adapted to practical black-box threat models. It is often the case that internal states/configurations and operating mechanism of public ML systems are not revealed to the practitioners (e.g., Google Cloud Vision API). Accordingly, the only mode of interaction with the system is via submitting inputs and receiving the corresponding predicted outputs.

To boost the practicality of such approaches, a few recent works have introduced a new class of threat models that exploit either a surrogate of the target model [31] or a gradient-free attack method [8, 1]. However, adversarial attacks that exploit a surrogate of the target model tend to yield low success rate if the surrogate is inaccurate. On the other hand, while attacks that use zeroth-order gradient estimation [8] are often more effective, they require a large number of queries to obtain an accurate estimate. Thus they are usually not economically efficient, especially in query-limited settings due to budget constraints.

To mitigate the above limitations of the existing literature,

this paper introduces a new perspective in the design of black-box adversarial attacks: We propose a general attack framework based on an operator splitting method, the alternating direction method of multipliers (ADMM), which integrates with both zeroth-order (ZO) optimization and Bayesian optimization (BO). Furthermore, unlike previous works which for ease of optimization often assume a specific distortion metric between an input and its perturbed version, our proposed framework is amenable to a broad family of distortion metrics including those previously used in the literature.

**Our Contributions:**

- We propose a general black-box adversarial attack framework via ADMM, including zeroth-order ADMM (ZO-ADMM) and ADMM with Bayesian optimization (BO-ADMM). We exploit a promising ZO-ADMM with random gradient estimation (RGE) [26] to design efficient black-box attacks that generalize the previous ZO coordinate descent based black-box attacks [8] and sidestep the notoriously intensive query complexity of attacks based on coordinate-wise random gradient estimation. Besides, we integrate the ADMM with BO for higher query efficiency in black-box settings (Section 4).
- We further generalize our formulation to accommodate various bounded  $\ell_p$ -norm-ball distortion metrics and their linear spans in the metric space (see Section 4.1). Such an extension is highly non-trivial to be incorporated into existing formulations of other black-box attacks, which are often heavily customized towards a specific norm-ball (e.g.,  $\ell_2$  or  $\ell_\infty$ ) for distortion metrics.
- Our framework is also made flexible to robustly accommodate for various threat models of the black-box attack (Section 5), which includes both score-based and decision-based settings. The former allows the attacker to have access to a vector of assessment scores for all output candidates (soft labels). And the latter only provides the system’s final decision on the most probable output (hard labels).
- Finally, we demonstrate the efficiency of our proposed framework on a variety of real-world image classification datasets such as MNIST, CIFAR-10 and ImageNet. The empirical results consistently show that our framework perform competitively to existing works in terms of the attack success rate while achieving a significant reduction on the query complexity (Section 6).

**2. Related Works**

The vulnerability of DNNs was first studied in the seminal works [2, 38], which were followed by a series of white-box threat models [15, 6, 28, 45, 46] that assume full access of the target model’s internal parameters/configurations. However, such internal knowledge of the target model is often not revealed and the adversary can only interact with it via submitting input queries and receiving feedback on potential outputs. Therefore, in the remaining of this section, we will

summarize recent advances on black-box adversarial attacks and discuss their limitations in comparison to our proposed framework.

**2.1. Black-box Attack with Surrogate Model**

A black-box attack using surrogate model is essentially a transfer attack [31] in which the adversary trains a DNN with data labeled by the target model. The resulting DNN is then exploited as a surrogate of the target model for which we can apply any state-of-the-art white-box attacks without requiring full access to internal states and operating mechanisms of the target model. Such attacks however depend heavily on the quality of training a surrogate model that closely resembles the true target model [27]. As a result, transfer attack tends to yield low success rate in data-intensive domains (e.g., ImageNet) for which it is hard to find a qualified surrogate.

**2.2. Black-box Attacks with Gradient Estimation**

Another approach to explore black-box attacks is to use gradient estimation via zeroth-order optimization (ZOO) [8]. They make queries to the model and estimate the output gradients with respect to the corresponding inputs, and then apply the state-of-the-art C&W attack method [6] to generate adversarial examples. However, this method is very computationally intensive as it requires a large number of queries per iteration to generate an accurate gradient estimation. Alternatively, the work [29] aims to estimate output gradient via greedy local search. At each iteration, the proposed technique perturbs only a small subset of input component. Such local search technique is very computationally efficient but it does not explicitly minimize the distortion between the original input and its perturbed version, the crafted noises often appear more visible. The work [18] investigates the more realistic threat models by defining the query-limited setting, the partial information setting, and the label-only setting. Three attacks methods are proposed based on the Natural Evolutionary Strategies and Monte Carlo approximation. But it only puts limits on the  $\ell_\infty$  norm instead of minimizing a certain  $\ell_p$  norm. Based on [18], the work [19] further investigates to utilize the prior information including the time-dependent priors (i.e., successive gradients are in fact heavily correlated) and the data-dependent priors (i.e., images tend to exhibit a spatially local similarity ) for higher query efficiency.

**2.3. Other Black-box Attacks**

In addition to the aforementioned works, there are also other black-box attacks [5, 39, 9, 16] under different practical settings, which are explored very recently. Among those, the notable boundary method [5] implements a decision-based attack, which starts from a very large adversarial perturbation (thus causing an immediate misclassification) and

tries to reduce the perturbation (i.e., minimize the distortion) through a random walk while remaining adversarial via staying on the boundary between the misclassified class and the true class. However, it suffers from high computational complexity due to a huge number of queries needed to decrease the distortion and it also has no guarantee on the convergence. Different from [5], the work [9] formulates the hard-label black-box attack as a real-valued optimization problem which is usually continuous and can be solved by the zeroth-order optimization algorithm. Similarly, [16] addresses the problem of finding a universal (image-agnostic) perturbation in the hard-label black-box setting.

In this paper, we will instead introduce an interesting reformulation of adversarial black-box attack based on ADMM, including ZO-ADMM [26] that enjoys the operator splitting advantage of ADMM and BO-ADMM that reduces the query complexity with the aid of Gaussian process.

