



Square Attack: A Query-Efficient Black-Box Adversarial Attack via Random Search

Maksym Andriushchenko^{1(✉)}, Francesco Croce², Nicolas Flammarion¹,
and Matthias Hein²

¹ EPFL, Lausanne, Switzerland
maksym.andriushchenko@epfl.ch

² University of Tübingen, Tübingen, Germany

Abstract. We propose the *Square Attack*, a score-based black-box l_2 - and l_∞ -adversarial attack that does not rely on local gradient information and thus is not affected by gradient masking. Square Attack is based on a randomized search scheme which selects localized square-shaped updates at random positions so that at each iteration the perturbation is situated approximately at the boundary of the feasible set. Our method is significantly more query efficient and achieves a higher success rate compared to the state-of-the-art methods, especially in the untargeted setting. In particular, on ImageNet we improve the average query efficiency in the untargeted setting for various deep networks by a factor of at least 1.8 and up to 3 compared to the recent state-of-the-art l_∞ -attack of Al-Dujaili & O’Reilly (2020). Moreover, although our attack is *black-box*, it can also outperform gradient-based *white-box* attacks on the standard benchmarks achieving a new state-of-the-art in terms of the success rate. The code of our attack is available at <https://github.com/max-andr/square-attack>.

1 Introduction

Adversarial examples are of particular concern when it comes to applications of machine learning which are safety-critical. Many defenses against adversarial examples have been proposed [1, 5, 7, 23, 32, 40, 56] but with limited success, as new more powerful attacks could break many of them [4, 10, 12, 35, 57]. In particular, gradient obfuscation or masking [4, 35] is often the reason why seemingly robust models turn out to be non-robust in the end. Gradient-based attacks are most often affected by this phenomenon (white-box attacks but also black-box attacks based on finite difference approximations [35]). Thus it is important to

M. Andriushchenko and F. Croce—Equal contribution.

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-58592-1_29) contains supplementary material, which is available to authorized users.

have attacks which are based on different principles. Black-box attacks have recently become more popular [8, 36, 46] as their attack strategies are quite different from the ones employed for adversarial training, where often PGD-type attacks [32] are used. However, a big weakness currently is that these black-box attacks need to query the classifier too many times before they find adversarial examples, and their success rate is sometimes significantly lower than that of white-box attacks.

In this paper we propose Square Attack, a score-based adversarial attack, i.e. it can query the probability distribution over the classes predicted by a classifier but has no access to the underlying model. The Square Attack exploits random search¹ [41, 43] which is one of the simplest approaches for black-box optimization. Due to a particular sampling distribution, it requires significantly fewer queries compared to the state-of-the-art black-box methods (see Fig. 1) in the score-based threat model while outperforming them in terms of *success rate*, i.e. the percentage of successful adversarial examples. This is achieved by a combination of a particular initialization strategy and our square-shaped updates. We motivate why these updates are particularly suited to attack neural networks and provide convergence guarantees for a variant of our method. In an extensive evaluation with untargeted and targeted attacks, three datasets (MNIST, CIFAR-10, ImageNet), normal and robust models, we show that Square Attack outperforms state-of-the-art methods in the l_2 - and l_∞ -threat model.




Fig. 1. Avg. number of queries of successful untargeted l_∞ -attacks on three ImageNet models for three score-based black-box attacks. Square Attack outperforms all other attacks by large margin

2 Related Work

We discuss black-box attacks with l_2 - and l_∞ -perturbations since our attack focuses on this setting. Although attacks for other norms, e.g. l_0 , exist [16, 36], they are often algorithmically different due to the geometry of the perturbations.

l_2 - and l_∞ -Score-Based Attacks. Score-based black-box attacks have only access to the score predicted by a classifier for each class for a given input. Most of such attacks in the literature are based on gradient estimation through finite differences. The first papers in this direction [6, 27, 51] propose attacks which approximate the gradient by sampling from some noise distribution around the point. While this approach can be successful, it requires many queries of the classifier, particularly in high-dimensional input spaces as in image classification. Thus, improved techniques reduce the dimension of the search space via using the principal components of the data [6], searching for perturbations in

¹ It is an iterative procedure different from random sampling inside the feasible region.

the latent space of an auto-encoder [50] or using a low-dimensional noise distribution [28]. Other attacks exploit evolutionary strategies or random search, e.g. [3] use a genetic algorithm to generate adversarial examples and alleviate gradient masking as they can reduce the robust accuracy on randomization- and discretization-based defenses. The l_2 -attack of [25] can be seen as a variant of random search which chooses the search directions in an orthonormal basis and tests up to two candidate updates at each step. However, their algorithm can have suboptimal query efficiency since it adds at every step only small (in l_2 -norm) modifications, and suboptimal updates cannot be undone as they are orthogonal to each other. A recent line of work has pursued black-box attacks which are based on the observation that successful adversarial perturbations are attained at corners of the l_∞ -ball intersected with the image space $[0, 1]^d$ [2, 34, 44]. Searching over the corners allows to apply discrete optimization techniques to generate adversarial attacks, significantly improving the query efficiency. Both [44] and [2] divide the image according to some coarse grid, perform local search in this lower dimensional space allowing componentwise changes only of $-\epsilon$ and ϵ , then refine the grid and repeat the scheme. In [2] such a procedure is motivated as an estimation of the gradient signs. Recently, [34] proposed several attacks based on evolutionary algorithms, using discrete and continuous optimization, achieving nearly state-of-the-art query efficiency for the l_∞ -norm. In order to reduce the dimensionality of the search space, they use the “tiling trick” of [28] where they divide the perturbation into a set of squares and modify the values in these squares with evolutionary algorithms. A related idea also appeared earlier in [22] where they introduced black rectangle-shaped perturbations for generating adversarial occlusions. In [34], as in [28], both size and position of the squares are fixed at the beginning and not optimized. Despite their effectiveness for the l_∞ -norm, these discrete optimization based attacks are not straightforward to adapt to the l_2 -norm. Finally, approaches based on Bayesian optimization exist, e.g. [45], but show competitive performance only in a low-query regime.

Different Threat and Knowledge Models. We focus on l_p -norm-bounded adversarial perturbations (for other perturbations such as rotations, translations, occlusions in the black-box setting see, e.g., [22]). Perturbations with *minimal* l_p -norm are considered in [13, 50] but require significantly more queries than norm-bounded ones. Thus we do not compare to them, except for [25] which has competitive query efficiency while aiming at small perturbations.

