

# $\sigma$ -ZERO: GRADIENT-BASED OPTIMIZATION OF $\ell_0$ -NORM ADVERSARIAL EXAMPLES

Anonymous authors

Paper under double-blind review

## ABSTRACT

Evaluating the adversarial robustness of deep networks to gradient-based attacks is challenging. While most attacks focus  $\ell_2$ -norm and  $\ell_\infty$ -norm constraints to craft input perturbations, only a few have investigated sparse  $\ell_1$ -norm and  $\ell_0$ -norm attacks. In particular,  $\ell_0$ -norm attacks remain the least studied due to the inherent complexity of optimizing over a non-convex and non-differentiable constraint. However, evaluating the robustness of these attacks might unveil weaknesses otherwise left untested with conventional  $\ell_2$  and  $\ell_\infty$  attacks. In this work, we propose a novel  $\ell_0$ -norm attack, called  $\sigma$ -zero, which leverages an ad-hoc differentiable approximation of the  $\ell_0$  norm to facilitate gradient-based optimization. Extensive evaluations on MNIST, CIFAR10, and ImageNet datasets, involving robust and non-robust models, show that  $\sigma$ -zero can find adversarial examples with minimal  $\ell_0$  distortion, outperforming competing methods minimum  $\ell_0$ -norm adversarial examples without requiring any time-consuming hyperparameter tuning, and that it outperforms all competing attacks in terms of success rate and scalability.

## 1 INTRODUCTION

Early research has unveiled that Deep Neural Networks (DNNs) are fooled by adversarial examples, i.e., slightly-perturbed inputs optimized to cause misclassifications (Biggio et al., 2013; Szegedy et al., 2014a; Goodfellow et al., 2015). In turn, this has demanded the need for more careful reliability assessments of such models. Most of the gradient-based attacks proposed to evaluate adversarial robustness of DNNs focus on optimizing adversarial examples under different  $\ell_p$ -norm constraints. In particular, while convex  $\ell_1$ ,  $\ell_2$ , and  $\ell_\infty$  norms have been widely studied (Chen et al., 2018; Croce & Hein, 2021a), only a few  $\ell_0$ -norm attacks have been considered so far. The main reason may be is that ad-hoc heuristics need to be adopted to compute efficient projections on the  $\ell_0$  norm, overcoming issues related to its non-convexity and non-differentiability. Although this task is challenging and computationally expensive, attacks based on the  $\ell_0$  norm have the potential to reveal uncovered issues in DNNs that may not be evident in other norm-based attacks (Carlini & Wagner, 2017; Croce & Hein, 2021a). For instance, these attacks, known for perturbing a minimal fraction of input features, can be used to determine the most sensitive characteristics that disproportionately influence the model’s decision-making process. Furthermore, they offer a different and relevant threat model to inspect for benchmarking benchmark existing defenses. Consequently, developing efficient algorithms for generating  $\ell_0$  adversarial examples is thus a crucial area of research that requires further exploration for a more accurate robustness evaluation to improve current adversarial robustness evaluations.

However Unfortunately, current implementations of  $\ell_0$  attacks exhibit a largely suboptimal tradeoff between their success rate and efficiency, i.e., they are either accurate but slow, or fast but inaccurate. In particular, the accurate ones resort to the use of complex projections to find smaller input perturbations but suffer from time or memory limitations, hindering their scalability to larger networks or high-dimensional data (Brendel et al., 2019a; C  saire et al., 2021). Other attacks execute faster, but their output solution is typically inaccurate and largely suboptimal as they rely on heuristic approaches and imprecise approximations to bypass the difficulties of optimizing the  $\ell_0$  norm, leading to overestimating adversarial robustness (Matyasko & Chau, 2021; Pintor et al., 2021). However, all existing strategies are often slow to converge because they require a large number of queries

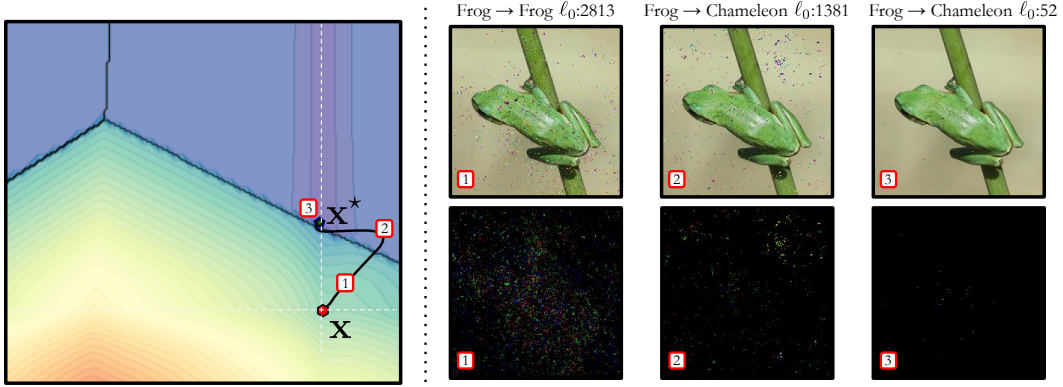


Figure 1: The leftmost plot illustrates shows an instance of  $\sigma$ -zero’s execution on a bi-dimensional two-dimensional problem. Here, the red dot represents and the green star respectively represent the initial point  $x$ , while and the green star indicates the corresponding adversarial example  $x^*$  that deceives the model. The Our gradient-based attack, powered by gradient descent, seeks to find this adversarial example while minimizing the number of perturbed features (i.e., the  $\ell_0$  norm of the perturbation). Gray lines surrounding  $x$  demarcate regions where the  $\ell_0$  distance between the adversarial example and  $x$  would be minimum norm is minimized. Complementary, the The rightmost plot showcases shows the adversarial examples generated images (top row) and the corresponding perturbations (bottom row) found by  $\sigma$ -zero during the three steps (top row) highlighted in the leftmost plot, alongside their prediction and  $\ell_0$  norm. The bottom row offers a look at the adversarial perturbations  $\delta$  yielded by.