### 3. Problem Formulation

In this work, we focus on adversarial attacks in the application of image classification with DNNs. In what follows, we first provide a general problem formulation for adversarial attack which is amenable to either white-box or black-box settings. Then, we will develop an efficient solution to the more interesting black-box setting where the adversary only has access to certain types of output of the DNN model (its internal structures and configurations are unknown to the adversary). Specifically, given a legitimate image  $\mathbf{x}_0 \in \mathbb{R}^d$  with its correct class label  $t_0$ , we aim to design an optimal adversarial perturbation  $\delta \in \mathbb{R}^d$  so that the perturbed example  $(\mathbf{x}_0 + \delta)$  is misclassified to target class  $t \neq t_0$  by the DNN model trained on legitimate images. The adversarial perturbation  $\delta$  can be obtained by solving the problem of the generic form,

$$\begin{aligned} & \underset{\delta}{\text{minimize}} && f(\mathbf{x}_0 + \delta, t) + \gamma D(\delta) \\ & \text{subject to} && (\mathbf{x}_0 + \delta) \in [0, 1]^d, \|\delta\|_\infty \leq \epsilon, \end{aligned} \quad (1)$$

where  $f(\mathbf{x}, t)$  denotes an attack loss incurred by misclassifying  $(\mathbf{x}_0 + \delta)$  to target class  $t$ ,  $D(\delta)$  is a distortion function that controls perceptual similarity between a legitimate image and an adversarial example, and  $\|\cdot\|_\infty$  signifies the  $\ell_\infty$  norm. In problem (1), the ‘hard’ constraints ensure that the perturbed noise  $\delta$  at each pixel (normalized to  $[0, 1]$ ) is imperceptible up to a predefined  $\epsilon$ -tolerant threshold, and the non-negative parameter  $\gamma$  places emphasis on the distortion. Furthermore, in the above problem, we mainly set  $D(\delta) = \|\delta\|_2^2$ , which is motivated by the superior performance of the outstanding C&W  $\ell_2$  adversarial attack. We highlight that  $D(\delta)$  can take other forms of  $\ell_p$  norms as discussed in Section 4.1.

The problem (1) is the general form of the problem in [18, 19] which does not consider the  $D(\delta)$  term. The advantage is that we are able to minimize the  $\ell_p$  distortion

after the adversarial perturbation is obtained, thus keeping the perturbation imperceptible. More specifically, if  $\epsilon$  is too small, we may not be able to obtain a successful adversarial example. Thus, we need to increase  $\epsilon$  to achieve a successful adversarial attack. But since  $\epsilon$  only limits the largest element of the perturbation, the whole perturbation over the image might be relatively large and easy to be recognized in case of large  $\epsilon$ . Thus, the  $D(\delta)$  term in problem (1) helps to minimize the  $\ell_p$  distortion of the whole perturbation, keeping it unnoticeable.

In the remaining of this section, we will discuss possible choices for the loss function  $f(\mathbf{x}, t)$ . Note that, without loss of generality, we only focus on targeted attack with designated target class  $t$  to mislead the DNN since the untargeted attack version can be easily implemented similar to the targeted attack [6]. We also emphasize that in the black-box setting, the gradients of  $f(\mathbf{x}, t)$  can not be obtained directly as it does in the white-box setting. The form of the loss function  $f(\mathbf{x}, t)$  depends on the constrained information in different black-box feedback settings. In particular, the definition of score-based (Section 3.1) and decision-based (Section 3.2) attacks as well as their loss functions will be discussed in the following subsections.

#### 3.1. Score-based Attack

In the score-based attack setting, the adversaries are able to make queries to DNN to obtain the soft labels (i.e., scores or probabilities of an image belonging to different classes), while information on gradients are not available. The loss function of problem (1) in the score-based attack is:

$$f(\mathbf{x}_0 + \delta, t) = \max\{\max_{j \neq t} \{\log P(\mathbf{x}_0 + \delta)_j\} - \log P(\mathbf{x}_0 + \delta)_t, -\kappa\}, \quad (2)$$

which is motivated by [6] and yields the best known performance among white-box attacks.  $P(\mathbf{x})_j$  denotes the target model’s prediction score or probability of the  $j$ -th class, and  $\kappa$  is a confidence parameter which is usually set to zero. Basically, this implies  $f(\mathbf{x}_0 + \delta, t) = 0$  if  $P(\mathbf{x}_0 + \delta)_t$  is the largest among all classes, which means the perturbation  $\delta$  has successfully made the target model misclassified  $\mathbf{x}_0 + \delta$  to target class  $t$ . Otherwise, it will be larger than zero. Note that in Eqn. (2) the log probability  $\log P(\mathbf{x})$  is used instead of directly using the actual probability  $P(\mathbf{x})$ . This is based on the observation that the output probability distribution tends to have one dominating class, making the query on the probability/score less effective. The utilization of the log operator can help to reduce the effect of the dominating class while it preserves the probability order for all classes.

#### 3.2. Decision-based Attack

Different from the score-based attack, the decision-based attack is more challenging in that the adversaries can only

make queries to get the hard-labels instead of the soft-labels. Let  $H(\mathbf{x})_i$  denote the hard-label decision.  $H(\mathbf{x})_i = 1$  if the decision for  $\mathbf{x}$  is label  $i$ , and 0 otherwise. We also have  $\sum_{i=1}^K H(\mathbf{x})_i = 1$  for all  $K$  classes. Then the loss function of problem (1) in the decision-based attack is specified as

$$f(\mathbf{x}_0 + \boldsymbol{\delta}, t) = \max_{j \neq t} H(\mathbf{x}_0 + \boldsymbol{\delta})_j - H(\mathbf{x}_0 + \boldsymbol{\delta})_t, \quad (3)$$

Therefore,  $f(\mathbf{x}_0 + \boldsymbol{\delta}, t) \in \{-1, 1\}$ , and the attacker succeeds if  $f(\mathbf{x}_0 + \boldsymbol{\delta}, t) = -1$ . The loss function (3) is *nonsmooth* with *discrete outputs*. The decision-based attack is therefore more challenging because existing combinatorial optimization methods become almost ineffective or inapplicable.

## 4. A General Black-box Adversarial Attack Framework

This section develops a general black-box adversarial attack framework for both the score-based and decision-based attacks by leveraging ADMM and gradient-free optimization. We will show that the proposed attack framework yields the following benefits: a) an efficient splitting between the black-box loss function and the adversarial distortion function, b) generalization to various  $\ell_p$  norm involved hard/soft constraints, and c) compatibility to different gradient-free operations. By introducing an auxiliary variable  $\mathbf{z}$ , problem (1) can be rewritten in the favor of ADMM-type methods [4, 47],

$$\begin{aligned} & \underset{\boldsymbol{\delta}, \mathbf{z}}{\text{minimize}} && f(\mathbf{x}_0 + \boldsymbol{\delta}, t) + \gamma D(\mathbf{z}) + \mathcal{I}(\mathbf{z}) \\ & \text{subject to} && \mathbf{z} = \boldsymbol{\delta}, \end{aligned} \quad (4)$$

where  $\mathcal{I}(\mathbf{z})$  is the indicator function given by,

$$\mathcal{I}(\mathbf{z}) = \begin{cases} 0 & (\mathbf{x}_0 + \mathbf{z}) \in [0, 1]^d, \|\mathbf{z}\|_\infty \leq \epsilon, \\ \infty & \text{otherwise.} \end{cases} \quad (5)$$

