Logit Margin Matters: Improving Transferable Targeted Adversarial Attack by Logit Calibration

Anonymous Author(s)
Affiliation
Address
email

Abstract

Previous works have extensively studied the transferability of adversarial samples in untargeted black-box scenarios. However, it still remains challenging to craft the targeted adversarial examples with higher transferability than non-targeted ones. Recent studies reveal that the traditional Cross-Entropy (CE) loss function is insufficient to learn transferable targeted perturbations due to the issue of vanishing gradient. In this work, we provide a comprehensive investigation of the CE function and find that the logit margin between the targeted and non-targeted classes will quickly obtain saturated in CE, which largely limits the transferability. Therefore, in this paper, we devote to the goal of enlarging logit margins and propose two simple and effective logit calibration methods, which are achieved by downscale the logits with a temperature factor and an adaptive margin, respectively. Both of them can effectively encourage the optimization to produce larger logit margins and lead to higher transferability. Besides, we show that minimizing the cosine distance between the adversarial examples and the targeted classifier can further improve the transferability, which is benefited from downscale logits via L2-normalization. Experiments conducted on the ImageNet dataset validate the effectiveness of the proposed methods, which outperforms the state-of-the-art methods in black-box targeted attacks. The source code of our method is available at Link.

1 Introduction

In the past decade, deep neural networks (DNNs) have achieved remarkable success in various fields, *e.g.*, image classification [24], image segmentation [19], and object detection [23]. However, Goodfellow *et al.* [5] reveal that the DNNs are vulnerable to adversarial attacks, in which adding imperceptible disturbances to the input can lead the DNNs to make an incorrect prediction. Following, many approaches [3, 4, 1, 27, 29] have been proposed to construct more destructive adversarial samples for investigating the vulnerability of the DNNs. [5, 18] also show that the adversarial samples are transferable across different networks, raising a more critical robustness threat under the black-box scenarios. Therefore, it is vital of exploring the vulnerability of the DNNs, which is very useful for designing robust DNNs.

Currently, most of the works [3, 29, 16, 10, 28, 6] have been devoted to the untargeted black-box attack, in which adversarial examples are crafted to fool unknown surrogate models to predict unspecified incorrect labels. For example, [3, 29] leveraged input-level transformation or augmentation to improve the non-targeted transferability. [10] proposed a powerful intermediate feature-level attack. [28, 6] demonstrated that backpropagating more gradients through the skip-connections can increase the transferability. Despite the success in non-targeted cases, the targeted transferability remains challenging, which requires eliciting the black-box models into a pre-defined target category label.

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

For learning the transferable adversarial samples in untargeted cases, most methods have leveraged the Cross-Entropy (CE) as the loss function. However, [15, 30] recently showed that the CE loss is insufficient for learning the adversarial perturbation in the targeted case due to the issue of vanishing gradient. To deal with this issue, [15] adopt the Poincaré distance to increase the gradient magnitude during the optimization adaptively. [30] demonstrated that an effortless logit loss equal to the negative value of the targeted logits could alleviate the gradient issue and achieve surprisingly strong targeted transferability. Besides, [30] also showed that optimizing with more iterations can significantly increase the targeted transferability.

In this study, we take a closer look at the vanishing gradient issue in the CE loss function and find that the logit margin between the targeted and non-targeted classes will quickly get saturated during the optimization. Moreover, this issue will influence the performance of the perturbations and thus essentially limit the transferability. Specifically, along with the training iterations in CE, we observe that the logits of the targeted and non-targeted classes increase rapidly in the first few iterations. However, after reaching the peak, the logit margin between the targeted and non-targeted classes will get saturated, and further training will decrease the logits simultaneously to maintain this margin. This phenomenon is mainly due to the fact that the softmax function in CE will approximately output the probability of the target class to 1 when reaching the saturated margin (*e.g.*, 10). Thus, it raises the problem that the transferability will not be further increased even optimized with more iterations. While in practice, we are encouraged to increase the transferability by maximizing both the logit for the targeted class and its margin against other non-targeted classes to cross the decision boundaries of other black-box models.

In this paper, we devote to enlarging logit margins to alleviate the above saturation issue in CE. Inspired by the temperature-scaling used in the knowledge distillation [8], a higher temperature T will produce a softer probability distribution over different classes. We firstly leverage this scaling technique into the targeted adversarial attack to calibrate the logits. Then the logits margin between the targeted and non-targeted classes will not be saturated after only a few iterations and will keep improving the transferability. On the other aspect, instead of using a constant T, we further explored an adaptive margin-based calibration by scaling the logits based on the logit margin of the target class and the highest non-target class. In addition, we also investigate the effectiveness of calibrating the targeted logit into the unit length feature space by L2-normalization, which is equivalent to minimizing the angle between the adversarial examples and the targeted classifier.

Finally, we conduct experiments on the ImageNet dataset to validate the effectiveness of the logits calibration for crafting transferable targeted adversarial examples. Experimental results demonstrated that the calibration of the logits helps achieve a higher attack success rate than other state-of-the-art methods. Additionally, we tested the logit calibration in Generative Adversarial Networks (GANs)-based TTP method [21] to train the target-class-specific generators, which is also beneficial for increasing the transferability in the resource-intensive method.

2 Related Works

In this section, we give a brief introduction of the related works from the following two aspects: *untargeted black-box attacks* and *targeted attacks*.

