Research Article

Boosting Adversarial Attacks on Neural Networks with Better Optimizer

Heng Yin , Hengwei Zhang , Jindong Wang , and Ruiyu Dou

State Key Laboratory of Mathematical Engineering and Advanced Computing, Zhengzhou 450001, Henan, China

Correspondence should be addressed to Hengwei Zhang; wby_zzmy_henan@163.com

Received 5 March 2021; Revised 19 April 2021; Accepted 26 May 2021; Published 7 June 2021

Academic Editor: Ting Chen

Copyright © 2021 Heng Yin et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Convolutional neural networks have outperformed humans in image recognition tasks, but they remain vulnerable to attacks from adversarial examples. Since these data are crafted by adding imperceptible noise to normal images, their existence poses potential security threats to deep learning systems. Sophisticated adversarial examples with strong attack performance can also be used as a tool to evaluate the robustness of a model. However, the success rate of adversarial attacks can be further improved in black-box environments. Therefore, this study combines a modified Adam gradient descent algorithm with the iterative gradient-based attack method. The proposed Adam iterative fast gradient method is then used to improve the transferability of adversarial examples. Extensive experiments on ImageNet showed that the proposed method offers a higher attack success rate than existing iterative methods. By extending our method, we achieved a state-of-the-art attack success rate of 95.0% on defense models.

1. Introduction

In image recognition tasks, convolutional neural networks are able to classify images with an accuracy approaching that of humans [1–4]. However, researchers have found that neural networks are also vulnerable to adversarial examples. Szegedy et al. [5] first proposed the concept of adversarial examples: images added with small perturbations, which cause neural network models to output incorrect classifications with high confidence. These adversarial perturbations are often indistinguishable to the human eyes (in other words, there is no obvious visual difference between the adversarial examples and the original images).

Adversarial attacks can be categorized into white- and black-box attacks. A variety of techniques can be used to generate adversarial examples and perform white-box attacks, depending on the model structure and corresponding parameters [6–10]. In addition, adversarial examples are generally transferable as data generated for one model may be able to fool other models. This facilitates black-box attacks, in which the structure and parameters of the model are not available, for various neural networks [11]. Goodfellow et al. [6] suggested that different models learn similar

decision boundaries during the same image classification tasks and obtain similar parameters. These properties make it easier to generalize adversarial examples to different models.

Although adversarial examples are generally transferable, the optimal approach of improving this transferability needs to be examined further. Due to a balance between attack performance and transferability, basic iterative attacks are often more effective than single-step attacks in white-box environments and weaker in black-box environments. In white-box attacks, iterative methods excessively fit specific network parameters, achieving a high success rate but preventing generalization to other models [9]. We consider that this is the result of overfitting since attack performance for adversarial examples in white- and black-box environments is similar to the neural network performance on training and test sets.

Unlike white-box attacks, black-box attacks are more consistent with actual attack-defense environments and are primarily performed in one of the following three ways. (1) Decision-based attacks are conducted using information collected from the network. Although access to model structure and parameters is unavailable, the attacker can

generate adversarial examples by executing multiple queries against the model [12–15]. (2) In substitute-model techniques, the attacker can input images into the model to obtain an output label. Then, the substitute model can be trained to imitate the target model and generate adversarial examples [16]. (3) Transferability attacks are conducted by improving the transferability of adversarial examples generated in a white-box setting [8, 9]. The first two black-box attacks require large quantities of model queries, which is impractical in some cases (e.g., online platforms often limit the number of queries). As such, this study investigates black-box attacks using adversarial data with strong transferability. We investigate a state-of-the-art optimizer Adam for improving black-box adversarial attacks.

In this paper, we propose the Adam iterative fast gradient method (AI-FGM) to improve the transferability of adversarial examples among different models [17]. Inspired by the fact that Adam is better than momentum in optimization, we adopt Adam optimizer into the iterative gradient-based attack, so as to accelerate data update in dimensions with small gradients and achieve better convergence, by applying the second momentum term and a decreasing step size. As shown in Figure 1, being different from the momentum method that just accumulates the gradients of data points along the optimization path, our Adam method can also accumulate the square of the gradients. Such accumulation might help obtain an adaptive update direction, which leads to a better local minimum. Besides, a variable step size is also used to avoid oscillating.

Compared with existing gradient-based methods, the proposed method offers improved attack performance in black-box settings. This approach was tested on multiple networks, including adversarially trained networks, achieving higher attack success rates on black-box models. In addition, we combined our approach with advanced methods and attacked an ensemble of multiple networks, which further improved the transferability of adversarial examples.