The augmented Lagrangian of the reformulated problem (4) is given by

$$\begin{aligned} \mathcal{L}(\mathbf{z}, \boldsymbol{\delta}, \mathbf{u}) &= \gamma D(\mathbf{z}) + \mathcal{I}(\mathbf{z}) + f(\mathbf{x}_0 + \boldsymbol{\delta}, t) \\ &+ \mathbf{u}^T (\mathbf{z} - \boldsymbol{\delta}) + \frac{\rho}{2} \|\mathbf{z} - \boldsymbol{\delta}\|_2^2, \end{aligned} \quad (6)$$

where  $\mathbf{u}$  is Lagrangian multiplier, and  $\rho > 0$  is a given penalty parameter. It can be further transformed as below,

$$\begin{aligned} \mathcal{L}(\mathbf{z}, \boldsymbol{\delta}, \mathbf{u}) &= \gamma D(\mathbf{z}) + \mathcal{I}(\mathbf{z}) + f(\mathbf{x}_0 + \boldsymbol{\delta}, t) \\ &+ \frac{\rho}{2} \left\| \mathbf{z} - \boldsymbol{\delta} + \frac{1}{\rho} \mathbf{u} \right\|_2^2 - \frac{1}{2\rho} \|\mathbf{u}\|_2^2. \end{aligned} \quad (7)$$

The ADMM algorithm [4] splits optimization variables into *two* blocks and adopts the following iterative scheme,

$$\mathbf{z}^{k+1} = \arg \min_{\mathbf{z}} \mathcal{L}(\mathbf{z}, \boldsymbol{\delta}^k, \mathbf{u}^k), \quad (8)$$

$$\boldsymbol{\delta}^{k+1} = \arg \min_{\boldsymbol{\delta}} \mathcal{L}(\mathbf{z}^{k+1}, \boldsymbol{\delta}, \mathbf{u}^k), \quad (9)$$

$$\mathbf{u}^{k+1} = \mathbf{u}^k + \rho(\mathbf{z}^{k+1} - \boldsymbol{\delta}^{k+1}), \quad (10)$$

where  $k$  denotes the iteration index. In problem (8), we minimize  $\mathcal{L}(\mathbf{z}, \boldsymbol{\delta}, \mathbf{u})$  over  $\mathbf{z}$  given parameters  $\boldsymbol{\delta}^k$  and  $\mathbf{u}^k$ . In problem (9), we minimize  $\mathcal{L}(\mathbf{z}, \boldsymbol{\delta}, \mathbf{u})$  over  $\boldsymbol{\delta}$  given  $\mathbf{z}^{k+1}$  from the previous step and  $\mathbf{u}^k$ . Then, the Lagrangian multiplier  $\mathbf{u}$  is updated in Eqn. (10). The major advantage of this ADMM-type algorithm is that it allows us to split the original complex problem into sub-problems, each of which can be solved more efficiently or even analytically. In what follows, we solve problems (8) and (9) respectively.

### 4.1. $\mathbf{z}$ -step

Problem (8) can be rewritten as

$$\begin{aligned} & \underset{\mathbf{z}}{\text{minimize}} && D(\mathbf{z}) + \frac{\rho}{2\gamma} \|\mathbf{z} - \mathbf{a}\|_2^2 \\ & \text{subject to} && (\mathbf{x}_0 + \mathbf{z}) \in [0, 1]^d, \|\mathbf{z}\|_\infty \leq \epsilon, \end{aligned} \quad (11)$$

where  $\mathbf{a} = \boldsymbol{\delta}^k - (1/\rho)\mathbf{u}^k$ . We set  $D(\mathbf{z}) = \|\mathbf{z}\|_2^2$  [6]. Problem (11) can be decomposed elementwise as below,

$$\begin{aligned} & \underset{z_i}{\text{minimize}} && \left( z_i - \frac{\rho}{2\gamma + \rho} a_i \right)^2 \\ & \text{subject to} && ([\mathbf{x}_0]_i + z_i) \in [0, 1], |z_i| \leq \epsilon, \end{aligned} \quad (12)$$

where  $[\mathbf{x}]_i$  (or  $x_i$ ) denotes the  $i$ -th element of  $\mathbf{x}$ . The solution to problem (12) is then given by

$$[\mathbf{z}^{k+1}]_i = \begin{cases} \min\{1 - [\mathbf{x}_0]_i, \epsilon\} & \frac{\rho}{2\gamma + \rho} a_i > \min\{1 - [\mathbf{x}_0]_i, \epsilon\} \\ \max\{-[\mathbf{x}_0]_i, -\epsilon\} & \frac{\rho}{2\gamma + \rho} a_i < \max\{-[\mathbf{x}_0]_i, -\epsilon\} \\ \frac{\rho}{2\gamma + \rho} a_i & \text{otherwise.} \end{cases} \quad (13)$$

**Generalization to various  $\ell_p$  norms.** In problem (11), in addition to the worst-case perturbation constraint  $\|\mathbf{z}\|_\infty \leq \epsilon$ , it is a common practice to set  $D(\mathbf{z}) = \|\mathbf{z}\|_2^2$  to measure the similarity between the legitimate image and the adversarial example. If  $D(\mathbf{z})$  takes other  $\ell_p$  norms such as  $\|\mathbf{z}\|_0$ ,  $\|\mathbf{z}\|_1$ , or even  $\ell_p$  norm combinations like  $\|\mathbf{z}\|_1 + \frac{\beta}{2} \|\mathbf{z}\|_2^2$  for  $\beta \geq 0$ , we are still able to obtain the solutions with minor modifications in the  $\mathbf{z}$ -step. This ability is highly non-trivial for other black-box attacks, which are often heavily customized to minimize a specific  $\ell_p$  norm for distortion measure. Due to space limitation, we explicitly show the  $\mathbf{z}$ -step solutions of  $D(\mathbf{z}) = \|\mathbf{z}\|_0$ ,  $D(\mathbf{z}) = \|\mathbf{z}\|_1$  and  $D(\mathbf{z}) = \|\mathbf{z}\|_1 + \frac{\beta}{2} \|\mathbf{z}\|_2^2$  derived with proximal operators [32] in the supplementary material.