(i.e., forward and backward passes), or they output suboptimal solutions. It thus remains an open challenge to develop a scalable and compelling method for assessing the robustness of DNNs against sparse perturbation attacks perturbations with minimum  $\ell_0$  norm.

To tackle these issues, in this work we propose a novel attack technique, namely  $\sigma$ -zero, which iteratively increases promotes the sparsity of the adversarial perturbations, minimizing their  $\ell_0$  norm (see Fig. 1 and Sect. 2). The underlying idea is to leverage an a differentiable approximation of the actual  $\ell_0$  norm, which is better suited to gradient-based optimizers. Specifically, we employ the approximation method initially introduced by Osborne et al. (2000b), and more recently adopted by Cinà et al. (2022) for staging energy-latency poisoning attacks. This method offers an unbiased, differentiable estimation of the true  $\ell_0$  norm and possesses the crucial property of differentiability, making it particularly well-suited for optimization techniques that rely on gradients, allowing us to optimize it via gradient descent.

Our experiments (Sect. 3) provide compelling evidence of the remarkable performance of our attack. We evaluate  $\sigma$ -zero on several benchmark datasets, including MNIST, CIFAR10, and Imagenet ImageNet, considering baseline and robust models from Robustbench (Croce et al., 2021). We compare its performance with state-of-the-art attacks, showing that  $\sigma$ -zero achieves better or comparable results in terms of attack success rate and perturbation size, while requiring significantly less computational resources being significantly faster and without requiring any sophisticated and time-consuming hyperparameter tuning. Overall, our approach encompasses two fundamental characteristics for a proficient adversarial attack, i.e., effectiveness and scalability, making it a catalyst for significant advancements in developing novel models with improved robustness, as well as better robustness evaluation tools.

## 2 $\sigma$ -ZERO: MINIMUM $\ell_0$ -NORM ADVERSARIAL EXAMPLES WITH DIFFERENTIABLE APPROXIMATION

Most of the existing attacks rely on heuristic approaches to bypass the difficulties of optimizing the norm. They either start with perturbing a small set of features and then enlarging this set until they find an adversarial example or, given a successful adversarial manipulation, try to cut the number of manipulated features

~~progressively (Carlini & Wagner, 2017; Pintor et al., 2021). Alternatively, some approaches resort to the use of imprecise approximations of the penalty-loss function or complex projection operators to run the attack (Chen et al., 2018; Matyasko & Chau, 2021; Brendel et al., 2019a). However, these strategies are often time-consuming because they require a large number of queries (i.e., forward and backward passes), or they output suboptimal solutions.~~

~~In the following, we present a~~ We present here  $\sigma$ -zero, our gradient-based approach ~~that exploits to finding minimum  $\ell_0$ -norm approximations. Starting from a general description of the adversarial examples. We start by describing the considered~~ threat model and ~~attack goal, we then~~ give a formal overview of the proposed attack and its algorithmic implementation.

**Threat Model.** We assume that the attacker has complete access to the target model, including its architecture and trained parameters, and exploits its gradient for staging white-box untargeted attacks. This setting is useful for worst-case evaluation of the adversarial robustness of DNN models, providing empirical upper bounds on the performance degradation that may incur when they are attacked, and it is the usual setting adopted also in previous work related to gradient-based adversarial robustness evaluations (Carlini & Wagner, 2017; Croce et al., 2021; Pintor et al., 2021).

**Problem Formulation.** In this work, we seek untargeted minimum  $\ell_0$ -norm adversarial perturbations that steer the model’s decision towards misclassification. To this end, let  $\mathbf{x} \in \mathcal{X} = [0, 1]^d$  be a  $d$ -dimensional input sample,  $y \in \mathcal{Y} = \{1, \dots, l\}$  its associated true label, and  $f : \mathcal{X} \times \Theta \mapsto \mathcal{Y}$  the target model, parameterized by  $\theta \in \Theta$ . While  $f$  outputs the predicted label, we will also use  $f_k$  to denote the continuous-valued output (logit) for class  $k \in \mathcal{Y}$ . The goal of our attack is to find the minimum  $\ell_0$ -norm adversarial noise  $\delta^*$  such that the corresponding adversarial example  $\mathbf{x}^* = \mathbf{x} + \delta^*$  is misclassified by  $f$ . This is formalized as the following optimization problem:

$$\delta^* \in \arg \min_{\delta} \quad \|\delta\|_0, \quad (1)$$

$$\text{s.t.} \quad f(\mathbf{x} + \delta, \theta) \neq y, \quad (2)$$

$$\mathbf{x} + \delta \in [0, 1]^d, \quad (3)$$

where  $\|\cdot\|_0$  denotes the  $\ell_0$  norm, which counts the number of non-zero dimensions. The hard-constraint in Equation 2 ensures that the perturbation  $\delta$  induces the target model  $f$  to misclassify the perturbed sample  $\mathbf{x}^*$ . Finally, Equation 3 represents a box constraint, ensuring that the adversarial example  $\mathbf{x}^*$  lies in  $[0, 1]^d$ . Note that when the source point  $\mathbf{x}$  is already misclassified by  $f$ , the trivial solution to the above minimization problem is  $\delta^* = \mathbf{0}$ .

Contrary to the  $\ell_1$ ,  $\ell_2$ ,  $\ell_\infty$  norms, when considering the  $\ell_0$  norm the problem becomes intractable with standard methods. The  $\ell_0$  norm is indeed non-differentiable, thus unsuitable for gradient-based optimization. To address this issue, we exploit the  $\ell_0$ -norm approximation function proposed by Osborne et al. (2000b), and defined as:

$$\hat{\ell}_0(\mathbf{x}) = \sum_{i=1}^d \frac{x_i^2}{x_i^2 + \sigma}, \quad \sigma > 0, \quad \hat{\ell}_0(\mathbf{x}) \in [0, d], \quad (4)$$

with  $\sigma$  being a hyperparameter controlling its approximation quality. When  $\sigma$  tends to zero, the approximation becomes more accurate. However, an increasingly accurate approximation could lead to the same optimization limits of the  $\ell_0$  norm.