2.1 Untargetd Black-box Attacks

After [25] exposed the vulnerability of deep neural networks, many attack methods [29, 4] have been proposed to craft highly transferable adversaries in the non-targeted scenario. We first briefly review several seminal gradient-based attack methods that focus on enhancing the transferability against black-box models. These methods also can be leveraged to improve the transferability in the targeted scenario.

Iterative-Fast Gradient Sign Method (I-FGSM) [14] is an iterative version of FGSM [5], which adds a small perturbation with a small step size α in the gradient direction iteratively:

$$\hat{x}_0 = x, \quad \hat{x}_{i+1} = \hat{x}'_i + \alpha \cdot \text{sign}(\nabla_{\hat{x}} J(\hat{x}'_i, y)), \tag{1}$$

where $\hat{x}'_i$ denotes the intermediate adversarial image in the i_{th} iteration, $\alpha = \epsilon/T$ makes the adversaries be constrained by an upper-bound ϵ through the l_p-norm, and T is number of iterations.

Momentum Iterative-FGSM (MI-FGSM) [3] integrates a momentum term into I-FGSM to accumulate the update direction of previous steps for achieving better transferablily. The iteration process can be summarized as:

$$g_{i+1} = \mu \cdot g_i + \frac{\nabla_{\hat{x}} J(\hat{x}_i', y)}{||\nabla_{\hat{x}} J(\hat{x}_i', y)||_1}, \quad \hat{x}_{i+1} = \hat{x}_i' + \alpha \cdot \text{sign}(g_{i+1}), \tag{2}$$

where g_i is the momentum term of gradient in the i_{th} iteration and μ is a decay factor.

Translation Invariant-FGSM (TI-FGSM) [4] is different from the above methods, which calculate the average gradient on multiple translated images for the update. This technique can prevent the attack from overfitting to the specific white-box model. To improve the computational efficiency, the TI-FGSM proposes to convolve the gradient of the untranslated image with a Gaussian kernel instead of calculating the gradient on translated images, as follows:

$$g_{i+1} = \mu \cdot g_i + \frac{W * \nabla_x J\left(f\left(x_i^{adv}\right), y\right)}{\left\|W * \nabla_x J\left(f\left(x_i^{adv}\right), y\right)\right\|_1}, \quad x_{i+1}^{adv} = x_i^{adv} + \alpha \cdot \text{sign}\left(g_{i+1}\right), \tag{3}$$

where W is a Gaussian kernel used for smoothing gradient and $*$ is the operator of convolution.

Diverse Input-FGSM (DI-FGSM) [29] is conceptually similar to TI-FGSM, but it is designed based on the input transformations with random resizing and padding operations rather than smoothing gradients. DI-FGSM can be represented as:

$$\begin{aligned} T(\hat{x}_i', p) &= \begin{cases} \text{RandPad}(\text{RandResize}(\hat{x}_i')), & \kappa \leq p \\ \hat{x}_i', & \kappa > p \end{cases}, \\ \hat{x}_{i+1} &= \hat{x}_i' + \alpha \cdot \text{sign}(\nabla_{\hat{x}} J(T(\hat{x}_i', p), y)), \end{aligned} \tag{4}$$

where $\kappa \in \mathcal{U}(0, 1)$ controls the using of the transformed image or the original one for gradient calculation. $T(\hat{x}_i', p)$ is implemented with the probability p at each iteration, which will bring more randomness since augmentation parameters are not fixed during iterations.

2.2 Targeted Attacks

Targeted attacks are different from non-targeted attacks, which need to change the decision to a specific target class. [13] integrates the above non-targeted attack methods into targeted attacks to craft targeted adversarial examples. However, the performance is limited because it is insufficient to fool the black-box model only by maximizing the probability of the target class.

Po+Trip [15] proposes the Poincaré space as a metric space and utilizes Triplet loss to improve targeted transferability by forcing adversarial example toward the target label and away from the ground-truth label. The Poincaré distance can alleviate the vanishing gradient in the cross-entropy loss and adaptively adjust the magnitude of gradient, which can be described as:

$$\begin{aligned} L_{Po}(\boldsymbol{x}', t) &= d(\boldsymbol{u}, \boldsymbol{v}) = \operatorname{arccosh}(1 + \delta(\boldsymbol{u}, \boldsymbol{v})), \\ \delta(\boldsymbol{u}, \boldsymbol{v}) &= \frac{2 \cdot \|\boldsymbol{u} - \boldsymbol{v}\|_2^2}{\left(1 - \|\boldsymbol{u}\|_2^2\right)\left(1 - \|\boldsymbol{v}\|_2^2\right)}, \boldsymbol{u} = \frac{z\left(\boldsymbol{x}'\right)}{\|z(\boldsymbol{x}')\|_1}, \boldsymbol{v} = \max\{\mathbb{1}_t - \xi, 0\}, \end{aligned} \tag{5}$$

where $z(\boldsymbol{x}')$ is the output logits of the surrogate model, $\boldsymbol{x}'$ is the perturbed image, $\mathbb{1}_t$ is the one-hot encoding related to the target class t and ξ is a small constant (0.0001). For further improving performance, a new triplet loss is proposed to reduce the distance between the output of adversarial example and the target label, and to increase the margin from the true label y, as below:

$$\begin{aligned} L_{\text{Trip}}\left(t, z(\boldsymbol{x}'), y\right) &= \left[D\left(z\left(\boldsymbol{x}'\right), \mathbb{1}_t\right) - D\left(z\left(\boldsymbol{x}'\right), \mathbb{1}_y\right) + \gamma\right]_+, \\ D\left(z\left(\boldsymbol{x}'\right), \mathbb{1}_y\right) &= 1 - \frac{\left|z\left(\boldsymbol{x}'\right) \cdot \mathbb{1}_y\right|}{\left\|z\left(\boldsymbol{x}'\right)\right\|_2 \left\|\mathbb{1}_y\right\|_2}. \end{aligned} \tag{6}$$

The overall loss function of Po+Trip is defined as $L_{all} = L_{Po}(\boldsymbol{x}', t) + \lambda \cdot L_{\text{Trip}}\left(t, z(\boldsymbol{x}'), y\right)$, where λ controls the weight of the triplet loss.