2. Related Works

2.1. Adversarial Attack Methods. A variety of techniques have been proposed to generate transferable adversarial examples, which can be divided into two categories: optimizer-based methods and data-augmentation-based methods. Specifically, Goodfellow et al. proposed the fast gradient sign method (FGSM) to generate adversarial examples with very low calculation costs [6]. By extending FGSM, Alexey et al. developed the iterative fast gradient sign method (I-FGSM) [7]. Dong et al. proposed the momentum iterative fast gradient sign method (MI-FGSM) to improve the transferability of adversarial examples in black-box environments [8]. Lin et al. proposed Nesterov iterative fast gradient sign method (NI-FGSM) and adopted Nesterov accelerated gradient into the iterative attacks [18]. Xie first applied data augmentation in the generation of adversarial examples and proposed the diverse inputs method (DIM) [9]. Dong et al. proposed a translation-invariant attack (TI-FGSM) and improved the transferability of adversarial

FIGURE 1: Schematic optimization path of MI-FGSM and the proposed AI-FGM. MI-FGSM accumulates the gradients of data points along the optimization path, while Adam accumulates the gradients and the square of the gradients. MI-FGSM adopts an invariable step size, while AI-FGM adopts a decayed step size that leads to better convergence.

examples on defense models by a large margin [19]. Lin et al. proposed the scale-invariant attack method (SIM) and combined it with existing methods, resulting in SI-NI-TI-DIM (the currently strongest black-box attack method) [18].

2.2. Adversarial Defense Methods. Multiple defense mechanisms have been proposed to protect deep learning models from the threat of adversarial examples [20–29]. Among these, adversarial training is the most effective way to improve model robustness [6, 30]. In this process, adversarial examples are generated and added to the training set to participate in the model training procedure. Since normal adversarial training remains vulnerable to adversarial examples, Tramèr et al. proposed ensemble adversarial training to improve robustness further [31]. In this process, adversarial data generated by multiple models are included in the training set for a single model, thereby producing a more robust classifier.

3. Methodology

This section provides a detailed introduction to our proposed methodology. Let x denote an input image, and y denote the corresponding ground-truth label. The term θ represents network parameters, and $J(\theta, x, y)$ describes a loss function, typically the cross-entropy loss. The primary objective is to generate an adversarial example x^* to fool the model by maximizing $J(\theta, x, y)$, such that the prediction label $y^{\text{pre}} \neq y$. In this paper, we used an L_∞ norm bound to limit adversarial perturbations, such that $\|x^* - x\|_\infty \leq \epsilon$, where ϵ refers to the size of the perturbation. The generation of adversarial examples can, therefore, be converted into the following optimization problem:

$$\begin{aligned} & \arg \max_{x^*} J(\theta, x^*, y), \\ & \text{s.t. } \|x^* - x\|_{\infty} \leq \varepsilon. \end{aligned} \quad (1)$$

3.1. Attack Methods Based on Gradient. In this section, several techniques that are used to solve the above optimization problem are introduced briefly:

- (i) Fast gradient sign method (FGSM) [6]: As one of the simplest techniques, it seeks adversarial examples in the direction of the gradient of the loss function with respect to the input image x . The method can be expressed as

$$x^* = x + \varepsilon \cdot \text{sign}(\nabla_x J(\theta, x, y)). \quad (2)$$

- (ii) Iterative fast gradient sign method (I-FGSM) [7]: This algorithm is an iterative version of FGSM. The approach involves dividing the FGSM gradient operation into multiple steps that can be expressed as follows:

$$\begin{aligned} x_0^* &= x, \\ x_{i+1}^* &= x_i^* + \alpha \cdot \text{sign}(\nabla_x J(\theta, x_i^*, y)), \\ x_{i+1}^* &= \text{Clip}(x_{i+1}^*, x - \varepsilon, x + \varepsilon). \end{aligned} \quad (3)$$

where α denotes the step size of each iteration and $\alpha = (\varepsilon/T)$, in which T denotes the number of iterations. The Clip function was included to limit the adversarial example x^* within the ε neighborhood of the original image x and satisfy the L_{∞} norm constraint. I-FGSM is more effective than FGSM in white-box environments but less effective in black-box environments. In other words, I-FGSM exhibits poor transferability. This iterative attack method is also known as projected gradient descent (PGD) if the algorithm is added by a random initialization on x [20].

- (iii) Momentum iterative fast gradient sign method (MI-FGSM) [8]: In this method, a momentum term is applied to the iterative process to escape from poor local maxima. This produces adversarial examples with more transferability, which can be expressed as [8]

$$\begin{aligned} x_0^* &= x, \\ g_0 &= 0, \\ g_{i+1} &= \mu \cdot g_i + \frac{\nabla_x J(\theta, x_i^*, y)}{\|\nabla_x J(\theta, x_i^*, y)\|_1}, \\ x_{i+1}^* &= x_i^* + \alpha \cdot \text{sign}(g_{i+1}), \\ x_{i+1}^* &= \text{Clip}(x_{i+1}^*, x - \varepsilon, x + \varepsilon), \end{aligned} \quad (4)$$

where μ denotes the decay factor of the momentum term, and g_i denotes the weighted accumulation of gradients in the first i rounds of iterations.