Finally, similarly to previous work (Carlini & Wagner, 2017; Rony et al., 2021a; Szegedy et al., 2014b), we transform the hard-constraint in Equation 2 in a soft-constraint. The resulting optimization problem therefore becomes:

$$\delta^* \in \arg \min_{\delta} \quad \mathcal{L}(\mathbf{x} + \delta, y, \theta) + \frac{1}{d} \hat{\ell}_0(\delta) \quad (5)$$

$$\text{s.t.} \quad \mathbf{x} + \delta \in [0, 1]^d, \quad (6)$$

where we substituted the  $\|\delta\|_0$  with the approximation  $\hat{\ell}_0(\delta)$  and normalize it with respect to the number of features  $d$  to ensure that its value is within the interval  $[0, 1]$ . The loss  $\mathcal{L}$  is defined as:

$$\mathcal{L}(\mathbf{x}, y, \theta) = \max \left( f_y(\mathbf{x}, \theta) - \max_{k \neq y} f_k(\mathbf{x}, \theta), 0 \right) + \mathbb{I}(f(\mathbf{x}, \theta) = y). \quad (7)$$

**Algorithm 1:**  $\sigma$ -zero Attack Pseudocode.

---

**Input:**  $\mathbf{x} \in [0, 1]^d$ , input sample;  $y$ , true class label;  $\theta$ , target model;  $N$ , number of iterations;  
 $\sigma$ ,  $\ell_0$ -approximation parameter;  $\eta_0$ , initial step size;  $\tau_0$ , initial sparsity threshold.

**Output:**  $\mathbf{x}^*$ , minimum  $\ell_0$  norm adversarial example.

```

1  $\delta \leftarrow \mathbf{0}$ ;  $\delta^* \leftarrow \delta$ ;  $\tau \leftarrow \tau_0$ ;  $\eta \leftarrow \eta_0$  ▷ initialization.
2 for  $i$  in  $1, \dots, N$  do
3    $\nabla \mathbf{g} \leftarrow \nabla_{\delta} [\mathcal{L}(\mathbf{x} + \delta, y, \theta) + \frac{1}{d} \hat{\ell}_0(\delta, \sigma)]$  ▷ gradient computation.
4    $\nabla \mathbf{g} \leftarrow \nabla \mathbf{g} / \|\nabla \mathbf{g}\|_{\infty}$  ▷ gradient normalization.
5    $\delta \leftarrow \text{clip}(\mathbf{x} - [\delta - \eta \cdot \nabla \mathbf{g}]) - \mathbf{x}$  ▷  $\delta$  update.
6    $\delta \leftarrow \Pi_{\tau}(\delta)$  ▷ zeroing  $\delta$  components below  $\tau$ .
7    $\eta = \text{cosine\_annealing}(\eta_0, i)$  ▷  $\eta$  update.
8   If  $\mathcal{L}(\mathbf{x} + \delta, y, \theta) \leq 0$   $\tau += 0.01 \cdot \eta$  else  $\tau -= 0.01 \cdot \eta$  ▷  $\tau$  update.
9 end
10 If  $\mathcal{L}(\mathbf{x} + \delta, y, \theta) \leq 0 \wedge \|\delta\|_0 < \|\delta^*\|_0$   $\delta^* \leftarrow \delta$  ▷  $\delta^*$  update.
11 return  $\mathbf{x}^* \leftarrow \mathbf{x} + \delta^*$ 

```

---

The first term in  $\mathcal{L}$  represents the logit difference, which is positive when the sample is correctly assigned to the true class  $y$ , and clipped to zero when it is misclassified (Carlini & Wagner, 2017). The second term merely adds 1 to the loss if the sample is correctly classified.<sup>1</sup> This ensures that the loss term  $\mathcal{L}$  is 0 only when an adversarial example is found, and higher than 1 otherwise. This in turn implies that the loss term  $\mathcal{L}$  is always higher than the  $\ell_0$ -norm term in Equation 5 (as the latter is bounded in  $[0, 1]$ ), when no adversarial example is found. Accordingly, it is not difficult to see that the feasible solutions of this problem only correspond to minimum-norm adversarial examples. It is also worth remarking that, conversely to the objective function proposed by Carlini & Wagner (2017), our objective does not require tuning the tradeoff between minimizing the loss and reducing the perturbation size to find minimum-norm adversarial examples, thereby avoiding a computationally-expensive line search for each input sample. In fact, the proposed objective function inherently induces an *alternate* optimization process between the loss term and the  $\ell_0$ -norm penalty, as shown in the Appendix (see Figure 4). In particular, when the sample is not adversarial, the attack algorithm mostly aims to decrease the loss term to find an adversarial example, while increasing the perturbation size. Conversely, when an adversarial example is found, the loss term is cropped to zero, and the perturbation size is gradually reduced.

**Solution Algorithm.** Given that the approximation function  $\hat{\ell}_0$  in Equation 4 is differentiable, we derive a custom gradient-based algorithm for solving Equation 5 and Equation 6. Our ~~algorithm, given in attack, detailed in Algorithm 1~~, is fast, not memory-demanding, and easy to implement. ~~The algorithm~~ It starts by initializing the adversarial perturbation  $\delta$  to a zero vector  $\mathbf{0}$  ( $\delta = \mathbf{0}$ ) (line 1). Subsequently, ~~we compute it computes~~ the gradient of the ~~attacker’s~~ objective function in Equation 5 with respect to  $\delta$  (line 3), and ~~normalize it to guarantee convergence in our optimization process (Pintor et al., 2021) normalizes it to speed up convergence (Rony et al., 2018; Pintor et al., 2021)~~. We then ~~use gradient descent~~ update  $\delta$  to minimize the ~~attacker’s objective~~, expressed in (-). ~~The objective via gradient descent, while also accounting for the~~ box constraints in Equation 6 ~~are also taken into account~~ through the usage of the clip operator. ~~Importantly, we impose~~ (line 5). ~~We enforce~~ sparsity in  $\delta$  by clipping to 0 all the components lower than ~~a fixed the current sparsity~~ threshold  $\tau$  ~~to zero~~ (Line 6). This step is necessary since the  $\hat{\ell}_0$  approximation is not exact, and might result in some values being closer to zero but not precisely zero. We therefore encourage the attack to focus only on ~~more the most~~ influential features, discarding less significant contributions. ~~The~~ We then decrease the step size  $\eta$  by following a cosine-annealing schedule (Rony et al., 2018; Pintor et al., 2021), and adjust the sparsity threshold  $\tau$  dynamically. In particular, if the current sample is adversarial, we increase  $\tau$  to promote sparser perturbations; otherwise, we decrease  $\tau$  to reduce  $\mathcal{L}$ . The variations of  $\tau$  are also iteratively reduced following the same cosine-annealing schedule of the step size. The above