### 4.2. $\boldsymbol{\delta}$ -step

Problem (9) can be written as

$$\underset{\boldsymbol{\delta}}{\text{minimize}} \quad f(\mathbf{x}_0 + \boldsymbol{\delta}, t) + \frac{\rho}{2} \|\boldsymbol{\delta} - \mathbf{b}\|_2^2, \quad (14)$$

where  $\mathbf{b} = \mathbf{z}^{k+1} + (1/\rho)\mathbf{u}^k$ . In the white-box setting, since the gradients of  $f(\mathbf{x}_0 + \boldsymbol{\delta}, t)$  are directly accessible, gradient descent method like stochastic gradient descent (SGD) or

Adam can be applied straight-forwardly. However, in black-box settings, the gradients of  $f(\mathbf{x}_0 + \boldsymbol{\delta}, t)$  are unavailable. Thus, to overcome this difficulty, we adopt two derivative-free methods: the random gradient estimation (RGE) method [12] and the Bayesian optimization [3] corresponding to ZO-ADMM and BO-ADMM, respectively.

#### 4.2.1 Random gradient estimation

In the black-box setting, the gradient of  $f(\mathbf{x}_0 + \boldsymbol{\delta}, t)$  is estimated through random gradient estimation (RGE),

$$\hat{\nabla} f(\boldsymbol{\delta}) = (d/(\nu Q)) \sum_{j=1}^Q [f(\boldsymbol{\delta} + \nu \mathbf{u}_j) - f(\boldsymbol{\delta})] \mathbf{u}_j, \quad (15)$$

where  $d$  is the number of optimization variables,  $\nu > 0$  is a smoothing parameter,  $\{\mathbf{u}_j\}$  denote independent and identically distributed (i.i.d.) random direction vectors drawn from a uniform distribution over a unit sphere, and  $Q$  is the number of random direction vectors. It has been shown in [26] that a large  $Q$  reduces the gradient estimation error and improves the convergence of ZO-ADMM. However, we find that a moderate size of  $Q$  is sufficient to provide a good trade-off between estimation error and query complexity, e.g.,  $Q = 20$  in our experiments. We also highlight that the RGE in (15) only requires  $O(Q)$  query complexity instead of  $O(dQ)$  caused by coordinate-wise gradient estimation used in [8]. Note that the natural evolutionary strategy (NES) [18] uses a central difference based gradient estimator requiring  $2Q$  queries. By contrast, RGE uses a forward difference based random gradient estimator, yielding  $Q + 1$  query counts, leading to higher query efficiency.

With the aid of RGE, the solution to problem (14) can now be obtained via stochastic gradient descent-like methods. However, it suffers from extremely high iteration and function query complexity due to the non-linearity of  $f$  as well as the iterative nature of ADMM. To sidestep this computational bottleneck, we propose the use of the linearized ADMM algorithm [36] in ZO-ADMM with RGE, and thus it enjoys dual advantages of gradient-free operation and linearization of the loss function. By linearization, the loss function  $f(\mathbf{x}_0 + \boldsymbol{\delta}, t)$  in problem (14) is replaced with its first-order Taylor expansion plus a regularization term (known as Bregman divergence), that is,  $\hat{\nabla} f(\boldsymbol{\delta}^k + \mathbf{x}_0, t)^T (\boldsymbol{\delta} - \boldsymbol{\delta}^k) + \frac{1}{2} \|\boldsymbol{\delta} - \boldsymbol{\delta}^k\|_{\mathbf{G}}^2$ , where  $\mathbf{G}$  is a pre-defined positive definite matrix, and  $\|\mathbf{x}\|_{\mathbf{G}}^2 = \mathbf{x}^T \mathbf{G} \mathbf{x}$ . We choose  $\mathbf{G} = \eta_k \mathbf{I}$  where  $1/\eta_k > 0$  is a decaying parameter, e.g.,  $\eta_k = \alpha \sqrt{k}$  for a given constant  $\alpha > 0$ . The Bregman divergence term is used to stabilize the convergence of  $\boldsymbol{\delta}$ .

Combining linearization and RGE, problem (14) now takes the following form:

$$\underset{\boldsymbol{\delta}}{\text{minimize}} \quad \begin{aligned} & (\hat{\nabla} f(\boldsymbol{\delta}^k + \mathbf{x}_0, t))^T (\boldsymbol{\delta} - \boldsymbol{\delta}^k) \\ & + \frac{\eta_k}{2} \|\boldsymbol{\delta} - \boldsymbol{\delta}^k\|_2^2 + \frac{\rho}{2} \|\boldsymbol{\delta} - \mathbf{b}\|_2^2, \end{aligned} \quad (16)$$

which yields a quadratic programming problem with the following closed-form solution:

$$\boldsymbol{\delta}^{k+1} = (1/(\eta_k + \rho)) \left( \eta_k \boldsymbol{\delta}^k + \rho \mathbf{b} - \hat{\nabla} f(\boldsymbol{\delta}^k + \mathbf{x}_0, t) \right). \quad (17)$$

Note that Eqn. (17) can be calculated with only one step of gradient estimation, which is a significant improvement on query efficiency compared with solving problem (14) using gradient descent method with thousands of random estimations. The convergence of the linearized ADMM for non-convex problems is proved in [25].

#### 4.2.2 Bayesian Optimization

In addition to RGE, BO is an alternative approach to solve problem (14). We model  $l(\boldsymbol{\delta}) := f(\mathbf{x}_0 + \boldsymbol{\delta}, t) + \frac{\rho}{2} \|\boldsymbol{\delta} - \mathbf{b}\|_2^2$  as a Gaussian process with a *prior* distribution  $l(\cdot) \sim \mathcal{N}(\mu_0, \kappa(\cdot, \cdot))$ , where  $\mu_0 = 0$  in practice and  $\kappa(\cdot, \cdot)$  is a positive definite kernel [33]. Consider a finite collection of *noisy observations*  $\mathcal{D}_k = \{y_1, \dots, y_k\}$ , where  $y_i \sim \mathcal{N}(l(\boldsymbol{\delta}^i), \sigma_n^2)$ , and  $\sigma_n^2$  is the noise variance. The *posterior probability* of a new function  $l(\boldsymbol{\delta})$  evaluation given  $\mathcal{D}_k$  is a Gaussian distribution with mean  $\mu$  and variance  $\sigma$ , that is  $l(\boldsymbol{\delta})|\mathcal{D}_k \sim \mathcal{N}(\mu, \sigma^2)$ , where

$$\mu = \boldsymbol{\kappa}^T [\mathbf{K} + \sigma_n^2 \mathbf{I}]^{-1} \mathbf{y}, \quad (18)$$

$$\sigma^2 = \kappa(\boldsymbol{\delta}, \boldsymbol{\delta}) - \boldsymbol{\kappa}^T [\mathbf{K} + \sigma_n^2 \mathbf{I}]^{-1} \boldsymbol{\kappa}, \quad (19)$$

$K_{ij} = \kappa(\boldsymbol{\delta}^i, \boldsymbol{\delta}^j)$ ,  $\boldsymbol{\kappa}$  is a vector of covariance terms between  $\{\boldsymbol{\delta}^i\}_{i=1}^k$  and  $\boldsymbol{\delta}$ , namely,  $\kappa_i = \kappa(\boldsymbol{\delta}^i, \boldsymbol{\delta})$ .