3.2. Adam Iterative Fast Gradient Method. The generation of adversarial examples is similar to the training of neural networks; both processes can be viewed as an optimization problem. Specifically, the adversarial example can be viewed as the training set, and the black-box model can be viewed as a testing set. From this point of view, the transferability of the adversarial examples is similar to the generalization of the models. Therefore, we can apply the methods used to improve the generalization of the models to the generation of adversarial examples. Many methods are proposed to improve the generalization of neural network models, which can be divided into two categories: better optimization and data augmentation. Correspondingly, these two kinds of methods can be used in the adversarial attacks, and there have been attempts to do so, for example, MI-FGSM and DIM [8, 9]. Based on the analysis above, we aim to improve the transferability of adversarial examples with Adam optimizer since it performs well in the training of neural networks.

Adam [17] is an optimization algorithm combining momentum [32] and RMSProp [33]. In the iteration process, Adam accumulates not only the gradients of the loss function with respect to the input image but also the square of the gradients. This is done to accelerate loss function ascent in dimensions with small gradients. In addition, Adam uses a decreasing step size in each iteration to achieve better convergence.

In contrast, both I-FGSM and MI-FGSM adopt a constant step size. In the later stages of the algorithm, oscillations occur near the local maximum, which do not converge well. To solve this problem, we have improved the Adam algorithm by normalizing the step size sequence (defined in Algorithm 1) and have applied it to the iterative gradient method. The proposed Adam iterative fast gradient method (AI-FGM) is summarized in Algorithm 1.

Specifically, the gradient $\nabla_x J(\theta, x_t^*, y)$ in each iteration is normalized by its own L_1 distance, defined in Algorithm 1, because the scale of these gradients differs widely in each iteration [8]. Similar to MI-FGSM, m_t accumulates gradients of the first t iterations with a decay factor β_1 , defined in Algorithm 1. The result can be considered the first momentum. The term v_t represents the second momentum, which accumulates the squares of gradients for the first t iterations with a decay factor β_2 , defined in Algorithm 1. It is noteworthy that g_t^2 denotes an elementwise square $g_t \odot g_t$, where $\odot$ represents the Hadamard product [32]. The terms β_1 and β_2 are typically defined in the range (0, 1). The update direction of input x is defined in Algorithm 1, where the stability coefficient δ is set to avoid a zero in the denominator. The s_t term is advantageous as it prompts x to escape from local maxima and accelerates the updating of x in dimensions with small gradients.

Learning rate decay is often used in the training of neural networks. In this study, the step size used in the Adam optimizer is continually reduced to help improve the convergence of the algorithm. If β_1 and β_2 are set appropriately, the value of $\sqrt{1 - \beta_2^t} / (1 - \beta_1^t)$ will decrease with the increase of t . Thus, a decreasing sequence can be generated when t ranges from 1 to T . The sequence can then be normalized to obtain the weight of step sizes in each iteration, relative to the total step size α . This can be used to acquire a set of exponential decay step sizes controlled by β_1 and β_2 , as defined in Algorithm 1.

Existing techniques typically use the sign function to satisfy the L_∞ norm limitation. However, if the sign function was applied in Algorithm 1, the update direction in our method would be equivalent to that of MI-FGSM. Hence, we constrain adversarial examples within the L_∞ norm bound by the Clip function and apply the step size and update direction within the corresponding L_2 norm bound, defined in Algorithm 1. The relationship between the L_∞ norm bound (ϵ) and the L_2 norm bound (α) is defined in Algorithm 1, where N represents the dimension of the input image x .

3.3. Attacking an Ensemble of Networks. The proposed method was also used to attack an ensemble of networks. If an adversarial example poses a threat to multiple networks, it is far more likely to transfer to other models [11]. We followed the ensemble strategy proposed by Dong et al., in which multiple models are attacked by fusing network logits [8]. Specifically, in order to attack an ensemble of K models, logits were fused as follows:

$$l(x) = \sum_{k=1}^K \omega_k l_k(x), \quad (5)$$

where $l_k(x)$ refers to the logit output of the k -th model, ω_k refers to the ensemble weight ($\omega_k \geq 0$), and $\sum_{k=1}^K \omega_k = 1$.