<sup>1</sup>While a sigmoid approximation may be adopted to overcome the non-differentiability of the  $\mathbb{I}$  term at the decision boundary, we simply set its gradient to zero *everywhere*, without any impact on the experimental results.

process is repeated for  $N$  iterations, and if during each iteration, we find a better solution that is adversarial and has a lower  $\ell_0$  norm, we update the optimal perturbation  $\delta^*$  to the current minimum (line 10). Finally, ~~we return~~ the best adversarial perturbation  $\delta^*$  identified during the optimization process ~~is returned~~ (line 11). In conclusion, the main contributions behind  $\sigma$ -zero are: (i) the idea of exploiting the numerically-stable approximation of the  $\ell_0$  norm by Osborne et al. (2000b) to design a novel loss function (Equation 5), which enables simultaneously searching for an adversarial example while minimizing the  $\ell_0$  norm of the perturbation (i.e., a non-trivial task given the non-convexity of this norm); and (ii) the introduction of the sparsity threshold  $\tau$  and its dynamic adjustment policy which, along with gradient normalization and step size annealing, help find very sparse adversarial perturbations faster. The combination of our novel formulation with the aforementioned optimization tricks yields a very fast and reliable  $\ell_0$ -norm attack algorithm, which does not even require specific hyperparameter tuning, as we will show in our experimental results.

### 3 EXPERIMENTS

We report the extensive evaluation of the proposed  $\sigma$ -zero attack to compare its performance and efficiency with other state-of-the-art  $\ell_0$  attacks, considering sixteen baseline and robust models and three different datasets.

#### 3.1 EXPERIMENTAL SETUP

**Datasets.** We conduct experiments on three popular datasets used for benchmarking adversarial robustness: MNIST (LeCun & Cortes, 2005), CIFAR10 (Krizhevsky, 2009) and ImageNet (Krizhevsky et al., 2012). We use a random subset of 1000 test samples from ImageNet to evaluate attacks performance on it, while we consider the entire test set for MNIST and CIFAR10. For the MNIST and CIFAR10 datasets we used a batch size of 32, while for ImageNet we opted for a batch size of 16.

**Attacks.** We compare  $\sigma$ -zero ~~to several attacks that are considered to be against the following~~ state-of-the-art ~~in~~, minimum-norm attacks, in their  $\ell_0$ -norm threat models. Specifically, we evaluate the -variant of: the Elastic-Net Attacks (EAD) attack (Chen et al., 2018) as a baseline for sparse attacks in; the variants: the Voting Folded Gaussian Attack (VFGA) attack (Césaire et al., 2021), the Primal-Dual Proximal Gradient Descent (PDPGD) attack (Matyasko & Chau, 2021), the Brendel & Bethge (BB) attack (Brendel et al., 2019a), including also its variant with adversarial initialization (BBadv), and the Fast Minimum Norm (FMN) attack (Pintor et al., 2021), and the sparse random search. We also consider two state-of-the-art  $\ell_1$ -norm attacks as additional baselines, i.e., the Elastic-Net (EAD) attack (Chen et al., 2018) and SparseFool (Modas et al., 2019), along with two further  $\ell_0$ -norm attacks, i.e., the  $\ell_0$ -norm Projected Gradient Descent (PGD- $\ell_0$ ) attack (Croce & Hein, 2019) and the Sparse Random Search (Sparse-RS) attack (Croce et al., 2022). ~~Compared to the remaining approaches, the latter fixes a bound on the~~<sup>2</sup> Compared to minimum-norm attacks, PGD- $\ell_0$  and Sparse-RS aim to maximize misclassification confidence within a given maximum number of features that can be perturbed modifiable features  $k$ . To not negatively influence its performance, we perform a binary search over  $k$  to optimize the success rate and minimize the number of manipulated features, reporting the best results. Thus, to ensure a fair comparison with minimum-norm attacks, as suggested by Rony et al. (2021b), we tune their perturbation budget  $k$  by performing a sample-wise binary search to find minimum-norm adversarial examples. Further details are reported in the Appendix. Finally, we configure all attacks to manipulate input values separately, without constraining the manipulations to individual pixels; e.g., on CIFAR10, the number of modifiable inputs is thus  $3 \times 32 \times 32 = 3072$ .