We choose the kernel function  $\kappa(\cdot, \cdot)$  as the ARD Matérn 5/2 kernel [35, 33],

$$\kappa(\mathbf{x}, \mathbf{y}) = \theta_0^2 \exp(-\sqrt{5}r) (1 + \sqrt{5}r + \frac{5}{3}r^2), \quad (20)$$

$$r^2 = \sum_{i=1}^d (x_i - y_i)^2 / \theta_i^2, \quad (21)$$

where  $\{\theta_i\}_{i=0}^d$  are hyperparameters. Note that  $\kappa(\boldsymbol{\delta}, \boldsymbol{\delta}) = \theta_0^2$ .

To determine the hyper-parameters  $\boldsymbol{\theta} = \{\{\theta_i\}_{i=0}^d, \sigma_n^2\}$ , we minimize the negative log marginal likelihood  $\log p(\mathcal{D}_k | \boldsymbol{\theta})$  [33],

$$\underset{\boldsymbol{\theta}}{\text{minimize}} \quad L(\boldsymbol{\theta}) \triangleq \frac{1}{2} \log |\mathbf{K} + \sigma_n^2 \mathbf{I}| \quad (22)$$

$$+ \frac{1}{2} \mathbf{y}^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y},$$

where  $\mathbf{y} = [y_1 \ y_2 \ \dots \ y_k]^T$ . This can be achieved by a standard gradient descent routine  $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \eta \partial L / \partial \boldsymbol{\theta}$  with a sufficiently small learning rate  $\eta$ .

In the setting of BO, the solution to problem (14) is often acquired by maximizing the *expected improvement* (EI). The

EI acquisition function is defined as [33]

$$\begin{aligned}
\delta^{k+1} &= \arg \max \text{EI}(\delta) \\
&= \arg \max \mathbb{E}_{l(\delta) | \mathcal{D}_k} [(l^+ - l(\delta)) \mathcal{I}(l(\delta) \leq l^+)], \\
&= \arg \max_{\delta} (l^+ - \mu) \Phi \left( \frac{l^+ - \mu}{\sigma} \right) + \sigma \phi \left( \frac{l^+ - \mu}{\sigma} \right),
\end{aligned} \tag{23}$$

where  $l^+$  denotes the best observed value, and  $\mathcal{I}(l(\delta) \leq l^+) = 1$  if  $l(\delta) \leq l^+$ , and 0 otherwise.  $\Phi$  and  $\phi$  denote the CDF and PDF of the standard normal distribution, respectively. We refer readers to the supplementary material for the detailed derivation of Eq. (23). We obtain  $\delta_{k+1}$  through the projected gradient descent method,

$$\begin{aligned}
\hat{\delta}^{(k+1)} &= \delta^{(k)} + \eta \nabla_{\delta=\delta^{(k)}} \text{EI}(\delta), \\
\delta^{(k+1)} &= \text{Proj}_{(\mathbf{x}_0 + \delta) \in [0, 1]^d, \|\delta\|_{\infty} \leq \epsilon} \left( \hat{\delta}^{(k+1)} \right).
\end{aligned} \tag{24}$$

The projection is introduced to ensure the feasibility of the next query point in BO.

## 5. Customized Score-based and Decision-based Black-box Attacks

For the score-based black-box attack, problem (1) with loss function (2) can be naturally solved through the general ADMM framework.

In the decision-based black-box attack, the form of the loss function (3) is non-smooth with discrete outputs. To overcome the discontinuity in Eqn. (3), a smoothing version of (3), denoted by  $f_{\mu}$  with smoothing parameter  $\mu > 0$  [14, 30], is taken into consideration,

$$f_{\mu}(\mathbf{x}_0 + \delta, t) = \mathbb{E}_{\mathbf{u} \in U_b} [f(\mathbf{x}_0 + \delta + \mu \mathbf{u}, t)], \tag{26}$$

where  $U_b$  is a uniform distribution within the unit Euclidean ball, or  $\mathbf{u}$  can follow a standard Gaussian distribution [18]. The rationale behind the smoothing technique in (26) is that the convolution of two functions, i.e.,  $\int_{\mathbf{u}} f(\mathbf{x}_0 + \delta + \mu \mathbf{u}, t) p(\mathbf{u}) d\mathbf{u}$ , is at least as smooth as the smoothest of the two original functions [11]. Therefore, when  $p$  is the density of a random variable with respect to Lebesgue measure, the loss function (26) is then smooth. In practice, we consider an empirical Monte Carlo approximation of (26)

$$f_{\mu}(\mathbf{x}_0 + \delta, t) \approx \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}_0 + \delta + \mu \mathbf{u}_i, t), \tag{27}$$

where  $\{\mathbf{u}_i\}$  are  $N$  i.i.d. samples drawn from  $U_b$ . With the smoothed loss function as expressed in Eqn. (27), problem (1) can be solved by the proposed general ADMM framework. To initialize ADMM, we initialize the perturbation  $\delta$  so that the initial perturbed image belongs to the target class, yielding a benefit of reducing query complexity compared to the initialization with an arbitrary image [5].

## 6. Performance Evaluation

In this section, the experimental results of the score-based and decision-based black-box attacks are demonstrated. We compare the proposed ADMM-based framework with various attack methods on three image classification datasets, MNIST [23], CIFAR-10 [20] and ImageNet [10]. The results of state-of-the-art white-box attack (i.e., C&W attack) are also provided for reference.

We train two networks for MNIST and CIFAR-10 datasets, respectively, which can achieve 99.5% accuracy on MNIST and 80% accuracy on CIFAR-10. The model architecture has four convolutional layers, two max pooling layers, two fully connected layers and a softmax layer. For ImageNet, we utilize a pre-trained Inception v3 network [37] instead of training our own model, which can achieve 96% top-5 accuracy. All experiments are conducted on machines with NVIDIA GTX 1080 TI GPUs.

### 6.1. Evaluation on MNIST and CIFAR-10

In the evaluation on MNIST and CIFAR-10, 200 correctly classified images are selected from MNIST and CIFAR-10 test datasets, respectively. For each image, the target labels are set to the other 9 classes and a total of 1800 attacks are performed for each attack method.