Logits [30] adopts a simple and straightforward idea by directly maximizing the target logit to pull the adversarial examples close to the target class for addressing the vanishing gradient problem in existing transferable attacks, which can be expressed as:

$$L_{Logit} = -z_t(\boldsymbol{x}'), \tag{7}$$

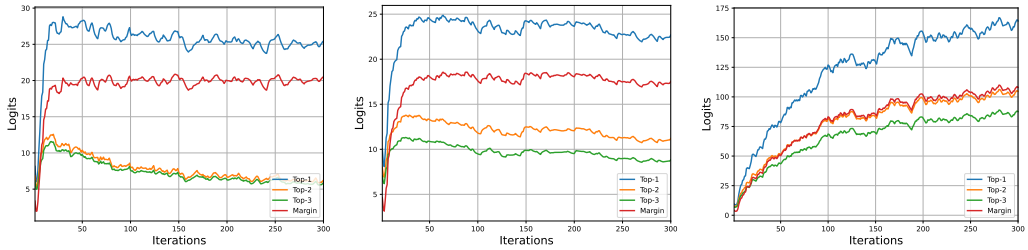

(a) Cross-Entropy (b) Po+Trip (c) Logit

Figure 1: The average Top-3 logits and logit margin of 50 adversarial samples trained by the Cross-Entropy, Po+Trip and Logit loss functions for crafting the ResNet-50.

where $z_t(\cdot)$ is the output logits of the target class.

Since targeted transferability is extremely difficult, many studies employ resource-intensive approaches, which train target class-specific models (auxiliary classifiers or generative models) on additional large-scale data, as an auxiliary. For example, the FDA methods [12, 11] use the intermediate feature distributions of CNNs to boost the targeted transferability by training class-specific auxiliary classifiers to model layer-wise feature distributions. The GAP [22] trains a generative model for crafting targeted adversarial examples. Subsequently, [20] uses a relativistic training objective to train the generative model for improving attack performance and cross-domain transferability. Recently, the TTP [21] utilizes the global and local distribution matching for training target class-specific generators for obtaining high targeted transferability. However, the TTP requires actual data samples from the target class and brings expensive training costs. Different from the above methods, we introduce two simple and effective logit calibration methods into the CE loss function, which achieve competitive performance without additional data and training.

3 Method

3.1 Problem Definition

Given a white-box surrogate model $\mathbb{F}_s$ and an input x not from the targeted class t, our primary goal is to learn an imperceptible perturbation δ that can fool the $\mathbb{F}_s$ to output the target t for $\hat{x} = x + \delta$. Besides, the prediction of $\hat{x}$ will also be t when feeding to other unknown black-box surrogate models. The l_∞-norm is usually used to constrained the δ within an upper-bound ϵ, denoted as $||\delta||_\infty \leq \epsilon$.

For the surrogate model $\mathbb{F}_s$, we denoted the feature for the final classification layer of the input x as $\phi(x)$. The logit z_i of a category i is computed by $z_i = W_i^T \phi(x) + b_i$, the W_i and b_i are the classifier weights and bias. The corresponding probability p_i after the softmax is $p_i = \frac{e^{z_i}}{\sum e^{z_j}}$.

3.2 Logit Margin

When successfully attacked the $\mathbb{F}_s$, the logit z_t of the target class will be higher than the logits z_{nt} of any other non-target class in the classification task. Their logit margins can be computed by,

$$G(\phi(\hat{x})) = z_t - z_{nt} = W_t^T \phi(\hat{x}) + b_t - W_{nt}^T \phi(\hat{x}) + b_{nt}. \tag{8}$$

[15, 30] showed that it is insufficient to obtain transferable targeted adversarial samples that are only close to the target class while not away from true class and other non-targeted classes. Based on this property, it encourages us to continually enlarge this logit margin to increase the separation between the targeted and other non-targeted classes.

To have a better understanding of the relationship between the logit margins and the targeted transferability, we visualize the average Top-3 logits (1 targeted class and other two non-targeted classes) of 50 random adversarial samples trained for crafted ResNet50 by the CE, Po+Trip, and the Logit loss functions with MI, DI and TI following [30]. We also compute the average logit margin of the targeted class against the Top-20 non-targeted classes. The logit and margin are shown in Figure 1, and the transferability from ResNet50 to VGG16 is plot in Figure 2.

From Figure 1, we can observe that the logits of the targeted class and the Top-2 non-targeted classes increase rapidly in the first few iterations for the CE and Po+Trip loss, as well as their logit margins. When reaching the peak, the margin is saturated, and the logits start to decrease simultaneously to maintain the saturated margin. By comparing the CE and Po+Trip, the Po+Trip needs fewer more iterations to reach the saturated status and thus shows a marginal better transferability than CE, as shown in Figure 2. In comparison, the Logit loss function will keep increasing the logits of the targeted category and the logit margin. Therefore, the Logit loss function shows a much better targeted-attack success rate than CE and Po+Trip. On the other hand, the Logit loss also significantly increases the logits for other non-targeted classes.