In other cases the attacker has a different knowledge of the classifier. A more restrictive scenario, considered by *decision-based* attacks [8, 9, 11, 14, 24], is when the attacker can query only the decision of the classifier, but not the predicted scores. Other works use more permissive threat models, e.g., when the attacker already has a substitute model similar to the target one [15, 20, 39, 47, 52] and thus can generate adversarial examples for the substitute model and then transfer them to the target model. Related to this, [52] suggest to refine this approach by running a black-box gradient estimation attack in a subspace spanned by the gradients of substitute models. However, the gain in query efficiency given by

such extra knowledge does not account for the computational cost required to train the substitute models, particularly high on ImageNet-scale. Finally, [31] use extra information on the target data distribution to train a model that predicts adversarial images that are then refined by gradient estimation attacks.

3 Square Attack

In the following we recall the definitions of the adversarial examples in the threat model we consider and present our black-box attacks for the l_∞ - and l_2 -norms.

Algorithm 1: The Square Attack via random search

Input: classifier f , point $x \in \mathbb{R}^d$, image size w , number of color channels c , l_p -radius ϵ , label $y \in \{1, \dots, K\}$, number of iterations N
Output: approximate minimizer $\hat{x} \in \mathbb{R}^d$ of the problem stated in Eq. (1)

- 1 $\hat{x} \leftarrow \text{init}(x)$, $l^* \leftarrow L(f(x), y)$, $i \leftarrow 1$
- 2 **while** $i < N$ **and** \hat{x} is not adversarial **do**
- 3 $h^{(i)} \leftarrow$ side length of the square to modify (according to some schedule)
- 4 $\delta \sim P(\epsilon, h^{(i)}, w, c, \hat{x}, x)$ (see Alg. 2 and 3 for the sampling distributions)
- 5 $\hat{x}_{\text{new}} \leftarrow$ Project $\hat{x} + \delta$ onto $\{z \in \mathbb{R}^d : \|z - x\|_p \leq \epsilon\} \cap [0, 1]^d$
- 6 $l_{\text{new}} \leftarrow L(f(\hat{x}_{\text{new}}), y)$
- 7 **if** $l_{\text{new}} < l^*$ **then** $\hat{x} \leftarrow \hat{x}_{\text{new}}$, $l^* \leftarrow l_{\text{new}}$;
- 8 $i \leftarrow i + 1$
- 9 **end**

3.1 Adversarial Examples in the l_p -threat Model

Let $f : [0, 1]^d \rightarrow \mathbb{R}^K$ be a classifier, where d is the input dimension, K the number of classes and $f_k(x)$ is the predicted score that x belongs to class k . The classifier assigns class $\arg \max_{k=1, \dots, K} f_k(x)$ to the input x . The goal of an *untargeted*

attack is to change the correctly predicted class y for the point x . A point \hat{x} is called an *adversarial example* with an l_p -norm bound of ϵ for x if

$$\arg \max_{k=1, \dots, K} f_k(\hat{x}) \neq y, \quad \|\hat{x} - x\|_p \leq \epsilon \quad \text{and} \quad \hat{x} \in [0, 1]^d,$$

where we have added the additional constraint that \hat{x} is an image. The task of finding \hat{x} can be rephrased as solving the constrained optimization problem

$$\min_{\hat{x} \in [0, 1]^d} L(f(\hat{x}), y), \quad \text{s.t.} \quad \|\hat{x} - x\|_p \leq \epsilon, \quad (1)$$

for a loss L . In our experiments, we use $L(f(\hat{x}), y) = f_y(\hat{x}) - \max_{k \neq y} f_k(\hat{x})$.

The goal of *targeted* attacks is instead to change the decision of the classifier to a particular class t , i.e., to find \hat{x} so that $\arg \max_k f_k(\hat{x}) = t$ under the same constraints on \hat{x} . We further discuss the targeted attacks in Sup. E.1.

3.2 General Algorithmic Scheme of the Square Attack

Square Attack is based on *random search* which is a well known iterative technique in optimization introduced by Rastrigin in 1963 [41]. The main idea of the algorithm is to sample a random update δ at each iteration, and to add this update to the current iterate \hat{x} if it improves the objective function. Despite its simplicity, random search performs well in many situations [54] and does not depend on gradient information from the objective function g .

Many variants of random search have been introduced [33, 42, 43], which differ mainly in how the random perturbation is chosen at each iteration (the original scheme samples uniformly on a hypersphere of fixed radius). For our goal of crafting adversarial examples we come up with two sampling distributions specific to the l_∞ - and the l_2 -attack (Sect. 3.3 and Sect. 3.4), which we integrate in the classic random search procedure. These sampling distributions are motivated by both how images are processed by neural networks with convolutional filters and the shape of the l_p -balls for different p . Additionally, since the considered objective is non-convex when using neural networks, a good initialization is particularly important. We then introduce a specific one for better query efficiency.

Our proposed scheme differs from classical random search by the fact that the perturbations $\hat{x} - x$ are constructed such that for every iteration they lie on the boundary of the l_∞ - or l_2 -ball before projection onto the image domain $[0, 1]^d$. Thus we are using the perturbation budget almost maximally at each step. Moreover, the changes are localized in the image in the sense that at each step we modify just a small fraction of contiguous pixels shaped into **squares**. Our overall scheme is presented in Algorithm 1. First, the algorithm picks the side length $h^{(i)}$ of the square to be modified (step 3), which is decreasing according to an a priori fixed schedule. This is in analogy to the step-size reduction in gradient-based optimization. Then in step 4 we sample a new update δ and add it to the current iterate (step 5). If the resulting loss (obtained in step 6) is smaller than the best loss so far, the change is accepted otherwise discarded. Since we are interested in query efficiency, the algorithm stops as soon as an adversarial example is found. The time complexity of the algorithm is dominated by the evaluation of $f(\hat{x}_{\text{new}})$, which is performed at most N times, with N total number of iterations. We plot the resulting adversarial perturbations in Fig. 3 and additionally in Sup. E where we also show imperceptible perturbations.

We note that previous works [28, 34, 44] generate perturbations containing squares. However, while those use a fixed grid on which the squares are constrained, we optimize the position of the squares as well as the color, making our attack more flexible and effective. Moreover, unlike previous works, we motivate squared perturbations with the structure of the convolutional filters (see Sect. 4).