The implementations of C&W (white-box) attack [6] and ZOO (black-box) attack [8] are based on the GitHub code released by the authors<sup>1</sup>. For ZOO attack, we use ZOO-Adam with default Adam parameters. For the transfer attack [31], we apply C&W attack to the surrogate model with  $\kappa = 20$  to improve the attack transferability and 2,000 iterations in each binary search step. In the proposed ZO-ADMM attack<sup>1</sup>, the sampling number in random gradient estimation as defined in Eqn. (15),  $Q$ , is set to 20 and the sampling number for the decision-based smoothed loss function (27),  $N$ , is set to 10. We set  $\rho = 10$  and  $\gamma = 1$  for MNIST,  $\rho = 2000$  and  $\gamma = 10$  for CIFAR-10, and  $\rho = 1000$  and  $\gamma = 1$  for ImageNet.  $\epsilon$  is set to 1 for MNIST and CIFAR-10<sup>2</sup> and 0.05 for ImageNet. In Eq. (15), we set  $\nu = 0.5$  for three datasets. The parameter  $\mu$  in Eq. (27) is set to 1 for MNIST, 0.1 for CIFAR-10, and 0.01 for ImageNet.

The experimental results are shown in Table 1. Besides the attack success rate (ASR) and the  $\ell_p$  norms, we report the query number required to achieve the first successful attack, which characterizes how fast the generated adversarial perturbation can mislead DNNs. We observe that the transfer attack suffers from low ASR and large  $\ell_2$  distortion. Both the ZOO attack and the proposed ZO-ADMM attack with RGE can achieve high ASR and competitive  $\ell_2$  distortion close to the C&W white-box attack. Compared with the ZOO attack,


<sup>1</sup>Codes are available at [https://github.com/LinLabNEU/Blackbox\\_ADMM](https://github.com/LinLabNEU/Blackbox_ADMM).

<sup>2</sup>This setting is intended to make a fair comparison to the pure  $\ell_2$ -norm attack framework ZOO.

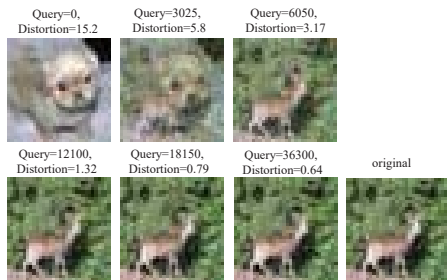
Table 1. Performance evaluation of adversarial attacks on MNIST and CIFAR-10.

Data set	Attack method		ASR	$\ell_1$ distortion	$\ell_2$ distortion	$\ell_\infty$ distortion	Query count on initial success	Reduction ratio on query count
MNIST	white -box	C&W white-box attack [6]	100%	22.14	1.962	0.5194	-	-
		Transfer attack (via C&W) [31]	30.6%	65.2	4.545	0.803	-	-
	score -based	ZOO attack [8]	98.8%	26.78	1.977	0.522	12,161	0.0 %
		score-based ZO-ADMM attack	98.3%	26.23	1.975	0.513	493.6	95.9%
		score-based BO-ADMM attack	87%	93.6	7.7	0.71	52.1	99.6%
	decision -based	boundary attack [5]	99%	32.9	2.21	0.563	25,328 <sup>a</sup>	0%
decision-based ZO-ADMM attack		100 %	30.48	2.166	0.548	7,603 <sup>a</sup>	62%	
CIFAR-10	white -box	C&W white-box attack [6]	100 %	11.7	0.332	0.0349	-	-
		Transfer attack (via C&W) [31]	8.5%	103.6	3.845	0.421	-	-
	score -based	ZOO attack [8]	97.6 %	15.2	0.361	0.0405	9982	0.0 %
		score-based ZO-ADMM attack	98.7 %	13.1	0.417	0.0392	421	95.7%
		score-based BO-ADMM attack	84.1%	148	5.29	0.62	46.3	99.6%
	decision -based	boundary attack [5]	100%	19.4	0.421	0.045	16,720 <sup>a</sup>	0%
decision-based ZO-ADMM attack		100%	17.25	0.415	0.0413	6,213 <sup>a</sup>	63%	

<sup>a</sup> As the decision-based attacks start from images in the target class, it achieves initial success immediately. Therefore, the query count on the initial success of the decision-based attack actually means the query number when it achieves the reported  $\ell_2$  distortion.



(a) An adversarial example evolution for MNIST starting from an image in the target class.



(b) An adversarial example evolution for CIFAR-10 starting from an image in the target class.

Figure 1. Adversarial examples generated by the proposed decision-based ZO-ADMM attack on MNIST and CIFAR-10.

the score-based ZO-ADMM attack requires fewer queries to obtain the first successful adversarial example. The query count in ZO-ADMM attack with RGE is reduced by 95.9% and 95.7% on MNIST and CIFAR-10, respectively. The reduction of query number is achieved by Eq. (17) in ZO-ADMM, which only requires one step of gradient estimation to solve the approximation problem instead of thousands of steps to solve the original problem. We also observe that the score-based BO-ADMM attack can achieve smaller query number compared with the RGE method, but it causes much larger  $\ell_p$  distortion. The reason is that BO-ADMM does not have very precise control for the perturbation. So it requires larger perturbation to mislead the DNN model. Although

BO-ADMM may have its limitations, we find that combining the advantage of BO- and ZO-ADMM can lead to more query-efficient attacks. Please refer to the BO-ZO-ADMM section in the appendix.

We notice that the decision-based ZO-ADMM attack achieves an  $\ell_2$  distortion slightly larger than the score-based attack with more queries as shown in Table 1. This is not surprising, since only the hard label outputs are available in the decision-based attack, which is more difficult to be optimized than the score-based attack. Although the  $\ell_2$  distortion is a bit larger, the perturbations are still visually indistinguishable. We compare the decision-based ZO-ADMM attack with the boundary attack [5]. As demonstrated in Table 1, the queries of the decision-based ZO-ADMM attack is about 60% less than that of the boundary attack to achieve the same level  $\ell_2$  distortion. We show the evolution of several adversarial examples in the decision-based attack versus the query number in Fig. 1. The decision-based attack starts from an image in the target class. Then it tries to decrease the  $\ell_2$  norm while keeping the classified label unchanged. After about 20,000 queries, the example is close to the original image with a satisfied  $\ell_2$  distortion.

## 6.2. Evaluation on ImageNet

We perform targeted and untargeted attacks in the score-based and decision-based settings on ImageNet. 100 correctly classified images are randomly selected. For each image in targeted attack, 9 random labels out of 1000 classes are selected to serve as the targets. We do not perform the transfer attack since it does not scale well to ImageNet due to training of the surrogate model. Instead, we provide the results of new baselines on ImageNet, including the query-limited attack as well as the label-only attack proposed in [18], and the bandit optimization based attack with time and data-dependent priors (named as Bandits<sub>TD</sub>) [19]. The query-limited and Bandits<sub>TD</sub> attacks are score-based attacks. The label-only attack is a decision-based attack.