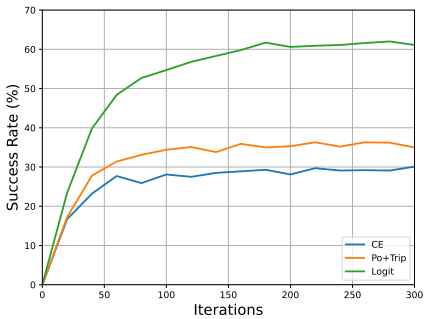

Figure 2: The targeted attack success rate (%) on VGG-16 by using the ResNet-50 as the surrogate model.

To further analyze why the CE loss function saturated to this logit margin and explore the effectiveness of increasing the margin during training, in the following sections, we will revisit the cross-entropy loss function and introduce the logit calibration to achieve this goal.

3.3 Revisiting the Cross-Entropy Loss

Firstly, our objective is to maximize the logit margin in Eq. 8. After computing the gradient w.r.t. to $\phi(\hat{x})$, we can get

$$\frac{\partial G}{\partial \phi(x)} = W_t - W_{nt}. \tag{9}$$

This gradient indicates that the adversarial feature $\phi(\hat{x})$ needs to move towards the target class while apart from those non-target classes. Next, we compute the gradient w.r.t. to $\phi(\hat{x})$ in the Cross-Entropy loss function

$$L_{CE} = -\log(p_t) = -z_t + \log(\sum e^{z_k}), \tag{10}$$

and get the gradient $\frac{\partial L_{ce}}{\partial \phi(\hat{x})}$ as

$$\begin{aligned} \frac{\partial L_{ce}}{\partial \phi(\hat{x})} &= -\frac{\partial z_t}{\partial \phi(\hat{x})} + \frac{1}{\sum e^{z_k}} \cdot \frac{\partial \sum e^{z_k}}{\partial \phi(\hat{x})} \\ &= -\frac{\sum e^{z_i}}{\sum e^{z_k}} \cdot \frac{\partial z_t}{\partial \phi(x)} + \frac{1}{\sum e^{z_k}} \sum e^{z_i} \frac{\partial z_i}{\partial \phi(\hat{x})} \\ &= \sum \frac{e^{z_i}}{\sum e^{z_j}} \cdot (\frac{\partial z_j}{\partial \hat{x}} - \frac{\partial z_t}{\partial \hat{x}}) \\ &= \sum -p_i(W_t - W_i). \end{aligned} \tag{11}$$

From Eq. 11, we actually can find the CE loss function is designed to adaptively optimize the $\phi(\hat{x})$ towards W_t and away from other W_i. However, after being optimized for several iterations, the p_i of the non-targeted class will quickly approximate to 0 and then significantly vanish the $W_t - W_i$.

Let's consider the case only with 2 classes (t and nt), we have the probabilities p_t and p_{nt} as:

$$p_t = \frac{e^{z_t}}{e^{z_t} + e^{z_{nt}}} = \frac{1}{1 + e^{-(z_t - z_{nt})}}, \tag{12}$$

$$p_{nt} = \frac{e^{z_{nt}}}{e^{z_t} + e^{z_{nt}}} = \frac{1}{1 + e^{(z_t - z_{nt})}}. \tag{13}$$

As shown in Figure 3, the p_t will get close to 1 when $z_t - z_{nt} > 6$ (*e.g.*, $p_{nt} \approx 2e^{-9}$ when $z_t - z_{nt} = 20$). In such a context, the gradient will significantly vanish. Recall that, in the CE loss function (Figure 1 (a)), the logit margin between Top-1 and Top-2 logits first increases rapidly but will reach saturated status when approaching a certain value. This further indicates that the optimization of the CE loss function is largely restrained when the logit margin reaches a certain value.

To this end, we raise the question *if we explicitly enforce the optimization to enlarge the logit margin ($z_t - z_{nt}$), could we get better transferable targeted adversarial samples?*

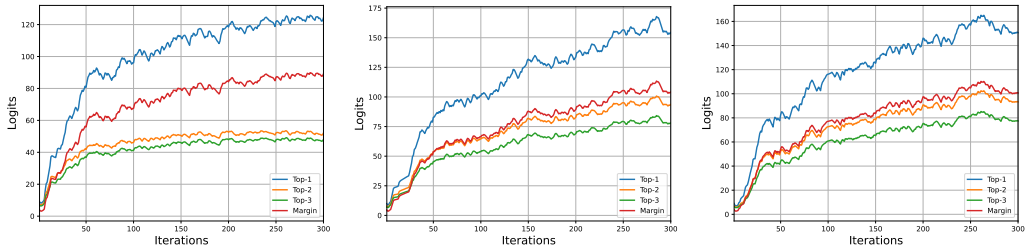

(a) T=5 (b) T=20 (c) Margin

Figure 4: The average Top-3 logits and logit margin of 50 adversarial samples after the logit calibration for crafting the ResNet-50.

To answer this, we propose to downscale the $z_t - z_{nt}$ by a factor s in the CE and extent the informative optimization for more iterations. Since in such circumstance, $z_t - z_{nt}$ will be enlarger by the factor s. Specifically, suppose that the optimization will be saturated when $z_t - z_{nt}$ reaches a certain value v. Using $z_t - z_{nt}$ and $\frac{z_t - z_{nt}}{s}$ in the CE will both approach the saturated value of v. Then, it is easy to infer that, for the latter case, $z_t - z_{nt}$ will be $v \times s$.