Size of the Squares. Given images of size $w \times w$, let $p \in [0, 1]$ be the percentage of elements of x to be modified. The length h of the side of the squares used is given by the closest positive integer to $\sqrt{p \cdot w^2}$ (and $h \geq 3$ for the l_2 -attack). Then, the initial p is the only free parameter of our scheme. With $N = 10000$

iterations available, we halve the value of p at $i \in \{10, 50, 200, 1000, 2000, 4000, 6000, 8000\}$ iterations. For different N we rescale the schedule accordingly.

3.3 The l_∞ -Square Attack

Initialization. As initialization we use vertical stripes of width one where the color of each stripe is sampled uniformly at random from $\{-\epsilon, \epsilon\}^c$ (c number of color channels). We found that convolutional networks are particularly sensitive to such perturbations, see also [53] for a detailed discussion on the sensitivity of neural networks to various types of high frequency perturbations.

Sampling Distribution. Similar to [44] we observe that successful l_∞ -perturbations usually have values $\pm\epsilon$ in all the components (note that this does not hold perfectly due to the image constraints $\hat{x} \in [0, 1]^d$). In particular, it holds

$$\hat{x}_i \in \{\max\{0, x_i - \epsilon\}, \min\{1, x_i + \epsilon\}\}.$$

Our sampling distribution P for the l_∞ -norm described in Algorithm 2 selects sparse updates of \hat{x} with $\|\delta\|_0 = h \cdot h \cdot c$ where $\delta \in \{-2\epsilon, 0, 2\epsilon\}^d$ and the non-zero elements are grouped to form a square. In this way, after the projection onto the l_∞ -ball of radius ϵ (Step 5 of Algorithm 1) all components i for which $\epsilon \leq x_i \leq 1 - \epsilon$ satisfy $\hat{x}_i \in \{x_i - \epsilon, x_i + \epsilon\}$, i.e. differ from the original point x in each element either by ϵ or $-\epsilon$. Thus $\hat{x} - x$ is situated at one of the corners of the l_∞ -ball (modulo the components which are close to the boundary). Note that all projections are done by clipping. Moreover, we fix the elements of δ belonging to the same color channel to have the same sign, since we observed that neural networks are particularly sensitive to such perturbations (see Sect. 4.3).

Algorithm 2: Sampling distribution P for l_∞ -norm

Input: maximal norm ϵ , window size h , image size w , color channels c
Output: New update δ

- 1 $\delta \leftarrow$ array of zeros of size $w \times w \times c$
- 2 sample uniformly
 $r, s \in \{0, \dots, w - h\} \subset \mathbb{N}$
- 3 **for** $i = 1, \dots, c$ **do**
- 4 $\rho \leftarrow \text{Uniform}(\{-2\epsilon, 2\epsilon\})$
- 5 $\delta_{r+1:r+h, s+1:s+h, i} \leftarrow \rho \cdot \mathbb{1}_{h \times h}$
- 6 **end**

3.4 The l_2 -Square Attack

Initialization. The l_2 -perturbation is initialized by generating a 5×5 grid-like tiling by squares of the image, where the perturbation on each tile has the shape described next in the sampling distribution. The resulting perturbation $\hat{x} - x$ is rescaled to have l_2 -norm ϵ and the resulting \hat{x} is projected onto $[0, 1]^d$ by clipping.




Fig. 2. Perturbation of the l_2 -attack

Sampling Distribution. First, let us notice that the adversarial perturbations typically found for the l_2 -norm tend to be much more localized than those for the l_∞ -norm [49], in the sense that large changes are applied on some pixels of the original image, while many others are minimally modified. To mimic this feature we introduce a new update η which has two “centers” with large absolute value and opposite signs, while the other components have lower absolute values as one gets farther away from the centers, but never reaching zero (see Fig. 2 for one example with $h = 8$ of the resulting update η). In this way the modifications are localized and with high contrast between the different halves, which we found to improve the query efficiency. Concretely, we define $\eta^{(h_1, h_2)} \in \mathbb{R}^{h_1 \times h_2}$ (for some $h_1, h_2 \in \mathbb{N}_+$ such that $h_1 \geq h_2$) for every $1 \leq r \leq h_1, 1 \leq s \leq h_2$ as

$$\eta_{r,s}^{(h_1, h_2)} = \sum_{k=0}^{M(r,s)} \frac{1}{(n+1-k)^2}, \text{ with } n = \left\lfloor \frac{h_1}{2} \right\rfloor,$$

and $M(r, s) = n - \max\{|r - \lfloor \frac{h_1}{2} \rfloor - 1|, |s - \lfloor \frac{h_2}{2} \rfloor - 1|\}$. The intermediate square update $\eta \in \mathbb{R}^{h \times h}$ is then selected uniformly at random from either

$$\eta = \left(\eta^{(h,k)}, -\eta^{(h,h-k)} \right), \text{ with } k = \lfloor h/2 \rfloor, \tag{2}$$

or its transpose (corresponding to a rotation of 90°).

Algorithm 3: Sampling distribution P for l_2 -norm

Input: maximal norm ϵ , window size h , image size w , number of color channels c , current image \hat{x} , original image x

Output: New update δ

```

1  $\nu \leftarrow \hat{x} - x$ 
2 sample uniformly  $r_1, s_1, r_2, s_2 \in \{0, \dots, w - h\}$ 
3  $W_1 := r_1 + 1 : r_1 + h, s_1 + 1 : s_1 + h, W_2 := r_2 + 1 : r_2 + h, s_2 + 1 : s_2 + h$ 
4  $\epsilon_{\text{unused}}^2 \leftarrow \epsilon^2 - \|\nu\|_2^2, \eta^* \leftarrow \nu / \|\eta\|_2$  with  $\eta$  as in (2)
5 for  $i = 1, \dots, c$  do
6    $\rho \leftarrow \text{Uniform}(\{-1, 1\})$ 
7    $\nu_{\text{temp}} \leftarrow \rho \eta^* + \nu_{W_1, i} / \|\nu_{W_1, i}\|_2$ 
8    $\epsilon_{\text{avail}}^i \leftarrow \sqrt{\|\nu_{W_1 \cup W_2, i}\|_2^2 + \epsilon_{\text{unused}}^2 / c}$ 
9    $\nu_{W_2, i} \leftarrow 0, \nu_{W_1, i} \leftarrow (\nu_{\text{temp}} / \|\nu_{\text{temp}}\|_2) \epsilon_{\text{avail}}^i$ 
10 end
11  $\delta \leftarrow x + \nu - \hat{x}$ 