4. Experiment

Extensive validation experiments were conducted to demonstrate the effectiveness of the proposed methodology. Experimental settings are provided in Section 4.1. The results of attacking a single network and some parameters that influence the results are discussed in Section 4.2. Thus, the results of attacking an ensemble of models are shown in Section 4.3. Finally, Section 4.4 exhibits the results of attacks by the combination of AI-FGM and existing methods (source code is available at https://github.com/YinHeng121/Adam_attack).

4.1. Experimental Setup

- (i) Data set: if the original images cannot be classified correctly by the network, it is meaningless to generate adversarial examples based on these data. Therefore, we selected 1,000 images belonging to the 1,000 categories from the ImageNet validation set,

all of which could be classified correctly by the tested networks. All images were scaled to $299 \times 299 \times 3$. As such, $N = 299 \times 299 \times 3$ in Section 3.2.

- (ii) Networks: seven different models were studied; four of which were normally trained networks (i.e., Inception-v3 (Inc-v3) [34], Inception-v4 (Inc-v4) [35], Inception-Resnet-v2 (IncRes-v2) [35], and Resnet-v2-101 (Res-101) [36]). The other three models were adversarially trained networks (i.e., ens3-adv-Inception-v3 (Inc-v3_{ens3}), ens4-adv-Inception-v3 (Inc-v3_{ens4}), and ens-adv-Inception-ResNet-v2 (IncRes-v2_{ens}) [31]).
- (iii) Other details: the proposed method was compared with the other methods discussed in Section 3.1. All the experiments described in this paper were based on the L_∞ norm. Thus, the momentum decay factor μ in Algorithm 1 was 1.0, and the stability coefficient δ in Algorithm 1 was 10^{-8} , as suggested in [17].

4.2. Attacking a Single Network. We first performed adversarial attacks on a single network with FGSM, I-FGSM, MI-FGSM, and AI-FGM. Moreover, the adversarial examples were generated on four normally trained networks and tested on all seven networks. The results are shown in Table 1, where the success rates are misclassification rates for the corresponding models, with adversarial examples used as input. The decay factors β_1 and β_2 in AI-FGM were set to 0.99 and 0.999, respectively. The maximum perturbation ϵ was 16. The number of iterations T for I-FGSM, PGD, MI-FGSM, and AI-FGM was 10. These methods are hereafter referred to as “iterative methods” without ambiguity. The effects of these parameter choices are discussed further in this section.

As shown in Table 1, all four iterative methods attacked a white-box model with a near 100% success rate. AI-FGM performed better than the other three methods on all black-box models. For example, for adversarial examples generated on Inc-v3, AI-FGM had a success rate of 60.7% on Inc-v4, while MI-FGSM, PGD, I-FGSM, and FGSM reached 54.3%, 18.9%, 27.5%, and 32.1%, demonstrating the effectiveness of the proposed method. An original image and the corresponding adversarial example generated for Inc-v3 by AI-FGM are shown in Figure 2.

4.2.1. Decay Factors. The terms β_1 and β_2 control not only the decay amplitude of the step sizes but also the accumulation intensity of gradients for m_t and v_t . As such, they have a direct impact on attack success rates. We applied a grid-search method to identify an optimal set of β_1 and β_2 values. In the experiments, the maximum perturbation ϵ was set to 16, and the number of iterations T was set to 10. The values of β_1 and β_2 ranged from 0 to 1. Notably, we not only chose the values of β_1 and β_2 uniformly from 0.1 to 0.9 but also selected values close to 0 and 1, as shown in Figure 3. This was done to provide a more comprehensive study on the effects of decay factors, as we propose that the

Input: A convolutional neural network f and the corresponding cross-entropy loss function $J(\theta, x, y)$; an original image x and the corresponding ground-truth label y ; the number of iterations T ; the iteration time step t ; the dimension of the input image N ; the size of the perturbation ε ; Adam decay factors β_1 and β_2 ; and a denominator stability factor δ .

Output: An adversarial example x^* , s.t. $\|x - x^*\|_\infty \leq \varepsilon$.

- (1) $m_0 = 0$, $v_0 = 0$, $x_0^* = x$, and $t = 0$;
- (2) $\alpha = \varepsilon \cdot \sqrt{N}$
- (3) while $t < T$ do:
- (4) $g_t = (\nabla_x J(\theta, x_t^*, y) / \|\nabla_x J(\theta, x_t^*, y)\|_1)$
- (5) $m_{t+1} = \beta_1 m_t + (1 - \beta_1) g_t$;
- (6) $v_{t+1} = \beta_2 \cdot v_t + (1 - \beta_2) \cdot g_t^2$;
- (7) $s_{t+1} = (m_{t+1} / \delta + \sqrt{v_{t+1}})$;
- (8) $\alpha_t = \alpha \cdot (\sqrt{1 - \beta_2^{t+1}} / 1 - \beta_1^{t+1}) / \sum_{i=0}^{T-1} (\sqrt{1 - \beta_2^{i+1}} / 1 - \beta_1^{i+1})$;
- (9) $x_{t+1}^* = x_t^* + \alpha_t \cdot (s_{t+1} / \|s_{t+1}\|_2)$;
- (10) $x_{t+1}^* = \text{Clip}(x_{t+1}^*, x - \varepsilon, x + \varepsilon)$;
- (11) $t = t + 1$;
- (12) end while
- (13) return $x^* = x_T^*$.