**Models.** We use a selection of both baseline and robust models to evaluate the attacks under different conditions. Our goal is to compare  $\sigma$ -zero on a vast set of models to ensure its broad effectiveness and to expose vulnerabilities that may not be revealed by other attacks (Croce & Hein, 2021a). For the MNIST dataset, we consider two adversarially-trained convolutional neural network (CNN) models; by Rony et al. (2021a), i.e., CNN-DDN (Rony et al., 2021a) and CNN-Trades (Rony et al., 2021a). These models have been trained to be robust to both  $\ell_2$  and  $\ell_\infty$  adversarial attacks, ~~and we.~~ We

<sup>2</sup>Sparse-RS is a gradient-free (black-box) attack, which only requires query access to the target model. We consider it as an additional baseline in our experiments, but it should not be considered a direct competitor of gradient-based attacks, as it works under much stricter assumptions (i.e., no access to input gradients).



denote them respectively with M1 and M2. For the CIFAR10 and ImageNet datasets, we employ state-of-the-art robust models from RobustBench (Croce et al., 2021). ~~Specifically, for~~ For CIFAR10, we adopt eight models. ~~The first model, denoted as,~~ denoted with C1-C10. C1 (Croce et al., 2021) ~~is a standard is a non-robust~~ WideResNet-28-10 model. ~~The second and third models, denoted as~~ C2 (Carmon et al., 2019) and C3 (Augustin et al., 2020) ~~respectively,~~ combine training data augmentation with adversarial training to ~~obtain models that are robust improve robustness~~ to  $\ell_\infty$  and  $\ell_2$  ~~adversarial attacks.~~ ~~The fourth model, attacks.~~ C4 (Engstrom et al., 2019) ~~is an adversarially trained model that is robust to  $\ell_2$ -norm attacks.~~ ~~The fifth model,~~ C5 (Gowal et al., 2021) ~~exploits generative models to artificially augment the original training set and improve adversarial robustness to generic  $\ell_p$ -norm attacks.~~ ~~The sixth model~~ C6 (Chen et al., 2020) is a robust ensemble model ~~from~~. ~~The seventh model,~~ C7 (Xu et al., 2023) is a recently proposed adversarial training defense ~~from that is~~ robust to  $\ell_2$  ~~-norm attacks.~~ ~~Finally, the eighth model, attacks.~~ C8 (Addepalli et al., 2022) ~~enforces diversity during data augmentation and~~ ~~combine combines~~ it with adversarial training. ~~We will refer to these models respectively with , , , , , and~~ ~~Finally, we also include the  $\ell_1$  robust models~~ C9 (Croce & Hein, 2021b) ~~and~~ C10 (Jiang et al., 2023) . For ImageNet, we consider a pretrained ~~vulnerable~~ ResNet-18 ~~denoted with~~ I1 (He et al., 2015) ~~and five adversarially-trained models that are robust,~~ ~~and five robust models~~ to  $\ell_\infty$ -attacks, ~~namely, denoted with~~ I2 (Engstrom et al., 2019), I3 (Wong et al., 2020), I4 (Salman et al., 2020), I5 (Hendrycks et al., 2021) ~~and~~. ~~We will denote these models with,~~ ~~,~~ ~~,~~ ~~,~~ ~~,~~ and I6 (Salman et al., 2020) .

**Hyperparameters.** We conduct our experiments using the default ~~configurations of hyperparameters used in the original implementation of the attacks from~~ AdversarialLib (Rony & Ben Ayed) and Foolbox (Rauber et al., 2017). <sup>3</sup> ~~We do not make any modifications to the hyperparameters of the attacks, except for the~~ ~~We only change the~~ number of steps ~~which we set to 1000.~~ ~~We make this modification to ensure a fair comparison among all attacks and to achieve better convergence of them, to ensure that all attacks reach convergence~~ (Pintor et al., 2022). ~~In the Appendix we also provide a comparison analysis of the attacks when considering a low-budget attacker, i.e., only 100 steps. By using a consistent number of steps for all attacks, we can better assess their relative performance and compare their performance in generating adversarial examples across different models and datasets.~~ ~~The only exception is~~ VFGA (Césaire et al., 2021) ~~which constitutes the only exception, as it terminates only once an adversarial example is obtained.~~ ~~We report additional results using 100 steps in the Appendix. As gradient-based attacks perform one forward and one backward pass in each step, we double the steps for Sparse-RS, which, being a gradient-free attack, only makes one forward pass per iteration. This ensures a fair comparison.~~ For  $\sigma$ -zero ~~we impose  $\tau = 0.3$  and  $\sigma = 0.001$ , and we held them constant across all models tested, providing evidence for the adaptability of,~~ ~~we set 1000 steps,  $\eta_0 = 1$ ,  $\tau_0 = 0.5$  and  $\sigma = 0.1$ . We keep the same configuration for all models and datasets, showing that no specific hyperparameter tuning is required for  $\sigma$ -zero across different problems.~~ ~~Additional analysis.~~ ~~Additional analyses~~ of the influence of the hyperparameters ~~on the performance of  $\sigma$ -zero~~ can be found in the Appendix. ~~Finally, the maximum number of steps is restricted to 1000, as for the other attacks.~~

**Evaluation Metrics.** For each attack, we report the Attack Success Rate (ASR), defined as the ratio of successfully attacked samples, and the median ~~of the  $\ell_0$  norm of the successful adversarial perturbations.~~ Additionally, we report  $ASR_{\epsilon, k}$ , which indicates the ASR of attacks with a fixed budget of  $\epsilon, k$  perturbed features. We also compare the computational effort of each attack considering their execution time, the average number of queries (i.e., the sum of #forwards and #backwards) needed to perform each attack, and the ~~VRAM consumption.~~ ~~Video Random Access Memory (VRAM) consumption.~~<sup>3</sup> We measure the execution time on a workstation with NVIDIA A100 Tensor Core GPU (40GB memory) and two Intel® Xeon® Gold 6238R processors. For measuring the memory consumption, we consider the maximum amount of VRAM used by each attack among all the batches, which is a minimum requirement to run it without failure. By assessing the performance of each attack across these various metrics, we can gain a more comprehensive understanding of their effectiveness and scalability.