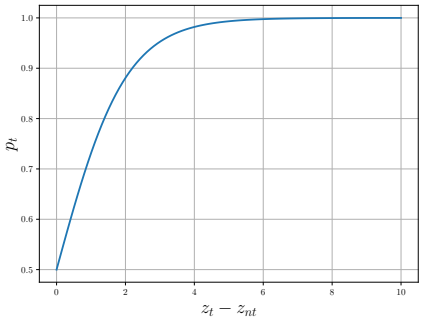

Figure 3: The probability of p_t under different $z_t - z_{nt}$.

3.4 Calibrating the Logits

To downscale the $z_t - z_{nt}$ during the optimization, we investigate three different types of logit calibrations in this study, *i.e.*, Temperature-based, Margin-based, and Angle-based.

3.4.1 Temperature-based

Inspired by the Temperature-scaling used in the Knowledge distillation [8], our first logit calibration directly downscale the logits by a constant temperature factor T,

$$z_i = \frac{z_i}{T}. \tag{14}$$

After introducing the T, the probability distribution $\boldsymbol{p}$ will be more softer over different classes. The corresponding gradient can be compute by:

$$\begin{aligned} \frac{\partial L_{ce}^T}{\partial \phi(\hat{x})} &= \frac{e^{z_j/T}}{\sum e^{z_j/T}} \cdot \frac{1}{T}\left(\frac{\partial z_j}{\partial x} - \frac{\partial z_t}{\partial \hat{x}}\right) \\ &= \sum -\hat{p}_i \frac{(W_t - W_i)}{T}. \end{aligned} \tag{15}$$

The $\hat{p}_i$ will not quickly approach to 0 after only a few iterations.

In Figure 4 (a)(b), we visualized the logits of using $T = 5$ and $T = 20$. We can find that targeted logits and the logit margin will keep increasing as the same as the Logit in Figure 1. Meanwhile, the trend of $T = 20$ is very similar with the Logit [30] and we show that the Logit loss function can be considered as a special case of calibrating the logits with a large T in the supplementary.

3.4.2 Margin-based

The previous Temperature-based logit calibration contains a hype-parameter T, which could be different for different surrogate model $\mathbb{F}_s$. To migrate this issue, we further introduce an adaptive margin-based logit calibration. Specifically, we calibrate the logits by using the margin between the Top-2 logits in each iteration, denoted as:

$$\tilde{z}_i = \frac{z_i}{\hat{z}_1 - \hat{z}_2}, \tag{16}$$

where $\hat{z}_1$ and $\hat{z}_2$ are the Top-1 and the Top-2 logit, respectively.

In this Margin-based logit calibration, we will enforce the p_t and $p_{\hat{1}}$ of the Top-1 non-target class at each iteration meet the following constraints:

$$p_t = \frac{1}{1 + \sum_{i \neq t} e^{-(\tilde{z}_t - \tilde{z}_i)}} < \frac{1}{1 + e^{-1}}, \tag{17}$$

$$p_{\hat{1}} = \frac{1}{e^{\tilde{z}_{\hat{1}} - \tilde{z}_t} + \sum_{i \neq t} e^{\tilde{z}_i - \tilde{z}_{\hat{1}}}} > \frac{1}{N-1}(1 - \frac{1}{1 + e^{-1}}). \tag{18}$$

Then, it can adaptively deal with the vanishing gradient issue in the original CE loss function. The logits and the margin is shown in Figure 4 (c).

3.4.3 Angle-based

On the other aspect, different W_t usually has a different norm. To further alleviate the influence of various norms, we calibrate the logit into the feature space with unit length by L2-normalization, $\frac{W_i^T \phi(\hat{x}) + b_i}{||W_i|| ||\phi(x)||}$. If omit the b_i, this calibration is compute the $cos(\theta)$ between $\phi(\hat{x})$ and W_i, and we term it as angle-based calibration. *Since, this angle-based calibration will bound each logit smaller than one.* Instead of using the CE loss function, we directly minimize the angle between the $\phi(\hat{x})$ and the targeted W_t. The optimization loss function is:

$$L_{cosine} = -\frac{W_t^T \phi(\hat{x})}{||W_t|| ||\phi(\hat{x})||}. \tag{19}$$

The angle-based classifiers have been widely using in Face-Recognition task [17, 2]. In the experiments, we evaluate the performance of using different logit calibration methods and also evaluate their mutual benefits.

4 Experiments

4.1 Experimental Setup

In this section, we evaluate the effectiveness of logit calibration for improving transferable targeted adversarial attack. Following the recent study [30], we conduct the experiments on the difficult ImageNet-Compatible Dataset[1]. This dataset contains 1,000 images with 1,000 unique class labels corresponded to the ImageNet dataset. We implement our methods based on the source code[2] provided by the Logit [30]. The same four diverse CNN models are used for evaluation, *i.e.*, ResNet-50 [7], DenseNet-121 [9], VGG-16 with Batch Normalization [24] and Inception-v3 [26]. The perturbation is bounded by $L_\infty \leq 16$. The TI [4], MI [3] and DI [29] were used for all attacks, and $||W||_1 = 5$ is set for TI. The I-FSGM is adopted for optimization with the $\alpha = 2$. The attacks were trained for 300 iterations on a NVIDIA-2080 Ti GPU. We run all the experiments for 5 times, and report the average number of successfully attacked samples.