```

Second, unlike l_∞ -constraints, l_2 -constraints do not allow to perturb each component independently from the others as the overall l_2 -norm must be kept smaller than ϵ . Therefore, to modify a perturbation $\hat{x} - x$ of norm ϵ with localized changes while staying on the hypersphere, we have to “move the mass” of $\hat{x} - x$ from one location to another. Thus, our scheme consists in randomly selecting two squared windows in the current perturbation $\nu = \hat{x} - x$, namely ν_{W_1} and




Fig. 3. Visualization of the adversarial perturbations and examples found by the l_∞ - and l_2 -versions of the Square Attack on ResNet-50

ν_{W_2} , setting $\nu_{W_2} = 0$ and using the budget of $\|\nu_{W_2}\|_2$ to increase the total perturbation of ν_{W_1} . Note that the perturbation of W_1 is then a combination of the existing perturbation plus the new generated η . We report the details of this scheme in Algorithm 3 where step 4 allows to utilize the budget of l_2 -norm lost after the projection onto $[0, 1]^d$. The update δ output by the algorithm is such that the next iterate $\hat{x}_{\text{new}} = \hat{x} + \delta$ (before projection onto $[0, 1]^d$ by clipping) belongs to the hypersphere $B_2(x, \epsilon)$ as stated in the following proposition.

Proposition 1. *Let δ be the output of Algorithm 3. Then $\|\hat{x} + \delta - x\|_2 = \epsilon$.*

4 Theoretical and Empirical Justification of the Method

We provide high-level theoretical justifications and empirical evidence regarding the algorithmic choices in Square Attack, with focus on the l_∞ -version (the l_2 -version is significantly harder to analyze).

4.1 Convergence Analysis of Random Search

First, we want to study the convergence of the random search algorithm when considering an L -smooth objective function g (such as neural networks with activation functions like softplus, swish, ELU, etc.) on the whole space \mathbb{R}^d (without projection²) under the following assumptions on the update δ_t drawn from the sampling distribution P_t :

$$\mathbb{E}\|\delta_t\|_2^2 \leq \gamma_t^2 C \text{ and } \mathbb{E}\langle \delta_t, v \rangle \geq \tilde{C}\gamma_t\|v\|_2, \forall v \in \mathbb{R}^d, \quad (3)$$

where γ_t is the step size at iteration t , $C, \tilde{C} > 0$ some constants and $\langle \cdot, \cdot \rangle$ denotes the inner product. We obtain the following result, similar to existing convergence rates for zeroth-order methods [21, 37, 38]:

Proposition 2. *Suppose that $\mathbb{E}[\delta_t] = 0$ and the assumptions in Eq. (3) hold. Then for step-sizes $\gamma_t = \gamma/\sqrt{T}$, we have*

$$\min_{t=0, \dots, T} \mathbb{E}\|\nabla g(x_t)\|_2 \leq \frac{2}{\gamma\tilde{C}\sqrt{T}} \left(g(x_0) - \mathbb{E}g(x_{T+1}) + \frac{\gamma^2 CL}{2} \right).$$

² Nonconvex constrained optimization under noisy oracles is notoriously harder [19].

This basically shows for T large enough one can make the gradient arbitrary small, meaning that the random search algorithm converges to a critical point of g (one cannot hope for much stronger results in non-convex optimization without stronger conditions).

Unfortunately, the second assumption in Eq. (3) does not directly hold for our sampling distribution P for the l_∞ -norm (see Sup. A.3), but holds for a similar one, P^{multiple} , where each component of the update δ is drawn uniformly at random from $\{-2\epsilon, 2\epsilon\}$. In fact we show in Sup. A.4, using the Khintchine inequality [26], that

$$\mathbb{E}\|\delta_t\|_2^2 \leq 4c\epsilon^2h^2 \text{ and } \mathbb{E}|\langle \delta_t, v \rangle| \geq \frac{\sqrt{2}c\epsilon h^2}{d} \|v\|_2, \forall v \in \mathbb{R}^d.$$

Moreover, while P^{multiple} performs worse than the distribution used in Algorithm 2, we show in Sect. 4.3 that it already reaches state-of-the-art results.

4.2 Why Squares?

Previous works [34, 44] build their l_∞ -attacks by iteratively adding square modifications. Likewise we change square-shaped regions of the image for both our l_∞ - and l_2 -attacks—with the difference that we can sample any square subset of the input, while the grid of the possible squares is fixed in [34, 44]. This leads naturally to wonder why squares are superior to other shapes, e.g. rectangles.

Let us consider the l_∞ -threat model, with bound ϵ , input space $\mathbb{R}^{d \times d}$ and a convolutional filter $w \in \mathbb{R}^{s \times s}$ with entries unknown to the attacker. Let $\delta \in \mathbb{R}^{d \times d}$ be the sparse update with $\|\delta\|_0 = k \geq s^2$ and $\|\delta\|_\infty \leq \epsilon$. We denote by $S(a, b)$ the index set of the rectangular support of δ with $|S(a, b)| = k$ and shape $a \times b$. We want to provide intuition why sparse square-shaped updates are superior to rectangular ones in the sense of reaching a maximal change in the activations of the first convolutional layer.

Let $z = \delta * w \in \mathbb{R}^{d \times d}$ denote the output of the convolutional layer for the update δ . The l_∞ -norm of z is the maximal componentwise change of the convolutional layer:

$$\begin{aligned} \|z\|_\infty &= \max_{u,v} |z_{u,v}| = \max_{u,v} \left| \sum_{i,j=1}^s \delta_{u-\lfloor \frac{s}{2} \rfloor+i, v-\lfloor \frac{s}{2} \rfloor+j} \cdot w_{i,j} \right| \\ &\leq \max_{u,v} \epsilon \sum_{i,j} |w_{i,j}| \mathbb{1}_{(u-\lfloor \frac{s}{2} \rfloor+i, v-\lfloor \frac{s}{2} \rfloor+j) \in S(a,b)}, \end{aligned}$$

where elements with indices exceeding the size of the matrix are set to zero. Note that the indicator function attains 1 only for the non-zero elements of δ involved in the convolution to get $z_{u,v}$. Thus, to have the largest upper bound possible on $|z_{u,v}|$, for some (u, v) , we need the largest possible amount of components of δ with indices in

$$C(u, v) = \left\{ (u - \lfloor \frac{s}{2} \rfloor + i, v - \lfloor \frac{s}{2} \rfloor + j) : i, j = 1, \dots, s \right\}$$

to be non-zero (that is in $S(a, b)$).