Table 1: Attack results with 1000 steps. For each attack, we report the corresponding ASR<sub>10</sub>, ASR<sub>50</sub>, ASR, median  $\|\delta\|_0$ , sample-level average execution time and the number of queries  $q$  ( $\times 1000$ ), and the maximum VRAM consumed during the execution. The symbol \* in VFGA suggests a potential overestimation due to using a smaller batch size for memory constraints.

| Attack         | Model | Performance          |                      |              |                | Computational effort |      |       | Model        | Performance          |                      |            |                | Computational effort |      |      |
|----------------|-------|----------------------|----------------------|--------------|----------------|----------------------|------|-------|--------------|----------------------|----------------------|------------|----------------|----------------------|------|------|
|                |       | ASR(%) <sub>10</sub> | ASR(%) <sub>50</sub> | ASR(%)       | $\ \delta\ _0$ | t(s)                 | q    | VRAM  |              | ASR(%) <sub>10</sub> | ASR(%) <sub>50</sub> | ASR(%)     | $\ \delta\ _0$ | t(s)                 | q    | VRAM |
| MNIST          |       |                      |                      |              |                |                      |      |       |              |                      |                      |            |                |                      |      |      |
| EAD            | M1    | 1.14                 | 53.66                | 100.0        | 49             | 0.47                 | 6.28 | 0.05  | M2           | 1.2                  | 55.57                | 100.0      | 48             | 0.5                  | 6.73 | 0.05 |
| VFGA           |       | 9.62                 | 82.42                | 99.98        | 27             | 0.05                 | 0.77 | 0.21  |              | 1.8                  | 39.33                | 99.95      | 57             | 0.05                 | 1.33 | 0.21 |
| PDPGD          |       | 2.97                 | 74.08                | 100.0        | 38             | 0.23                 | 2.0  | 0.04  |              | 3.31                 | 66.3                 | 100.0      | 42             | 0.23                 | 2.0  | 0.04 |
| BB             |       | 14.97                | 97.86                | 100.0        | 18             | 0.90                 | 2.99 | 0.05  |              | 25.95                | 91.62                | 100.0      | 18             | 0.74                 | 3.71 | 0.05 |
| BBadv          |       | 14.81                | 91.23                | 100.0        | 19             | 0.77                 | 2.01 | 0.07  |              | 14.42                | 40.88                | 100.0      | 89             | 0.71                 | 2.01 | 0.07 |
| PGD- $\ell_0$  |       | 12.22                | 99.84                | 100.0        | 19             | 1.15                 | 1.99 | 0.07  |              | 5.04                 | 90.17                | 100.0      | 24             | 1.42                 | 2.0  | 0.06 |
| Sparse-RS      |       | 12.57                | 83.74                | 100.0        | 25             | 3.45                 | 3.05 | 0.07  |              | 60.04                | 98.48                | 100.0      | 9.0            | 2.51                 | 2.44 | 0.06 |
| SPARSEFOOL     |       | 4.86                 | 6.76                 | 96.98        | 469            | 1.07                 | 0.18 | 0.06  |              | 0.93                 | 1.21                 | 91.68      | 463            | 2.87                 | 0.86 | 0.07 |
| FMN            |       | 7.29                 | 93.74                | 100.0        | 29             | 0.21                 | 2.0  | 0.04  |              | 10.86                | 91.84                | 99.41      | 24             | 0.22                 | 2.0  | 0.04 |
| $\sigma$ -zero |       | 19.6                 | <b>99.98</b>         | 100.0        | <b>16</b>      | 0.31                 | 2.0  | 0.04  | <b>61.57</b> | <b>100.0</b>         | 100.0                | <b>9.0</b> | 0.31           | 2.0                  | 0.04 |      |
| CIFAR10        |       |                      |                      |              |                |                      |      |       |              |                      |                      |            |                |                      |      |      |
| EAD            | C1    | 6.74                 | 21.33                | 100.0        | 126            | 2.32                 | 6.9  | 1.47  | C5           | 14.47                | 35.9                 | 100.0      | 74             | 10.76                | 5.55 | 9.92 |
| VFGA           |       | 48.58                | 93.41                | 99.99        | 11             | 0.17                 | 0.32 | 11.96 |              | 27.71                | 67.51                | 99.88      | 29             | 4.91*                | 1.02 | >40  |
| PDPGD          |       | 16.58                | 78.97                | 100.0        | 27             | 0.64                 | 2.0  | 1.31  |              | 17.68                | 40.89                | 100.0      | 69             | 3.96                 | 2.0  | 8.86 |
| BB             |       | <b>69.77</b>         | 99.79                | 100.0        | <b>7</b>       | 5.81                 | 2.76 | 1.47  |              | 13.46                | 17.14                | 17.94      | $\infty$       | 3.46                 | 2.08 | 9.93 |
| BBadv          |       | 69.58                | 99.86                | 100.0        | 7              | 4.57                 | 2.01 | 1.63  |              | -                    | -                    | -          | -              | -                    | -    | -    |
| PGD- $\ell_0$  |       | 31.82                | 85.21                | 100          | 18             | 6.46                 | 1.93 | 1.73  |              | -                    | -                    | -          | -              | -                    | -    | -    |
| Sparse-RS      |       | 59.64                | 98.59                | 100          | 9              | 3.76                 | 1.49 | 1.74  |              | -                    | -                    | -          | -              | -                    | -    | -    |
| SPARSEFOOL     |       | 11.19                | 11.19                | 56.56        | 3072           | 1.42                 | 0.37 | 1.57  |              | -                    | -                    | -          | -              | -                    | -    | -    |
| FMN            |       | 67.52                | 99.97                | 100.0        | 7              | 0.60                 | 2.0  | 1.3   |              | 27.97                | 68.38                | 100.0      | 29             | 3.91                 | 2.0  | 8.86 |
| $\sigma$ -zero |       | 80.84                | <b>100.0</b>         | 100.0        | <b>5</b>       | 0.83                 | 2.0  | 1.51  | <b>44.72</b> | <b>94.14</b>         | 100.0                | <b>12</b>  | 4.39           | 2.0                  | 9.92 |      |