Table 2. Performance evaluation of adversarial attacks on ImageNet.

Attack method		Untargeted attack			Targeted attack		
		ASR	Query count on initial success	Reduction ratio	ASR	Query count on initial success	Reduction ratio
score-based	C&W white-box attack [6]	100%	-	-	99%	-	-
	ZOO attack [8]	90%	15631	0.0%	78%	$2.11 \times 10^6$	0.0%
	Query-limited attack [18]	100%	4785	69.4%	98%	34128	98.4%
	Bandits <sub>TD</sub> attack [19]	94%	1259	92%	- <sup>a</sup>	-	-
	score-based ZO-ADMM attack	98%	891	94.3%	97%	16058	99.2%
decision-based	Label only [18]	- <sup>b</sup>	-	-	92%	$1.89 \times 10^6$ <sup>c</sup>	10.4%
	decision-based ZO-ADMM attack	100%	11742 <sup>c</sup>	24.9%	94%	$1.52 \times 10^6$ <sup>c</sup>	28%

<sup>a</sup> It mainly explores untargeted attack. <sup>b</sup> The label only attack mainly explores targeted attack. <sup>c</sup> The query count on initial success for the decision-based attack means the query number when it achieves the same  $\ell_2$  distortion with the ZOO attack on its initial success.

The experimental results are summarized in Table 2. For score-based attacks, we can observe that the score-based ZO-ADMM attack can achieve a high ASR with fewer queries than the other attacks. It reduces the query number on initial success by 94.3% and 99.2% for untargeted and targeted attacks, respectively, compared with the ZOO attack. For decision-based attacks, the ZO-ADMM attack can obtain a high ASR with fewer queries compared with the label-only attack or even the ZOO attack using score-based information. Some adversarial examples generated by the ZO-ADMM attack are demonstrated in the supplementary material. More experimental results including the comparison with Auto-Zoom [39] and the boundary method [5] method are demonstrated in the Appendix.

### 6.3. Convergence of the ZO-ADMM Attack

In Fig. 2, we demonstrate the convergence of the proposed ZO-ADMM targeted black-box attack, where the average  $\ell_2$  distortion of 9 targeted adversarial examples versus the query number is presented. As we can see, since we initialize the adversarial distortion from zeros, the score-based ZO-ADMM attack increases  $\ell_2$  distortion until a successful adversarial example is found. After that, it tries to decrease the  $\ell_2$  distortion but keeps the target label unchanged. For the decision-based attack, Fig 2 shows that the  $\ell_2$  distortion is initially large as ZO-ADMM starts from an image in the target class instead of the original image. The resulting  $\ell_2$  distortion then decreases as the query number increases. We highlight that the ZO-ADMM attack is able to reach the successful attack with hundreds of queries on MNIST or CIFAR-10 and tens of thousands of queries on ImageNet, which significantly outperforms the ZOO attack. Besides Fig. 2 demonstrating the  $\ell_2$  distortion versus query number, we present the  $\ell_2$  distortion versus ADMM iteration number in the supplementary material and similar results can be drawn.

### 6.4. Evaluation for Various $\ell_p$ Norms

In the previous experiments, we mainly consider the case of  $D(\mathbf{z}) = \|\mathbf{z}\|_2^2$  for a fair comparison with other white-box and black-box algorithms. However, we highlight that the ZO-ADMM method is able to optimize various  $\ell_p$  norms, not only  $\ell_2$  norm. In Table 3, we present the experimental




Figure 2. Convergence of the ZO-ADMM attack.

Table 3. Performance evaluation of the ZO-ADMM attacks on MNIST for different  $\ell_p$  norms.

Attack method	ASR	$\ell_0$	$\ell_1$	$\ell_2$
ZO-ADMM $\ell_0$	100%	<b>18.5</b>	12.6	9.72
ZO-ADMM $\ell_1$	100%	465	<b>10.5</b>	2.71
ZO-ADMM $\ell_2$	100%	483	22.09	<b>1.93</b>

results for different  $\ell_p$  norms when solving problem (11). Here we focus on three score-based black-box attacks with ZO-ADMM by minimizing the  $\ell_0$ ,  $\ell_1$  and  $\ell_2$  distortion, respectively. As we can see, our proposed method is well adapted to different  $\ell_p$  norms in the design of black-box adversarial examples.

## 7. Conclusion

In this paper, we propose a general framework to design norm-ball bounded black-box adversarial examples by leveraging an operator splitting method (namely, ADMM), together with the gradient-free operations including random gradient estimation and Bayesian optimization. The proposed framework can be applied to both score-based and decision-based settings. Compared to state-of-the-art black-box attacks, our approach achieves better query efficiency without losing the attack performance in terms of attack success rate as well as  $\ell_p$ -norm distortion.

## Acknowledgement

This work is partly supported by the National Science Foundation CNS-1932351.