Therefore, it is desirable to have the shape $S(a, b)$ of the perturbation δ selected so to maximize the number N of convolutional filters $w \in \mathbb{R}^{s \times s}$ which fit into the rectangle $a \times b$. Let \mathcal{F} be the family of the objects that can be defined as the union of axis-aligned rectangles with vertices on \mathbb{N}^2 , and $\mathcal{G} \subset \mathcal{F}$ be the squares of \mathcal{F} of shape $s \times s$ with $s \geq 2$. We have the following proposition:

Proposition 3. *Among the elements of \mathcal{F} with area $k \geq s^2$, those which contain the largest number of elements of \mathcal{G} have*

$$N^* = (a - s + 1)(b - s + 1) + (r - s + 1)^+ \quad (4)$$

of them, with $a = \lfloor \sqrt{k} \rfloor$, $b = \lfloor \frac{k}{a} \rfloor$, $r = k - ab$ and $z^+ = \max\{z, 0\}$.

This proposition states that, if we can modify only k elements of δ , then shaping them to form (approximately) a square allows to maximize the number of pairs (u, v) for which $|S(a, b) \cap C(u, v)| = s^2$. If $k = l^2$ then $a = b = l$ are the optimal values for the shape of the perturbation update, i.e. the shape is exactly a square.

Table 1. Ablation study of the l_∞ -Square Attack which shows how the individual design decisions improve the performance. The fourth row corresponds to the method for which we have shown convergence guarantees in Sect. 4.1. The last row corresponds to our final l_∞ -attack. c indicates the number of color channels, h the length of the side of the squares, so that “# random sign” c represents updates with constant sign for each color, while $c \cdot h^2$ updates with signs sampled independently of each other

Update shape	# random signs	Initialization	Failure rate	Avg. queries	Median queries
random	$c \cdot h^2$	vert. stripes	0.0%	401	48
random	$c \cdot h^2$	uniform rand.	0.0%	393	132
random	c	vert. stripes	0.0%	339	53
square	$c \cdot h^2$	vert. stripes	0.0%	153	15
rectangle	c	vert. stripes	0.0%	93	16
square	c	uniform rand.	0.0%	91	26
square	c	vert. stripes	0.0%	73	11

4.3 Ablation Study

We perform an ablation study to show how the individual design decisions for the sampling distribution of the random search improve the performance of l_∞ -Square Attack, confirming the theoretical arguments above. The comparison is done for an l_∞ -threat model of radius $\epsilon = 0.05$ on 1,000 test points for a ResNet-50 model trained normally on ImageNet (see Sect. 5 for details) with a query limit of 10,000 and results are shown in Table 1. Our sampling distribution is special in two aspects: i) we use localized update shapes in form of squares and ii) the update is constant in each color channel. First, one can observe that our

update shape “square” performs better than “rectangle” as we discussed in the previous section, and it is significantly better than “random” (the same amount of pixels is perturbed, but selected randomly in the image). This holds both for c (constant sign per color channel) and $c \cdot h^2$ (every pixel and color channel is changed independently of each other), with an improvement in terms of average queries of 339 to 73 and 401 to 153 respectively. Moreover, with updates of the same shape, the constant sign over color channels is better than selecting it uniformly at random (improvement in average queries: 401 to 339 and 153 to 73). In total the algorithm with “square- c ” needs more than $5\times$ less average queries than “random- $c \cdot h^2$ ”, showing that our sampling distribution is the key to the high query efficiency of Square Attack.

The last innovation of our random search scheme is the initialization, crucial element of every non-convex optimization algorithm. Our method (“square- c ”) with the vertical stripes initialization improves over a uniform initialization on average by $\approx 25\%$ and, especially, median number of queries (more than halved).

We want to also highlight that the sampling distribution “square- $c \cdot h^2$ ” for which we shown convergence guarantees in Sect. 4.1 performs already better in terms of the success rate and the median number of queries than the state of the art (see Sect. 5). For a more detailed ablation, also for our l_2 -attack, see Sup. C.

Table 2. Results of **untargeted** attacks on ImageNet with a limit of 10,000 queries. For the l_∞ -attack we set the norm bound $\epsilon = 0.05$ and for the l_2 -attack $\epsilon = 5$. Models: normally trained **I**: Inception v3, **R**: ResNet-50, **V**: VGG-16-BN. The Square Attack outperforms for both threat models all other methods in terms of success rate and query efficiency. The missing entries correspond to the results taken from the original paper where some models were not reported

Norm	Attack	Failure rate			Avg. queries			Med. queries		
		I	R	V	I	R	V	I	R	V
l_∞	Bandits [28]	3.4%	1.4%	2.0%	957	727	394	218	136	36
	Parsimonious [44]	1.5%	–	–	722	–	–	237	–	–
	DFO $_c$ -CMA [34]	0.8%	0.0%	0.1%	630	270	219	259	143	107
	DFO $_d$ -Diag. CMA [34]	2.3%	1.2%	0.5%	424	417	211	20	20	2
	SignHunter [2]	1.0%	0.1%	0.3%	471	129	95	95	39	43
	Square attack	0.3%	0.0%	0.0%	197	73	31	24	11	1
l_2	Bandits [28]	9.8%	6.8%	10.2%	1486	939	511	660	392	196
	SimBA-DCT [25]	35.5%	12.7%	7.9%	651	582	452	564	467	360
	Square attack	7.1%	0.7%	0.8%	1100	616	377	385	170	109

5 Experiments

In this section we show the effectiveness of the Square Attack. Here we concentrate on **untargeted** attacks since our primary goal is query efficient robustness evaluation, while the **targeted** attacks are postponed to the supplement. First,

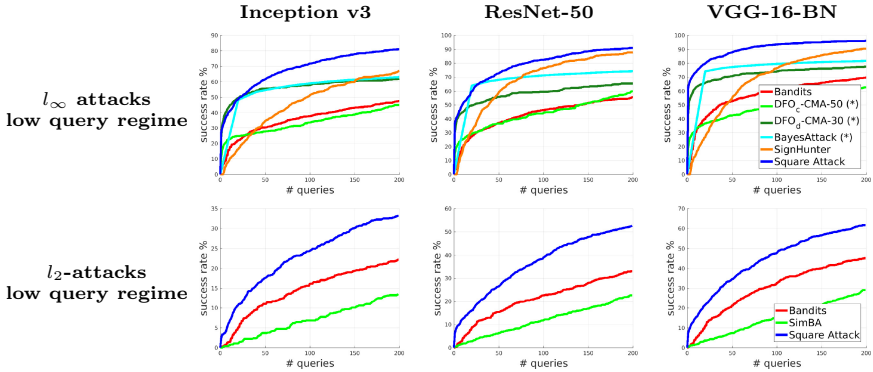


Fig. 4. Success rate in the low-query regime (up to 200 queries). * denotes the results obtained via personal communication with the authors and evaluated on 500 and 10,000 randomly sampled points for BayesAttack [45] and DFO [34] methods, respectively

we follow the standard setup [28,34] of comparing black-box attacks on three ImageNet models in terms of success rate and query efficiency for the l_∞ - and l_2 -untargeted attacks (Sect. 5.1). Second, we show that our *black-box* attack can even outperform *white-box* PGD attacks on several models (Sect. 5.2). Finally, in the supplement we provide more experimental details (Sup. B), a stability study of our attack for different parameters (Sup. C) and random seeds (Sup. D), and additional results including the experiments for targeted attacks (Sup. E).