| EAD            | C2    | 12.7                 | 30.38                | 100.0        | 90             | 1.92                 | 5.70 | 1.47  | C6           | 17.29                | 33.68                | 100.0      | 105            | 8.33                 | 5.37 | 5.39 |
| VFGA           |       | 28.98                | 75.37                | 99.92        | 24             | 0.59                 | 0.78 | 11.71 |              | 34.17                | 81.79                | 99.89      | 20             | 4.30*                | 0.62 | >40  |
| PDPGD          |       | 16.47                | 42.50                | 100.0        | 63             | 0.64                 | 2.0  | 1.32  |              | 21.37                | 48.96                | 99.82      | 51             | 2.15                 | 2.0  | 5.12 |
| BB             |       | 11.73                | 14.24                | 14.7         | $\infty$       | 0.63                 | 2.05 | 1.47  |              | 37.98                | 78.76                | 83.58      | 16             | 12.49                | 3.14 | 5.39 |
| BBadv          |       | 37.64                | 90.57                | 100          | 16             | 4.68                 | 2.01 | 1.64  |              | -                    | -                    | -          | -              | -                    | -    | -    |
| PGD- $\ell_0$  |       | 21.4                 | 56.85                | 100          | 39             | 5.79                 | 1.92 | 1.75  |              | -                    | -                    | -          | -              | -                    | -    | -    |
| Sparse-RS      |       | 31.02                | 62.81                | 90.78        | 27             | 6.6                  | 1.89 | 1.71  |              | -                    | -                    | -          | -              | -                    | -    | -    |
| SPARSEFOOL     |       | 18.31                | 18.77                | 56.39        | 3072           | 11.31                | 1.4  | 1.62  |              | -                    | -                    | -          | -              | -                    | -    | -    |
| FMN            |       | 28.43                | 74.7                 | 100.0        | 26             | 0.59                 | 2.0  | 1.31  |              | 33.3                 | 79.7                 | 100.0      | 21             | 2.05                 | 2.0  | 5.12 |
| $\sigma$ -zero |       | <b>47.15</b>         | <b>95.38</b>         | 100.0        | <b>11</b>      | 0.82                 | 2.0  | 1.53  | <b>49.29</b> | <b>97.14</b>         | 100.0                | <b>11</b>  | 2.71           | 2.0                  | 5.39 |      |
| EAD            | C3    | 9.21                 | 11.42                | 100.0        | 360            | 2.53                 | 5.62 | 1.89  | C7           | 9.38                 | 23.62                | 100.0      | 148            | 2.23                 | 5.8  | 2.15 |
| VFGA           |       | 21.82                | 66.5                 | 99.62        | 34             | 0.48                 | 0.94 | 16.53 |              | 22.79                | 56.72                | 99.81      | 39             | 3.15*                | 1.84 | >40  |
| PDPGD          |       | 13.96                | 45.15                | 100.0        | 55             | 1.12                 | 2.0  | 1.8   |              | 11.97                | 38.41                | 100.0      | 69             | 0.76                 | 2.0  | 2.0  |
| BB             |       | 21.32                | 56.78                | 58.64        | 33             | 2.31                 | 2.89 | 1.89  |              | 38.82                | 93.24                | 100.0      | 15             | 6.49                 | 2.87 | 2.16 |
| BBadv          |       | 31.64                | 96.31                | 100.0        | 17             | 3.92                 | 2.01 | 1.99  |              | -                    | -                    | -          | -              | -                    | -    | -    |
| PGD- $\ell_0$  |       | 17.33                | 58.82                | 100.0        | 39             | 10.31                | 1.93 | 2.30  |              | -                    | -                    | -          | -              | -                    | -    | -    |
| Sparse-RS      |       | 21.41                | 61.00                | 100.0        | 36             | 5.54                 | 2.26 | 2.20  |              | -                    | -                    | -          | -              | -                    | -    | -    |
| SPARSEFOOL     |       | 14.3                 | 21.22                | 98.74        | 3070           | 3.62                 | 0.46 | 1.90  |              | -                    | -                    | -          | -              | -                    | -    | -    |
| FMN            |       | 20.61                | 71.7                 | 100.0        | 33             | 1.08                 | 2.0  | 1.8   |              | 23.95                | 70.24                | 100.0      | 30             | 0.73                 | 2.0  | 2.0  |
| $\sigma$ -zero |       | <b>36.61</b>         | <b>97.55</b>         | 100.0        | <b>15</b>      | 1.41                 | 2.0  | 1.92  | <b>44.53</b> | <b>96.77</b>         | 100.0                | <b>12</b>  | 0.93           | 2.0                  | 2.15 |      |
| EAD            | C4    | 9.48                 | 11.14                | 100.0        | 398            | 2.57                 | 5.66 | 1.89  | C8           | 15.75                | 29.23                | 100.0      | 118            | 1.01                 | 5.32 | 0.41 |
| VFGA           |       | 30.5                 | 90.04                | 99.88        | 19             | 0.28                 | 0.52 | 16.53 |              | 29.55                | 74.15                | 99.54      | 25             | 0.17                 | 0.77 | 3.07 |
| PDPGD          |       | 15.5                 | 49.19                | 100.0        | 51             | 1.16                 | 2.0  | 1.8   |              | 19.43                | 41.0                 | 100.0      | 66             | 0.44                 | 2.0  | 0.36 |
| BB             |       | 16.32                | 31.03                | 31.36        | $\infty$       | 3.01                 | 2.37 | 1.89  |              | 38.64                | 91.83                | 100.0      | 15             | 10.90                | 2.93 | 0.41 |
| BBadv          |       | 37.06                | 99.11                | 100.0        | 14             | 4.51                 | 2.01 | 1.99  |              | 38.01                | 93.04                | 100.0      | 16             | 4.6                  | 2.01 | 0.54 |
| PGD- $\ell_0$  |       | 19.9                 | 70.04                | 100.0        | 33             | 8.97                 | 1.93 | 2.30  |              | 24.20                | 59.98                | 100        | 36             | 4.1                  | 1.9  | 0.56 |