ALGORITHM 1: Adam iterative fast gradient method.

TABLE 1: Attack success rates (%) for all seven networks included in the study.

	Attack	Inc-v3	Inc-v4	Inces-v2	Res-101	Inc-v3 _{ens3}	Inc-v3 _{ens4}	Inces-v2 _{ens}
Inc-v3	FGSM	72.2	32.1	31.7	32.3	10.3	10.6	4.2
	I-FGSM	100.0	27.5	23.0	20.9	6.2	4.7	1.8
	PGD	99.8	18.9	14.7	14.4	6.1	6.2	3.2
	MI-FGSM	100.0	54.3	50.6	44.0	13.9	13.3	6.5
	AI-FGM	100.0	60.7	55.8	50.2	17.0	16.8	8.5
Inc-v4	FGSM	39.4	64.9	29.5	33.0	12.2	11.1	5.0
	I-FGSM	43.8	100.0	27.4	23.4	6.1	6.3	2.4
	PGD	32.0	99.8	17.7	16.6	6.4	5.8	3.1
	MI-FGSM	69.9	100.0	57.9	54.1	19.6	17.7	8.7
	AI-FGM	72.9	100.0	60.1	57.1	21.7	19.5	10.3
Inces-v2	FGSM	37.6	31.8	57.9	31.4	13.7	12.0	6.8
	I-FGSM	46.1	35.0	99.4	30.4	7.3	6.7	4.4
	PGD	30.1	23.0	97.3	18.3	6.1	5.7	2.7
	MI-FGSM	73.5	69.3	99.5	60.0	27.0	23.0	16.7
	AI-FGM	74.6	71.1	99.5	61.8	31.3	25.7	20.5
Res-101	FGSM	38.3	33.0	30.2	79.3	14.6	13.3	6.4
	I-FGSM	35.1	28.3	25.1	99.5	8.4	6.7	3.7
	PGD	30.9	22.4	20.9	99.6	7.3	7.2	3.4
	MI-FGSM	60.0	55.3	50.6	99.5	22.9	19.8	11.3
	AI-FGM	64.0	57.7	54.0	99.5	27.2	24.1	15.4

The diagonal blocks indicate white-box attacks, while the off-diagonal blocks indicate black-box attacks.

FIGURE 2: Continued.

FIGURE 2: Classification of a normal image and corresponding adversarial example by Inc-v3 and Inc-v4. The first row shows the top 10 confidence distributions for a clean image, for which both models provided a correct prediction. The second row shows the top 10 confidence distributions for the adversarial example generated for Inc-v3 by AI-FGM. It is evident that the adversarial example successfully attacked Inc-v3 (white-box) and Inc-v4 (black-box) with high confidence. (a) Clean. (b) Adversarial.

FIGURE 3: Attack success rates (%) of adversarial examples generated for Inc-v3 with AI-FGM, applied to Inc-v3 (white-box), Inc-v4 (black-box and normally trained), and Inc-v3ens4 (black-box and adversarially trained) with β_1 and β_2 in the range of $(0, 1)$. (a) Inc-v3. (b) Inc-v4. (c) Inc-v3_{ens4}.

relationship between the success rate and the decay factors may be nonlinear. Adversarial examples were then generated on Inc-v3 with AI-FGM and used to perform attacks on Inc-v3, Inc-v4, and Inc-v3_{ens4}. It is evident from the figure that regardless of the values of β_1 and β_2 , attack success rates were always near 100% in white-box environments. Beyond that, attack success rates were more sensitive to β_1 for black-box models, regardless of whether networks were normally or

adversarially trained. Success rates were maximized when both β_1 and β_2 were close to 1.

4.2.2. The Number of Iterations. The effect of iteration quantities on success rates was studied by performing attacks using iterative methods. The maximum perturbation ϵ was set to 16, and the decay factors β_1 and β_2 were set to 0.99 and

0.999, respectively. The number of iterations T ranged from 1 to 20. Adversarial examples generated on Inc-v3 with I-FGSM, MI-FGSM, and AI-FGM were then used to attack Inc-v3 and Inc-v4, as shown in Figure 4.