## References

- [1] Arjun Nitin Bhagoji, Warren He, Bo Li, and Dawn Song. Exploring the space of black-box attacks on deep neural networks. *CoRR*, abs/1712.09491, 2017.
- [2] Battista Biggio, Igino Corona, Davide Maiorca, Blaine Nelson, Nedim Šrđić, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. Evasion attacks against machine learning at test time. In *Joint European conference on machine learning and knowledge discovery in databases*, pages 387–402, 2013.
- [3] Ilija Bogunovic, Jonathan Scarlett, Andreas Krause, and Volkan Cevher. Truncated variance reduction: A unified approach to bayesian optimization and level-set estimation. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 1507–1515. Curran Associates, Inc., 2016.
- [4] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.
- [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] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *Security and Privacy (SP), 2017 IEEE Symposium on*, pages 39–57. IEEE, 2017.
- [7] Pin-Yu Chen, Yash Sharma, Huan Zhang, Jinfeng Yi, and Cho-Jui Hsieh. Ead: elastic-net attacks to deep neural networks via adversarial examples. *arXiv preprint arXiv:1709.04114*, 2017.
- [8] Pin-Yu Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui Hsieh. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, pages 15–26. ACM, 2017.
- [9] Minhao Cheng, Thong Le, Pin-Yu Chen, Huan Zhang, Jinfeng Yi, and Cho-Jui Hsieh. Query-efficient hard-label black-box attack: An optimization-based approach. In *International Conference on Learning Representations*, 2019.
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. IEEE, 2009.
- [11] John C Duchi, Peter L Bartlett, and Martin J Wainwright. Randomized smoothing for stochastic optimization. *SIAM Journal on Optimization*, 22(2):674–701, 2012.
- [12] John C Duchi, Michael I Jordan, Martin J Wainwright, and Andre Wibisono. Optimal rates for zero-order convex optimization: The power of two function evaluations. *IEEE Transactions on Information Theory*, 61(5):2788–2806, 2015.
- [13] Ivan Evtimov, Kevin Eykholt, Earlene Fernandes, Tadayoshi Kohno, Bo Li, Atul Prakash, Amir Rahmati, and Dawn Song. Robust physical-world attacks on machine learning models. *arXiv preprint arXiv:1707.08945*, 2017.
- [14] Xiang Gao, Bo Jiang, and Shuzhong Zhang. On the information-adaptive variants of the admm: an iteration complexity perspective. *Journal of Scientific Computing*, 76(1):327–363, 2018.
- [15] Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *2015 ICLR*, arXiv preprint arXiv:1412.6572, 2015.
- [16] Thomas A Hogan and Bhavya Kailkhura. Universal hard-label black-box perturbations: Breaking security-through-obscurity defenses. *arXiv preprint arXiv:1811.03733*, 2018.
- [17] Weiwei Hu and Ying Tan. Black-box attacks against RNN based malware detection algorithms. *CoRR*, abs/1705.08131, 2017.
- [18] Andrew Ilyas, Logan Engstrom, Anish Athalye, and Jessy Lin. Black-box adversarial attacks with limited queries and information. In *Proceedings of the 35th International Conference on Machine Learning, ICML, July 2018*.
- [19] Andrew Ilyas, Logan Engstrom, and Aleksander Madry. Prior convictions: Black-box adversarial attacks with bandits and priors. *ICLR*, 2019.
- [20] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- [21] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533*, 2016.
- [22] Yann LeCun, Y Bengio, and Geoffrey Hinton. Deep learning. *521:436–44*, 05 2015.
- [23] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [24] Ji Lin, Chuang Gan, and Song Han. Defensive quantization: When efficiency meets robustness. *ICLR*, 2019.
- [25] Qinghua Liu, Xinyue Shen, and Yuantao Gu. Linearized admm for nonconvex nonsmooth optimization with convergence analysis. *IEEE Access*, 2019.
- [26] Sijia Liu, Jie Chen, Pin-Yu Chen, and Alfred O Hero. Zeroth-order online admm: Convergence analysis and applications. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84, pages 288–297, April 2018.
- [27] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into transferable adversarial examples and black-box attacks. *CoRR*, abs/1611.02770, 2016.
- [28] Seyed Mohsen Moosavi Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, number EPFL-CONF-218057, 2016.
- [29] Nina Narodytska and Shiva Prasad Kasiviswanathan. Simple black-box adversarial perturbations for deep networks. *CoRR*, abs/1612.06299, 2016.
- [30] Yurii Nesterov and Vladimir Spokoiny. Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 2(17):527–566, 2015.
- [31] Nicolas Papernot, Patrick D. McDaniel, and Ian J. Goodfellow. Transferability in machine learning: from phenom-

- ena to black-box attacks using adversarial samples. CoRR, abs/1605.07277, 2016.
- [32] Neal Parikh, Stephen Boyd, et al. Proximal algorithms. Foundations and Trends® in Optimization, 1(3):127–239, 2014.
- [33] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas. Taking the human out of the loop: A review of bayesian optimization. Proceedings of the IEEE, 104(1):148–175, 2016.
- [34] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K. Reiter. Adversarial generative nets: Neural network attacks on state-of-the-art face recognition. CoRR, abs/1801.00349, 2018.
- [35] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical bayesian optimization of machine learning algorithms. In Advances in neural information processing systems, pages 2951–2959, 2012.
- [36] Taiji Suzuki. Dual averaging and proximal gradient descent for online alternating direction multiplier method. In International Conference on Machine Learning, pages 392–400, 2013.
- [37] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 2818–2826, 2016.
- [38] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, and et al. Intriguing properties of neural networks. arXiv preprint arXiv:1312.6199, 2013.
- [39] Chun-Chen Tu, Paishun Ting, Pin-Yu Chen, Sijia Liu, Huan Zhang, Jinfeng Yi, Cho-Jui Hsieh, and Shin-Ming Cheng. Autozoom: Autoencoder-based zeroth order optimization method for attacking black-box neural networks. arXiv preprint arXiv:1805.11770, 2018.
- [40] Siyue Wang, Xiao Wang, Pu Zhao, Wujie Wen, David Kaeli, Peter Chin, and Xue Lin. Defensive dropout for hardening deep neural networks under adversarial attacks. In ICCAD '18, pages 71:1–71:8, New York, NY, USA, 2018. ACM.
- [41] Xiao Wang, Siyue Wang, Pin-Yu Chen, Yanzhi Wang, Brian Kulis, Xue Lin, and Sang Chin. Protecting neural networks with hierarchical random switching: Towards better robustness-accuracy trade-off for stochastic defenses. In Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI), pages 6013–6019, 2019.
- [42] Kaidi Xu, Hongge Chen, Sijia Liu, Pin-Yu Chen, Tsui-Wei Weng, Mingyi Hong, and Xue Lin. Topology attack and defense for graph neural networks: An optimization perspective. arXiv preprint arXiv:1906.04214, 2019.
- [43] Kaidi Xu, Sijia Liu, Gaoyuan Zhang, Mengshu Sun, Pu Zhao, Quanfu Fan, Chuang Gan, and Xue Lin. Interpreting adversarial examples by activation promotion and suppression. arXiv preprint arXiv:1904.02057, 2019.
- [44] Kaidi Xu, Sijia Liu, Pu Zhao, Pin-Yu Chen, Huan Zhang, Quanfu Fan, Deniz Erdogmus, Yanzhi Wang, and Xue Lin. Structured adversarial attack: Towards general implementation and better interpretability. In International Conference on Learning Representations, 2019.
- [45] Pu Zhao, Sijia Liu, Yanzhi Wang, and Xue Lin. An admm-based universal framework for adversarial attacks on deep neural networks. In ACM Multimedia 2018.
- [46] Pu Zhao, Siyue Wang, Cheng Gongye, Yanzhi Wang, Yungsi Fei, and Xue Lin. Fault sneaking attack: A stealthy framework for misleading deep neural networks. In Proceedings of the 56th Annual Design Automation Conference 2019, DAC '19, pages 165:1–165:6, New York, NY, USA, 2019. ACM.
- [47] Pu Zhao, Kaidi Xu, Sijia Liu, Yanzhi Wang, and Xue Lin. Admm attack: An enhanced adversarial attack for deep neural networks with undetectable distortions. In Proceedings of the 24th Asia and South Pacific Design Automation Conference, ASPDAC '19, pages 499–505, New York, NY, USA, 2019. ACM.