5.1 Evaluation on ImageNet

We compare the Square Attack to state-of-the-art score-based black-box attacks (without any extra information such as surrogate models) on three pretrained models in PyTorch (Inception v3, ResNet-50, VGG-16-BN) using 1,000 images from the ImageNet validation set. Unless mentioned otherwise, we use the code from the other papers with their suggested parameters. As it is standard in the literature, we give a budget of 10,000 queries per point to find an adversarial perturbation of l_p -norm at most ϵ . We report the *average* and *median* number of queries each attack requires to craft an adversarial example, together with the *failure rate*. All query statistics are computed only for successful attacks on the points which were originally correctly classified.

Tables 2 and 3 show that the Square Attack, despite its simplicity, achieves in all the cases (models and norms) the **lowest failure rate**, (<1% everywhere except for the l_2 -attack on Inception v3), and almost always requires **fewer queries** than the competitors to succeed. Figure 4 shows the progression of the success rate of the attacks over the first 200 queries. Even in the low query regime the Square Attack outperforms the competitors for both norms. Finally, we highlight that the only hyperparameter of our attack, p , regulating the size of the squares, is set for all the models to 0.05 for l_∞ and 0.1 for l_2 -perturbations.

l_∞ -Attacks. We compare our attack to Bandits [29], Parsimonious [44], DFO_c/DFO_d [34], and SignHunter [2]. In Table 2 we report the results of the l_∞ -attacks with norm bound of $\epsilon = 0.05$. The Square Attack always has the lowest failure rate, notably 0.0% in 2 out of 3 cases, and the lowest query consumption. Interestingly, our attack has median equal 1 on VGG-16-BN, meaning that the proposed initialization is particularly effective for this model.

The closest competitor in terms of the *average* number of queries is SignHunter [2], which still needs on average between 1.8 and 3 times more queries to find adversarial examples and has a higher failure rate than our attack. Moreover, the median number of queries of SignHunter is much worse than for our method (e.g. 43 vs 1 on VGG). We note that although DFO_c-CMA [34] is competitive to our attack in terms of *median* queries, it has a significantly higher failure rate and between 2 and 7 times worse average number of queries. Additionally, our method is also more effective in the low-query regime (Fig. 4) than other methods (including [45]) on all the models.

Table 3. Query statistics for untargeted l_2 -attacks computed for the points for which all three attacks are successful for fair comparison

Attack	Avg. queries			Med. queries		
	I	R	V	I	R	V
Bandits [28]	536	635	398	368	314	177
SimBA-DCT [25]	647	563	421	552	446	332
Square attack	352	287	217	181	116	80

l_2 -Attacks. We compare our attack to Bandits [28] and SimBA [25] for $\epsilon = 5$, while we do not consider SignHunter [2] since it is not as competitive as for the l_∞ -norm, and in particular worse than Bandits on ImageNet (see Fig. 2 in [2]).

As Table 2 and Fig. 4 show, the Square Attack outperforms by a large margin the other methods in terms of failure rate, and achieves the lowest median number of queries for all the models and the lowest average one for VGG-16-BN. However, since it has a significantly lower failure rate, the statistics of the Square Attack are biased by the “hard” cases where the competitors fail. Then, we recompute the same statistics considering only the points where all the attacks are successful (Table 3). In this case, our method improves by at least $1.5\times$ the average and by at least $2\times$ the median number of queries.

5.2 Square Attack Can Be More Accurate Than White-Box Attacks

Here we test our attack on problems which are challenging for both white-box PGD and other black-box attacks. We use for evaluation *robust accuracy*, defined as the worst-case accuracy of a classifier when an attack perturbs each input in some l_p -ball. We show that our algorithm outperforms the competitors both on

Table 4. On the robust models of [32] and [55] on MNIST l_∞ -Square Attack with $\epsilon = 0.3$ achieves state-of-the-art (SOTA) results outperforming white-box attacks

Model	Robust accuracy	
	SOTA	Square
Madry et al. [32]	88.13%	88.25%
TRADES [55]	93.33%	92.58%

state-of-the-art robust models and defenses that induce different types of gradient masking. Thus, our attack is useful to evaluate robustness without introducing adaptive attacks designed for each model separately.

Outperforming White-Box Attacks on Robust Models. The models obtained with the adversarial training of [32] and TRADES [55] are standard benchmarks to test adversarial attacks, which means that many papers have tried to reduce their robust accuracy, without limit on the computational budget and primarily via white-box attacks. We test our l_∞ -Square Attack on these robust models on MNIST at $\epsilon = 0.3$, using $p = 0.8$, 20k queries and 50 random restarts, i.e., we run our attack 50 times and consider it successful if any of the runs finds an adversarial example (Table 4). On the model of Madry et al. [32] Square Attack is only 0.12% far from the *white-box* state-of-the-art, achieving the second best result (also outperforming the 91.47% of SignHunter [2] by a large margin). On the TRADES benchmark [57], our method obtains a new SOTA of 92.58% robust accuracy outperforming the white-box attack of [17]. Additionally, the subsequent work of [18] uses the Square Attack as part of their *AutoAttack* where they show that the Square Attack outperforms other white-box attacks on 9 out of 9 MNIST models they evaluated. Thus, our black-box attack can be also useful for robustness evaluation of new defenses in the setting where gradient-based attacks require many restarts and iterations.