4.2 Comparison with other methods in single-model transfer

We first compare the proposed (temperature-based, margin-based and angle-based) logit calibration methods with the original CE, Po+Trip [15], and Logit [30] in the single-model transfer task. In this task, we take one model as the surrogate model for training, and test the targeted transferability in attacking other 3 models.

As shown in Table 1, the original CE loss function produces a worst performance than the Po+Trip and Logit. But after performing the logit calibration in the CE loss function, we can find a significant performance boost compared with the original CE. All the calibration methods can outperform the Logit, especially when using the ResNet50 and Dense121 as the surrogate. These results indicate that the logit margin can significantly influence the performance of the targeted transferability. On the

[1] https://github.com/cleverhans-lab/cleverhans/tree/master/cleverhans_v3.1.0/examples/nips17_adversarial_competition/dataset

[2] https://github.com/ZhengyuZhao/Targeted-Tansfer

Table 1: The average number (#) of successfully attacked targeted samples in the single-model transfer scenario. (# with 20/100/300 iterations are reported, and the highest one at 300 iterations is showed **bold**).

Attack	Surrogate Model: ResNet50			Surrogate Model: Dense121		
	→Dense121	→VGG16	→Inc-v3	→Res50	→VGG16	→Inc-v3
CE	269.6/401.6/427	173.8/276.2/291	23.2/40.8/45.8	122.6/171.8/184	85.8/104.6/109	15.6/23/28
Po+Trip	278.8/511/8/548	178.8/354.8/346.8	32.4/67.6/78	110.4/148.4/150	73.2/92.2/85.6	16.4/27.6/27.6
Logit	313.8/640/717.8	238.4/550.2/623.6	30.6/86/109.4	174/386/435.4	136.8/337.8/378	23.4/65.8/75.4
T=5	332.6/699.4/**778.2**	248.2/599/661.2	30.6/93.6/**121.8**	192.8/434/475	145.8/365.6/394.2	23/72.6/88
T= 10	316/685.4/770	236/584.8/**664.4**	27.8/94/116	178.8/431.8/**492.6**	134.2/367.8/**415.4**	21.8/77.4/88.4
Margin	332.8/658.4/764.8	230.6/585.6/657.2	29.6/94.6/121.6	187.8/428/472.4	145/365.2/414.4	25.4/78.6/**93.6**
Angle	389.2/724.6/771.8	292/607.4/651.6	43.6/107.2/110.6	206.2/431.6/477.6	165/356.8/392.6	30.4/77.2/89
Attack	Surrogate Model: VGG16			Surrogate Model: Inc-v3		
	→Res50	→Dense121	→Inc-v3	→Res50	→Dense121	→VGG16
CE	4.6/2.6/5.6	5.6/3/3.2	0.2/0.2/0.4	6.8/11.6/18	6/13/18.6	4/8.2/13.2
Po+Trip	6.8/6.2/7.2	6.6/5.8/5	0.4/0.4/0.4	9.6/15.8/16.8	6.4/17.4/24.6	7.2/12/18.2
Logit	34.4/98.6/115.6	35.4/119.6/138.6	3.2/10.2/**13.2**	5.8/11.4/20.4	6.4/19.4/29.6	5.6/15.4/27.6
T=5	31/69.6/69	32.8/76/78	2.2/8.6/8.2	6.8/17/20.6	4.6/19/**32.8**	4/15.8/25.6
T= 10	35.8/90.2/97.4	33.6/105/117.4	3.2/511.4/12.6	5.4/12.6/19.4	5.8/20/26.8	4.2/15.2/**28**
Margin	32.8/102.8/**120.4**	34.6/124.6/**145**	2.6/11/**13.2**	5.4/14.2/17	7.2/21/31.2	5.2/17/26.8
Angle	4.4/6.6/5.4	5.6/4.4/4.8	0/0/0.4	7.6/18.4/**26**	8.2/22/29.6	9/17.2/24.2

Table 2: The average number (#) of successfully attacked targeted samples by using different T in CE loss function. (# with 20/100/300 iterations are reported).

Attack	Surrogate Model: ResNet50			Surrogate Model: Dense121		
	→Dense121	→VGG16	→Inc-v3	→Res50	→VGG16	→Inc-v3
T=0.5	131.6/160.2/194.8	71.2/94.8/110	11.6/18.2/24.4	41.6/50.4/61.6	24.6/34.6/32.4	6.2/9.2/11.4
T=1	269.6/401.6/427	173.8/276.2/291	23.2/40.8/45.8	122.6/171.8/184	85.8/104.6/109	15.6/23/28
T=2	341.8/628.4/676.8	244.2/522.8/538.6	33/72.4/85	187.2/349.8/360.8	132/273.2/270	22/55.4/61.4
T=5	332.6/699.4/**778.2**	248.2/599/661.2	30.6/93.6/**121.8**	192.8/434/475	145.8/365.6/394.2	23/72.6/88
T= 10	316/685.4/770	236/584.8/**664.4**	27.8/94/116	178.8/431.8/**492.6**	134.2/367.8/**415.4**	21.8/77.4/**88.4**
T = 20	303.8/655.8/742.8	229/554.4/636.4	32.4/89.8/116	176.2/402.8/461.6	134/353.6/400.6	22.8/67.4/87.2
Attack	Surrogate Model: VGG16			Surrogate Model: Inc-v3		
	→Res50	→Dense121	→Inc-v3	→Res50	→Dense121	→VGG16
T=0.5	2/1/2	0.8/0.6/1	0/0/0	3/9.4/9.2	3.4/8.2/14.4	3.2/5.6/13
T=1	4.6/2.6/5.6	5.6/3/3.2	0.2/0.2/0.4	6.8/11.6/18	6/13/18.6	4/8.2/13.2
T=2	16/18.2/18.2	17.8/19/15.6	2/1.8/2.4	6.2/15.4/20	4/17.2/21.8	4.6/11.6/20.2
T=5	31/69.6/69	32.8/76/78	2.2/8.6/8.2	6.8/17/20.6	4.6/19/32.8	4/15.8/25.6
T= 10	35.8/90.2/97.4	33.6/105/117.4	3.2/511.4/12.6	5.4/12.6/19.4	5.8/20/26.8	4.2/15.2/28
T = 20	33.8/97.2/**111.2**	35.8/126.6/**137.6**	2.6/12/2/**13.4**	5.4/14.4/**22.6**	5.8/18.2/**30.8**	5.2/15.8/**23.8**