As shown in the figure, the success rates of black-box attacks for several iterative methods decreased with the increase of the iteration quantities. However, AI-FGM always outperformed I-FGSM and MI-FGSM for a given value of T .

4.2.3. Perturbation Size. The effect of adversarial perturbation size on attack success rates was studied by setting the number of iterations T to 10, with decay factors of 0.99 and 0.999. The size of perturbations ϵ ranged from 1 to 30. Adversarial examples generated on Inc-v3 with I-FGSM, MI-FGSM, and AI-FGM were used to attack Inc-v3 and Inc-v4, as shown in Figure 5.

It is evident from the figure that attack success rates can reach 100% in white-box environments. Success rates for black-box models increased steadily with increasing values of ϵ . AI-FGM always outperformed I-FGSM and MI-FGSM for a given value of ϵ . In other words, AI-FGM can achieve comparable black-box attack success rates with smaller perturbations.

4.3. Attacking an Ensemble of Networks. The experimental results presented above suggest that AI-FGM can increase the transferability of adversarial examples. In addition, the success rate of black-box attacks can be further improved by attacking an ensemble of networks. As discussed in Section 3.3, multiple models were attacked by fusing network logits. In the experiment, all seven networks described in Section 4.1 were used, and we generated adversarial examples on the ensemble of Inc-v3, Inc-v4, IncRes-v2, and Res-101, with FGSM, I-FGSM, MI-FGSM, NI-FGSM, and AI-FGM. Attacks were then performed on the other three defense models. Decay factors were set to 0.99 and 0.999; the number of iterations was 10; the maximum perturbation was 16; and each network had an equal ensemble weight of $\omega_k = (1/4)$. The corresponding results are shown in Table 2, and it is evident that AI-FGM was more effective than the other four methods on adversarially trained models.

4.4. Combination with Advanced Methods. In this subsection, we combined our AI-FGM with SI-NI-TI, SI-NI-DI, and SI-NI-TI-DI [18] and compared the black-box attack success rates of our extensions with the original methods under a single-model setting. The adversarial examples were generated on Inc-v3, with the number of iterations set to 10 and the maximum perturbation to 16. It is noteworthy that our combination with other methods is more like an improvement. For example, the combination of AI-FGM and SI-NI-TI was done by replacing the Nesterov optimizer in SI-NI-TI with Adam, thus resulting in SI-AI-TI. As shown in Table 3, our method SI-AI-TI-DI achieved an average attack success rate of 65.1%, surpassing the state-of-the-art attack by 10.7%.

Furthermore, we generated adversarial examples on the ensemble models using our SI-AI-TI-DI. As shown in

FIGURE 4: Attack success rates (%) of adversarial examples generated for Inc-v3 with I-FGSM, MI-FGSM, and AI-FGM, applied to Inc-v3 (white-box) and Inc-v4 (black-box) with T ranging from 1 to 20. It is noteworthy that the curves of Inc-v3 vs. I-FGSM, Inc-v3 vs. MI-FGSM, and Inc-v3 vs. AI-FGM overlap.

FIGURE 5: Attack success rates (%) of adversarial examples generated for Inc-v3 with I-FGSM, MI-FGSM, and AI-FGM, applied to Inc-v3 (white-box) and Inc-v4 (black-box) with ϵ ranging from 1 to 30. It is noteworthy that the curves of Inc-v3 vs. I-FGSM, Inc-v3 vs. MI-FGSM, and Inc-v3 vs. AI-FGM overlaps.

Table 4, we achieved an average attack success rate of 95.0% on adversarially trained models under the black-box setting, which raised a new security issue for the robust deep neural networks.

5. Discussion

It is commonly acknowledged that the training of neural network models is similar to the generation of adversarial

TABLE 2: Attack success rates (%) on adversarially trained networks under the ensemble-model setting.

Attack	Inc-v3 _{ens3}	Inc-v3 _{ens4}	Incres-v2 _{ens}	Average
FGSM	18.6	17.3	9.5	15.1
I-FGSM	20.1	16.9	11.4	16.1
MI-FGSM	49.4	43.2	25.9	39.5
NI-FGSM	46.7	41.6	23.0	37.1
AI-FGM (ours)	56.2	50.6	33.0	46.6

The adversarial examples are generated on the ensemble of Inc-v3, Inc-v4, IncRes-v2, and Res-101.

TABLE 3: Comparison of combined methods and original methods on attack success rates (%) against adversarially trained networks under the single-model setting.