Resistance to Gradient Masking. In Table 5 we report the robust accuracy at different thresholds ϵ of the l_∞ -adversarially trained models on MNIST of [32] for the l_2 -threat model. It is known that the PGD is ineffective since it suffers from gradient masking [48]. Unlike PGD and other black-box attacks, our Square Attack does not suffer from gradient masking and yields robust accuracy close to zero for $\epsilon = 2.5$, with only a single run. Moreover, the l_2 -version of SignHunter [2] fails to accurately assess the robustness because the method optimizes only over the extreme points of the l_∞ -ball of radius ϵ/\sqrt{d} embedded in the target l_2 -ball.

Table 5. l_2 -robustness of the l_∞ -adversarially trained models of [32] at different thresholds ϵ . PGD is shown with 1, 10, 100 random restarts. The black-box attacks are given a 10k queries budget (see the supplement for details)

ϵ_2	Robust accuracy						
	White-box			Black-box			
	PGD ₁	PGD ₁₀	PGD ₁₀₀	SignHunter	Bandits	SimBA	Square
2.0	79.6%	67.4%	59.8%	95.9%	80.1%	87.6%	16.7%
2.5	69.2%	51.3%	36.0%	94.9%	32.4%	75.8%	2.4%
3.0	57.6%	29.8%	12.7%	93.8%	12.5%	58.1%	0.6%

Table 6. l_∞ -robustness of Clean Logit Pairing (CLP), Logit Squeezing (LSQ) [30]. The Square Attack is competitive to white-box PGD with many restarts (R = 10,000, R = 100 on MNIST, CIFAR-10 resp.) and more effective than black-box attacks [2, 28]

ϵ_∞	Model	Robust accuracy				
		White-box		Black-box		
		PGD ₁	PGD _R	Bandits	SignHunter	Square
0.3	CLP _{MNIST}	62.4%	4.1%	33.3%	62.1%	6.1%
	LSQ _{MNIST}	70.6%	5.0%	37.3%	65.7%	2.6%
16/255	CLP _{CIFAR}	2.8%	0.0%	14.3%	0.1%	0.2%
	LSQ _{CIFAR}	27.0%	1.7%	27.7%	13.2%	7.2%

Attacking Clean Logit Pairing and Logit Squeezing. These two l_∞ defenses proposed in [30] were broken in [35]. However, [35] needed up to 10k restarts of PGD which is computationally prohibitive. Using the publicly available models from [35], we run the Square Attack with $p = 0.3$ and 20k query limit (results in Table 6). We obtain robust accuracy similar to PGD_R in most cases, but with a *single run*, i.e. without additional restarts. At the same time, although on some models Bandits and SignHunter outperform PGD₁, they on average achieve significantly worse results than the Square Attack. This again shows the utility of the Square Attack to accurately assess robustness.

6 Conclusion

We have presented a simple black-box attack which outperforms by a large margin the state-of-the-art both in terms of query efficiency and success rate. Our results suggest that our attack is useful *even in comparison to white-box attacks* to better estimate the robustness of models that exhibit gradient masking.

Acknowledgements. We thank L. Meunier and S. N. Shukla for providing the data for Fig. 6. M.A. thanks A. Modas for fruitful discussions. M.H and F.C. acknowledge support by the Tue.AI Center (FKZ: 01IS18039A), DFG TRR 248, project number 389792660 and DFG EXC 2064/1, project number 390727645.

References

1. Akhtar, N., Mian, A.: Threat of adversarial attacks on deep learning in computer vision: a survey. *IEEE Access* **6**, 14410–14430 (2018)
2. Al-Dujaili, A., O’Reilly, U.M.: There are no bit parts for sign bits in black-box attacks. In: *ICLR* (2020)
3. Alzantot, M., Sharma, Y., Chakraborty, S., Srivastava, M.: GenAttack: practical black-box attacks with gradient-free optimization. In: *Genetic and Evolutionary Computation Conference (GECCO)* (2019)

4. Athalye, A., Carlini, N., Wagner, D.A.: Obfuscated gradients give a false sense of security: circumventing defenses to adversarial examples. In: ICML (2018)
5. Bastani, O., Ioannou, Y., Lampropoulos, L., Vytiniotis, D., Nori, A., Criminisi, A.: Measuring neural net robustness with constraints. In: NeurIPS (2016)
6. Bhagoji, A.N., He, W., Li, B., Song, D.: Practical black-box attacks on deep neural networks using efficient query mechanisms. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11216, pp. 158–174. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01258-8_10
7. Biggio, B., Roli, F.: Wild patterns: ten years after the rise of adversarial machine learning. *Pattern Recogn.* **84**, 317–331 (2018)
8. Brendel, W., Rauber, J., Bethge, M.: Decision-based adversarial attacks: reliable attacks against black-box machine learning models. In: ICLR (2018)
9. Brunner, T., Diehl, F., Le, M.T., Knoll, A.: Guessing smart: biased sampling for efficient black-box adversarial attacks. In: ICCV (2019)
10. Carlini, N., Wagner, D.: Adversarial examples are not easily detected: bypassing ten detection methods. In: ACM Workshop on Artificial Intelligence and Security (2017)
11. Chen, J., Jordan, M.I., J., W.M.: HopSkipJumpAttack: a query-efficient decision-based attack (2019). arXiv preprint [arXiv:1904.02144](https://arxiv.org/abs/1904.02144)
12. Chen, P., Sharma, Y., Zhang, H., Yi, J., Hsieh, C.: EAD: elastic-net attacks to deep neural networks via adversarial examples. In: AAAI (2018)
13. Chen, P.Y., Zhang, H., Sharma, Y., Yi, J., Hsieh, C.J.: ZOO: zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In: 10th ACM Workshop on Artificial Intelligence and Security - AISec 2017. ACM Press (2017)
14. Cheng, M., Le, T., Chen, P.Y., Yi, J., Zhang, H., Hsieh, C.J.: Query-efficient hard-label black-box attack: an optimization-based approach. In: ICLR (2019)
15. Cheng, S., Dong, Y., Pang, T., Su, H., Zhu, J.: Improving black-box adversarial attacks with a transfer-based prior. In: NeurIPS (2019)
16. Croce, F., Hein, M.: Sparse and imperceivable adversarial attacks. In: ICCV (2019)
17. Croce, F., Hein, M.: Minimally distorted adversarial examples with a fast adaptive boundary attack. In: ICML (2020)
18. Croce, F., Hein, M.: Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In: ICML (2020)
19. Davis, D., Drusvyatskiy, D.: Stochastic model-based minimization of weakly convex functions. *SIAM J. Optim.* **29**(1), 207–239 (2019)
20. Du, J., Zhang, H., Zhou, J.T., Yang, Y., Feng, J.: Query-efficient meta attack to deep neural networks. In: ICLR (2020)
21. Duchi, J., Jordan, M., Wainwright, M., Wibisono, A.: Optimal rates for zero-order convex optimization: the power of two function evaluations. *IEEE Trans. Inf. Theory* **61**(5), 2788–2806 (2015)
22. Fawzi, A., Frossard, P.: Measuring the effect of nuisance variables on classifiers. In: British Machine Vision Conference (BMVC) (2016)
23. Gu, S., Rigazio, L.: Towards deep neural network architectures robust to adversarial examples. In: ICLR Workshop (2015)
24. Guo, C., Frank, J.S., Weinberger, K.Q.: Low frequency adversarial perturbation. In: UAI (2019)
25. Guo, C., Gardner, J.R., You, Y., Wilson, A.G., Weinberger, K.Q.: Simple black-box adversarial attacks. In: ICML (2019)
26. Haagerup, U.: The best constants in the Khintchine inequality. *Studia Math.* **70**(3), 231–283 (1981)