other aspect, we find that $T = 10$ has better performance than $T = 5$ on the VGG-16, suggesting that different models may need different T. Instead of finding the best T for a different model, the Margin-based calibration can solve the issue and reach the overall best transferability in all four models. However, we find that the Angle-based calibration is not working on the VGG16, which needs further investigation.

4.3 The influence of different T in CE

In this section, we evaluate the influence of using different T in the CE loss function. The results are reported in Table 2. From the Table, we can have the following observations. **1)** The scaling factor T has a significant influence on the targeted transferability. Specifically, there is a large decrease in performance when using a small $T = 0.5$. After increasing the T, we can observe the number of successfully attacked samples will increase. **2)** The optimal T for different model is different. For example, $T = 5$ can produce the overall best performance for ResNet50, Dense121, and Inception v3, while the VGG16 requires a large T to obtain better transferability. **3)** The performance are comparable when using $T = 5$ and $T = 10$ for ResNet50, Dense121, and Inception v3. This is because that we use I-FSGM for optimization, which only considers the sign of the gradients. **4)** Using a larger T, the performance will be similar as the Logit loss function (see Table 1). We provide a deep analysis of this phenomenon in the supplementary material.

4.4 The mutual benefits of different calibration methods

In this part, we evaluate the mutual benefits of combining different calibrations and can have the following findings. **1)** Combining the T=5 and Margin, there is no increase in performance compared

Table 3: The comparison of combining logit calibrations. (# with 20/100/300 iterations are reported)

Attack	Surrogate Model: ResNet50			Surrogate Model: Dense121		
	→Dense121	→VGG16	→Inc-v3	→Res50	→VGG16	→Inc-v3
T=5 + Margin	338.2/698.4/772	239.6/590/655.4	33.4/96/111	192.6/442.6/477.8	141.2/377.4/408.4	25.4/74.8/93.6
T=5 + Angle	345.2/742.6/**823.8**	256.2/664.8/**721.6**	35.8/104.6/**131.4**	202.6/526.6/**619.2**	158.2/450.2/**536.4**	23.4/92/**127.2**
Margin + Angle	344/708.4/781.4	242.6/601.8/673.8	35/103.6/125.8	198.8/465.6/527.4	152.2/392.8/445.4	27/82/99
Attack	Surrogate Model: VGG16			Surrogate Model: Inc-v3		
	→Res50	→Dense121	→Inc-v3	→Res50	→Dense121	→VGG16
T=5 + Margin	34.8/101.8/**114**	37.2/123.6/145.6	3/10.8/13	4.8/14.4/16	6.2/21.2/28.6	5/17.2/**28.4**
T=5 + Angle	21.6/25/23.4	23.6/25.6/23.2	1.6/1.4/1.6	5.2/16/24/4	5.8/19.6/**31.2**	5.4/16.6/24.6
Margin + Angle	33/98.4/111.2	35.4/126.4/**146**	2.6/12/**14**	6.4/14/**21**	5.6/17/**31.2**	5.4/15.4/26.2

Table 4: **Comparison with TTP [21] on Target Transferablity.** The averaged Top-1 targeted accuracy (%) across 10 targets are computed with 49.95K ImageNet validation samples. Perturbation budget: $l_\infty \leq 16$. * indicates the training surrogate model.

Dataset	Loss	ResNet50*	$VGG19_{BN}$	Dense121	ResNet152	WRN-50-2	Average
ImageNet	TTP	97.02*	78.15	81.64	80.56	78.25	83.12
	CE	97.15*	70.44	78.96	76.22	78.24	80.20
	CE (T=5)	**99.18***	**86.65**	**90.55**	**90.30**	**93.22**	**91.98**
Painting	TTP	96.63*	73.09	84.76	76.27	75.92	81.33
	CE (T=5)	**98.95***	**82.97**	**87.07**	**87.81**	**91.70**	**89.70**

with using one of them. This is because that the gradient directions of these two methods are very similar. **2)** Combining the T=5 and Angle, we can observe a further improvement when using ResNet50 and Dense121 as the surrogate model, *e.g.*, the transferable rate of "ResNet50 → Dense121" is increased to 823.8 with 300 iterations. Since Angle has poor performance on ResNet 50, the transferable rates of corresponding combinations are also low. **3)** Combining the Margin and Angle, there are only slight improvements on ResNet50 and Dense121, while it can alleviate the negative effects caused by the angle-based calibration.

4.5 Comparison with the TTP method

In this section, we further evaluate the proposed temperature-based logit calibration in the GAN-based targeted attacks. Following the setting in TTP [21], we sampled 50K images from the ImageNet training set and 50K images from the Painting dataset[3], which are used to train the targeted generators from different source domains. Instead of using the distribution matching and neighborhood similarity matching loss [21], we only use the cross-entropy function for training the targeted generators while keeping other settings identical. More training and evaluation details used by TTP can be referred to [21]. We used the ResNet50 as the surrogate model and reported the results in Table 4.