Attack	Inc-v3 _{ens3}	Inc-v3 _{ens4}	Incres-v2 _{ens}	Average
SI-NI-TI	54.7	52.3	34.4	47.1
SI-AI-TI (ours)	56.3	54.4	39.4	50.0
SI-NI-DI	39.5	37.2	19.3	32.0
SI-AI-DI (ours)	50.8	51.5	27.7	43.3
SI-NI-TI-DI	62.3	59.3	41.7	54.4
SI-AI-TI-DI (ours)	72.6	69.8	53.0	65.1

The adversarial examples are generated on the ensemble of Inc-v3, Inc-v4, Incres-v2, and Res-101.

TABLE 4: Comparison of SI-AI-TI-DI and SI-NI-TI-DI under the ensemble-model setting.

Attack	Inc-v3 _{ens3}	Inc-v3 _{ens4}	Incres-v2 _{ens}	Average
SI-NI-TI-DI	95.4	93.7	89.5	92.9
SI-AI-TI-DI (ours)	96.2	96.0	92.7	95.0

The adversarial examples are generated on the ensemble of Inc-v3, Inc-v4, Incres-v2, and Res-101.

examples, especially for gradient-based generation methods. Hence, techniques used in the training of neural networks, to improve model generalizability, can also be adopted to improve the transferability of adversarial examples. Since the Adam optimizer is often used in the training of neural networks, to improve convergence and achieve better performance on test sets, a second momentum term and decay step size in Adam were included in the AI-FGM algorithm. This was done to improve the transferability of adversarial examples. Furthermore, we suggest that other techniques (such as data augmentation) could be used to improve the performance of adversarial examples further in black-box environments.

6. Conclusions

In this study, we proposed the Adam iterative fast gradient method to improve the transferability of adversarial examples. Specifically, the Adam algorithm was modified to increase its suitability for the generation of adversarial examples. The proposed method improved both the iteration update direction and step size. The effectiveness of the proposed method was verified by an extensive series of

experiments with ImageNet. Compared with previous gradient-based adversarial example generation techniques, our method improved attack success rates in black-box environments. In addition, we further improved the transferability of adversarial examples by combining our approach with existing methods and attacking ensemble models, which achieved a state-of-the-art attack success rate against adversarially trained networks. We suggest the proposed method could be used as a reference for other iterative gradient-based methods. For example, data augmentation could be combined with AI-FGM to achieve better attack performance.

Data Availability

The public data (data set and models) can be downloaded from https://github.com/tensorflow/cleverhans/tree/master/examples/nips17_adverarial_competition/dataset, <https://github.com/tensorflow/models/tree/master/research/slim>, and https://github.com/tensorflow/models/tree/master/research/adv_imagenet_models. The source code of the proposed method in this paper is available at https://github.com/YinHeng121/Adam_attack.

Conflicts of Interest

The authors declare that there are no conflicts of interest regarding the publication of this paper.

Authors' Contributions

Heng Yin and Hengwei Zhang contributed equally to this work.

Acknowledgments

This work was supported by the National Key Research and Development Program of China, under grant no. 2017YFB0801900.

References

- [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Communications of the ACM*, vol. 60, no. 6, pp. 84–90, 2017.
- [2] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014, <http://arxiv.org/abs/1409.1556>.
- [3] C. Szegedy, "Going deeper with convolutions," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–9, Boston, MA, USA, June 2015.
- [4] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, Las Vegas, NV, USA, June 2016.
- [5] C. Szegedy, "Intriguing properties of neural networks," in *Proceedings of the 2013 International Conference on Learning Representations*, Banff, Canada, April 2013.
- [6] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," 2014, <http://arxiv.org/abs/1412.6572>.