27. Ilyas, A., Engstrom, L., Athalye, A., Lin, J.: Black-box adversarial attacks with limited queries and information. In: ICML (2018)
28. Ilyas, A., Engstrom, L., Madry, A.: Prior convictions: black-box adversarial attacks with bandits and priors. In: ICLR (2019)
29. Ilyas, A., Santurkar, S., Tsipras, D., Engstrom, L., Tran, B., Madry, A.: Adversarial examples are not bugs, they are features. In: NeurIPS (2019)
30. Kannan, H., Kurakin, A., Goodfellow, I.: Adversarial logit pairing (2018). arXiv preprint [arXiv:1803.06373](https://arxiv.org/abs/1803.06373)
31. Li, Y., Li, L., Wang, L., Zhang, T., Gong, B.: NATTACK: learning the distributions of adversarial examples for an improved black-box attack on deep neural networks. In: ICML (2019)
32. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. In: ICLR (2018)
33. Matyas, J.: Random optimization. *Autom. Remote Control* **26**(2), 246–253 (1965)
34. Meunier, L., Atif, J., Teytaud, O.: Yet another but more efficient black-box adversarial attack: tiling and evolution strategies (2019). arXiv preprint, [arXiv:1910.02244](https://arxiv.org/abs/1910.02244)
35. Mosbach, M., Andriushchenko, M., Trost, T., Hein, M., Klakow, D.: Logit pairing methods can fool gradient-based attacks. In: NeurIPS 2018 Workshop on Security in Machine Learning (2018)
36. Narodytska, N., Kasiviswanathan, S.: Simple black-box adversarial attacks on deep neural networks. In: CVPR Workshops (2017)
37. Nemirovsky, A.S., Yudin, D.B.: Problem Complexity and Method Efficiency in Optimization. Wiley-Interscience Series in Discrete Mathematics. Wiley, Hoboken (1983)
38. Nesterov, Y., Spokoiny, V.: Random gradient-free minimization of convex functions. *Found. Comput. Math.* **17**(2), 527–566 (2017). <https://doi.org/10.1007/s10208-015-9296-2>
39. Papernot, N., McDaniel, P., Goodfellow, I.: Transferability in machine learning: from phenomena to black-box attacks using adversarial samples (2016). arXiv preprint [arXiv:1605.07277](https://arxiv.org/abs/1605.07277)
40. Papernot, N., McDaniel, P., Wu, X., Jha, S., Swami, A.: Distillation as a defense to adversarial perturbations against deep networks. In: IEEE Symposium on Security & Privacy (2016)
41. Rastrigin, L.: The convergence of the random search method in the extremal control of a many parameter system. *Autom. Remote Control* **24**, 1337–1342 (1963)
42. Schrack, G., Choit, M.: Optimized relative step size random searches. *Math. Program.* **10**, 230–244 (1976). <https://doi.org/10.1007/BF01580669>
43. Schumer, M., Steiglitz, K.: Adaptive step size random search. *IEEE Trans. Automat. Control* **13**(3), 270–276 (1968)
44. Seungyong, M., Gaon, A., Hyun, O.S.: Parsimonious black-box adversarial attacks via efficient combinatorial optimization. In: ICML (2019)
45. Shukla, S.N., Sahu, A.K., Willmott, D., Kolter, Z.: Black-box adversarial attacks with Bayesian optimization (2019). arXiv preprint [arXiv:1909.13857](https://arxiv.org/abs/1909.13857)
46. Su, J., Vargas, D., Sakurai, K.: One pixel attack for fooling deep neural networks. *IEEE Trans. Evol. Comput.* **23**, 828–841 (2019)
47. Suya, F., Chi, J., Evans, D., Tian, Y.: Hybrid batch attacks: finding black-box adversarial examples with limited queries (2019). arXiv preprint, [arXiv:1908.07000](https://arxiv.org/abs/1908.07000)
48. Tramèr, F., Boneh, D.: Adversarial training and robustness for multiple perturbations. In: NeurIPS (2019)

49. Tsipras, D., Santurkar, S., Engstrom, L., Turner, A., Madry, A.: Robustness may be at odds with accuracy. In: ICLR (2019)
50. Tu, C.C., et al.: Autozoom: autoencoder-based zeroth order optimization method for attacking black-box neural networks. In: AAAI Conference on Artificial Intelligence (2019)
51. Uesato, J., O'Donoghue, B., Van den Oord, A., Kohli, P.: Adversarial risk and the dangers of evaluating against weak attacks. In: ICML (2018)
52. Yan, Z., Guo, Y., Zhang, C.: Subspace attack: exploiting promising subspaces for query-efficient black-box attacks. In: NeurIPS (2019)
53. Yin, D., Lopes, R.G., Shlens, J., Cubuk, E.D., Gilmer, J.: A Fourier perspective on model robustness in computer vision. In: NeurIPS (2019)
54. Zabinsky, Z.B.: Random search algorithms. In: Wiley Encyclopedia of Operations Research and Management Science (2010)
55. Zhang, H., Yu, Y., Jiao, J., Xing, E.P., Ghaoui, L.E., Jordan, M.I.: Theoretically principled trade-off between robustness and accuracy. In: ICML (2019)
56. Zheng, S., Song, Y., Leung, T., Goodfellow, I.J.: Improving the robustness of deep neural networks via stability training. In: CVPR (2016)
57. Zheng, T., Chen, C., Ren, K.: Distributionally adversarial attack. In: AAAI (2019)