From Table 4, we make the following findings. **1)** By using ImageNet as the training dataset, the TTP shows better transferability than the CE in attacking other black-box models. The average targeted accuracy of TTP is around 3% higher than that of CE. **2)** After downscale the logit by 5 in the CE loss function (CE (T=5)), we can observe a significant boost of the Top-1 targeted accuracy for all models, reaching the average targeted accuracy of 91.98% (ImageNet). **3)** For both ImageNet and Painting as the training source, the CE (T=5) can surpass the TTP by a large margin (91.98% vs. 83.12% & 89.70% vs. 81.33%). These experimental results demonstrate that the proposed temperate-based logit calibration is also effective in training generator-based targeted attackers. Note that, compared to TPP, our logit calibration has the benefit of without using any data from the target class.

5 Conclusion

In this study, we analyzed the logit margin in different loss functions for the transferable targeted attack, and find that the margin will quickly get saturated in the CE loss and thus limited the transferablity. To deal with this issue, we introduce to use logit calibrations in the CE loss function, including Temperature-based, Margin-based, and Angle-based. Experimental results verified the effectiveness of using the logit calibration in the CE loss function for crafting transferable targeted adversarial samples. The proposed logit calibration methods are simple and easy to implement, which can achieve state-of-the-art performance in transferable targeted attack.

[3] https://www.kaggle.com/c/painter-by-numbers

References

[1] Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *ICML*, 2019.

[2] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019.

[3] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *CVPR*, 2018.

[4] Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Evading defenses to transferable adversarial examples by translation-invariant attacks. In *CVPR*, 2019.

[5] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015.

[6] Yiwen Guo, Qizhang Li, and Hao Chen. Backpropagating linearly improves transferability of adversarial examples. In *NeurIPS*, 2020.

[7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

[8] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network (2015). *arXiv preprint arXiv:1503.02531*, 2015.

[9] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, 2017.

[10] Qian Huang, Isay Katsman, Horace He, Zeqi Gu, Serge Belongie, and Ser-Nam Lim. Enhancing adversarial example transferability with an intermediate level attack. In *CVPR*, 2019.

[11] Nathan Inkawhich, Kevin Liang, Binghui Wang, Matthew Inkawhich, Lawrence Carin, and Yiran Chen. Perturbing across the feature hierarchy to improve standard and strict blackbox attack transferability. *NeruIPS*, 2020.

[12] Nathan Inkawhich, Kevin J Liang, Lawrence Carin, and Yiran Chen. Transferable perturbations of deep feature distributions. *arXiv preprint arXiv:2004.12519*, 2020.

[13] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, 2016.

[14] Alexey Kurakin, Ian J Goodfellow, and Samy Bengio. Adversarial examples in the physical world. In *Artificial intelligence safety and security*, pages 99–112. Chapman and Hall/CRC, 2018.

[15] Maosen Li, Cheng Deng, Tengjiao Li, Junchi Yan, Xinbo Gao, and Heng Huang. Towards transferable targeted attack. In *CVPR*, 2020.

[16] Jiadong Lin, Chuanbiao Song, Kun He, Liwei Wang, and John E Hopcroft. Nesterov accelerated gradient and scale invariance for adversarial attacks. 2020.

[17] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphereface: Deep hypersphere embedding for face recognition. In *CVPR*, 2017.

[18] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into transferable adversarial examples and black-box attacks. In *ICLR*, 2016.

[19] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.

[20] Muhammad Muzammal Naseer, Salman H Khan, Muhammad Haris Khan, Fahad Shahbaz Khan, and Fatih Porikli. Cross-domain transferability of adversarial perturbations. *NeurIPS*, 2019.

[21] Muzammal Naseer, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Fatih Porikli. On generating transferable targeted perturbations. In *ICCV*, 2021.

[22] Omid Poursaeed, Isay Katsman, Bicheng Gao, and Serge Belongie. Generative adversarial perturbations. In *CVPR*, 2018.

[23] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NeurIPS*, 2015.

[24] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.

[25] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.

[26] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016.

[27] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. Ensemble adversarial training: Attacks and defenses. In *ICLR*, 2018.

[28] Dongxian Wu, Yisen Wang, Shu-Tao Xia, James Bailey, and Xingjun Ma. Skip connections matter: On the transferability of adversarial examples generated with resnets. In *ICLR*, 2020.

[29] Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, and Alan L Yuille. Improving transferability of adversarial examples with input diversity. In *CVPR*, 2019.

[30] Zhengyu Zhao, Zhuoran Liu, and Martha Larson. On success and simplicity: A second look at transferable targeted attacks. *NeurIPS*, 34, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [No]
   (c) Did you discuss any potential negative societal impacts of your work? [No] No negative effects
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See abstract.
   (b) Did you specify all the training details (*e.g.*, data splits, hyperparameters, how they were chosen)? [Yes] See Section 4.1.
   (c) Did you report error bars (*e.g.*, with respect to the random seed after running experiments multiple times)? [No] We report the average of running experiments 5 times.
   (d) Did you include the total amount of compute and the type of resources used (*e.g.*, type of GPUs, internal cluster, or cloud provider)? [Yes] See URL in abstract
4. If you are using existing assets (*e.g.*, code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No] The code and data are public

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No] No use of personally identifiable information or objectionable content.

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]