- [7] A. Kurakin, I. Goodfellow, and S. Bengio, "Adversarial examples in the physical world," 2016, <http://arxiv.org/abs/1607.02533>.
- [8] Y. Dong, "Boosting adversarial attacks with momentum," in *Proceedings of the 2018 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9185–9193, Salt Lake City, UT, USA, June 2018.
- [9] C. Xie, "Improving transferability of adversarial examples with input diversity," in *Proceedings of the 2019 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2725–2734, Long Beach, CA, USA, June 2019.
- [10] N. Carlini and D. Wagner, "Towards evaluating the robustness of neural networks," in *Proceedings of the 2017 IEEE Symposium on Security and Privacy*, pp. 39–57, San Jose, CA, USA, May 2017.
- [11] Y. Liu, X. Chen, C. Liu, and D. Song, "Delving into transferable adversarial examples and black-box attacks," 2016, <http://arxiv.org/abs/1611.02770>.
- [12] A. I. Ilyas, L. Engstrom, A. Athalye, and J. Lin, "Black-box adversarial attacks with limited queries and information," 2018, <http://arxiv.org/abs/1804.08598>.
- [13] D. Wierstra, T. Schaul, T. Glasmachers, Y. Sun, J. Peters, and J. Schmidhuber, "Natural evolution strategies," *Journal of Machine Learning and Research*, vol. 15, pp. 949–980, 2014.
- [14] M. Alzantot, Y. Sharma, S. Chakraborty, H. Zhang, C. J. Hsieh, and M. B. Srivastava, "GenAttack: practical black-box attacks with gradient-free optimization," in *Proceedings of the 2019 Genetic Evolution Computer Conference*, pp. 1111–1119, New York, NY, USA, July 2019.
- [15] W. Brendel, J. Rauber, and M. Bethge, "Decision-based adversarial attacks: reliable attacks against black-box machine learning models," 2017, <http://arxiv.org/abs/1712.04248>.
- [16] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, "Practical black-box attacks against machine learning," in *Proceedings of the 2017 Asia Conference on Computer and Communications Security*, pp. 506–519, Abu Dhabi, UAE, April 2017.
- [17] D. P. Kingma and J. Ba, "A method for stochastic optimization," 2014, <http://arxiv.org/abs/1412.6980>.
- [18] J. Lin, C. Song, K. He, L. Wang, and J. E. Hopcroft, "Nesterov accelerated gradient and scale invariance for adversarial attacks," in *Proceedings of the ICLR 2020*, Addis Ababa, Ethiopia, April 2020.
- [19] Y. Dong, T. Pang, H. Su, and J. Zhu, "Evading defenses to transferable adversarial examples by translation-invariant attacks," in *Proceedings of the 2019 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4307–4316, Long Beach, CA, USA, June 2019.
- [20] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," 2018, <http://arxiv.org/abs/1706.06083>.
- [21] N. Papernot, P. McDaniel, X. Wu, S. Jha, and A. Swami, "Distillation as a defense to adversarial perturbations against deep neural networks," in *Proceedings of the 2016 IEEE Symposium on Security and Privacy*, pp. 582–597, San Jose, CA, USA, May 2016.
- [22] J. Buckman, A. Roy, C. Raffel, and I. Goodfellow, "Thermometer encoding: one hot way to resist adversarial examples," in *Proceedings of the 2018 International Conference Learning Representations*, Vancouver, BC, Canada, May 2018.
- [23] C. Guo, M. Rana, M. Cissé, and L. Maaten, "Countering adversarial images using input transformations," 2017, <http://arxiv.org/abs/1711.00117>.
- [24] P. Samangouei, M. Kabkab, and R. Chellappa, "Defense-gan: protecting classifiers against adversarial attacks using generative models," 2018, <http://arxiv.org/abs/1805.06605>.
- [25] A. Levine and S. Feizi, "Wasserstein smoothing: Certified robustness against wasserstein adversarial attacks," in *Proceedings of the 2020 International Conference Artificial Intelligent Statistic Palermo*, Sicily, Italy, August 2020.
- [26] C. Mao, "Multitask learning strengthens adversarial robustness," 2020, <http://arxiv.org/abs/2007.07236>.
- [27] F. Wu, W. Yang, L. Xiao, and J. Zhu, "Adaptive wiener filter and natural noise to eliminate adversarial perturbation," *Electronics*, vol. 9, no. 10, 1634 pages, 2020.
- [28] J. M. Cohen, E. Rosenfeld, and Z. Kolter, "Certified adversarial robustness via randomized smoothing," 2019, <http://arxiv.org/abs/1902.02918>.
- [29] Y. Shi, C. Fan, L. Zou, C. Sun, and Y. Liu, "Unsupervised adversarial defense through tandem deep image priors," *Electronics*, vol. 9, no. 11, 1957 pages, 2020.
- [30] A. Kurakin, I. Goodfellow, and S. Bengio, "Adversarial machine learning at scale," 2016, <http://arxiv.org/abs/1611.01236>.
- [31] F. Tramèr, A. Kurakin, N. Papernot, D. Boneh, and P. McDaniel, "Ensemble adversarial training: attacks and defenses," 2017, <http://arxiv.org/abs/1705.07204>.
- [32] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, "On the importance of initialization and momentum in deep learning," in *Proceedings of the 2013 International Conference Machine Learning*, pp. 1139–1147, Atlanta, GA, USA, June 2013.
- [33] T. Tieleman and G. Hinton, "Lecture 6.5-RMSProp, coursera: neural networks for machine learning," 2012.
- [34] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2818–2826, Las Vegas, NV, USA, June 2016.
- [35] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. Alemi, "Inception-v4, inception-ResNet and the impact of residual connections on learning," in *Proceedings of the 2017 AAAI Conference Artificial Intelligence*, San Francisco, California, USA, February 2017.
- [36] K. He, X. Zhang, S. Ren, and J. Sun, "Identity mappings in deep residual networks," in *Proceedings of the 2016 Computer Vision*, pp. 630–645, Cham, Switzerland, September 2016.