| Sparse-RS      |       | 22.82                | 62.18                | 100          | 36             | 13.2                 | 2.26 | 2.20  |              | 31.51                | 67.82                | 98.46      | 27             | 9.84                 | 3.95 | 0.54 |
| SPARSEFOOL     |       | 15.52                | 40.86                | 93.82        | 3039           | 9.3                  | 1.56 | 1.90  |              | 23.18                | 26.54                | 51.80      | 3072           | 0.58                 | 0.33 | 0.51 |
| FMN            |       | 26.85                | 85.6                 | 100.0        | 23             | 1.09                 | 2.0  | 1.8   |              | 29.75                | 73.71                | 100.0      | 16             | 0.41                 | 2.0  | 0.36 |
| $\sigma$ -zero |       | <b>42.96</b>         | <b>99.15</b>         | 100          | <b>12</b>      | 1.39                 | 2.0  | 1.91  | <b>44.29</b> | <b>94.21</b>         | 100.0                | <b>13</b>  | 0.63           | 2.0                  | 0.51 |      |
| EAD            | C9    | 12.96                | 13.23                | 100.0        | 800            | 0.94                 | 4.89 | 0.65  | C10          | 23.94                | 24.78                | 100.0      | 768            | 1.04                 | 4.99 | 0.65 |
| VFGA           |       | 18.86                | 49.98                | 99.72        | 51             | 0.32                 | 1.25 | 4.44  |              | 33.61                | 69.47                | 99.83      | 28             | 0.25                 | 0.82 | 4.22 |
| PDPGD          |       | 15.95                | 35.13                | 100.0        | 75             | 0.41                 | 2.0  | 0.59  |              | 26.89                | 42.38                | 100        | 66             | 0.4                  | 2.0  | 0.60 |
| BB             |       | 14.13                | 22.91                | 27.64        | $\infty$       | 1.04                 | 2.25 | 0.65  |              | 24.72                | 27.98                | 29.50      | $\infty$       | 0.54                 | 2.09 | 0.65 |
| BBadv          |       | 19.93                | 72.43                | 100          | 34             | 5.28                 | 2.01 | 0.64  |              | 35.67                | 82.46                | 100        | 22             | 3.03                 | 2.01 | 0.65 |
| PGD- $\ell_0$  |       | 17.05                | 36.85                | 100.0        | 72             | 4.45                 | 1.92 | 0.72  |              | 28.2                 | 45.42                | 100.0      | 60             | 4.44                 | 1.85 | 0.70 |
| Sparse-RS      |       | 17.89                | 34.56                | 92.91        | 90             | 13.62                | 2.42 | 0.69  |              | 30.61                | 48.57                | 95.45      | 54             | 6.29                 | 2.03 | 0.68 |
| SPARSEFOOL     |       | 15.89                | 24.36                | 58.29        | 3072           | 1.63                 | 0.48 | 0.66  |              | 26.85                | 43.07                | 91.14      | 69             | 4.32                 | 1.49 | 0.66 |
| FMN            |       | 18.61                | 48.87                | 100          | 52             | 0.24                 | 2.0  | 0.60  |              | 32.63                | 62.96                | 100.0      | 34             | 0.35                 | 2.0  | 0.59 |
| $\sigma$ -zero |       | <b>21.49</b>         | <b>73.02</b>         | <b>100.0</b> | <b>32</b>      | 0.43                 | 2.0  | 0.71  | <b>37.27</b> | <b>82.92</b>         | 100.0                | <b>20</b>  | 0.42           | 2.0                  | 0.72 |      |
| ImageNet       |       |                      |                      |              |                |                      |      |       |              |                      |                      |            |                |                      |      |      |
| EAD            | I1    | 34.4                 | 36.3                 | 100.0        | 460            | 1.69                 | 6.06 | 1.21  | I4           | 56.2                 | 61.4                 | 100.0      | 0              | 1.41                 | 5.29 | 1.21 |
| VFGA           |       | 48.4                 | 72.4                 | 99.2         | 14             | 3.03*                | 1.05 | >40   |              | 61.6                 | 76.6                 | 99.3       | 1              | 3.44*                | 1.21 | >40  |
| FMN            |       | 48.7                 | 81.0                 | 100.0        | 12             | 0.62                 | 2.0  | 1.14  |              | 63.8                 | 78.7                 | 100.0      | 0              | 0.57                 | 2.0  | 1.14 |
| $\sigma$ -zero |       | <b>62.0</b>          | <b>95.9</b>          | 100.0        | <b>5</b>       | 0.81                 | 2.0  | 1.19  |              | <b>75.5</b>          | <b>92.8</b>          | 100.0      | 0              | 0.72                 | 2.0  | 1.27 |
| EAD            | I2    | 44.6                 | 51.0                 | 100.0        | 42             | 3.64                 | 5.67 | 4.36  | I5           | 26.5                 | 28.4                 | 100.0      | 981            | 3.53                 | 5.49 | 4.36 |
| VFGA           |       | 49.1                 | 63.4                 | 96.7         | 12             | 15.17*               | 2.35 | >40   |              | 36.6                 | 59.5                 | 97.9       | 31             | 12.12*               | 1.98 | >40  |
| FMN            |       | 50.9                 | 67.0                 | 100.0        | 9              | 1.21                 | 2.0  | 4.25  |              | 38.1                 | 67.7                 | 100.0      | 25             | 1.23                 | 2.0  | 4.25 |
| $\sigma$ -zero |       | <b>63.1</b>          | <b>87.4</b>          | 100.0        | <b>3</b>       | 1.43                 | 2.0  | 4.43  |              | <b>46.6</b>          | <b>86.9</b>          | 100.0      | <b>13</b>      | 1.51                 | 2.0  | 4.43 |
| EAD            | I3    | 55.1                 | 60.2                 | 100.0        | 0              | 3.53                 | 5.5  | 4.36  | I6           | 32.3                 | 33.5                 | 100.0      | 572            | 8.34                 | 5.34 | 5.67 |
| VFGA           |       | 62.2                 | 76.2                 | 98.8         | 1              | 10.12*               | 1.43 | >40   |              | 35.4                 | 46.5                 | 95.5       | 66             | 52.32*               | 3.95 | >40  |
| FMN            |       | 64.1                 | 79.5                 | 100.0        | 0              | 1.22                 | 2.0  | 4.25  |              | 35.6                 | 47.2                 | 100.0      | 58             | 3.15                 | 2.0  | 5.54 |
| $\sigma$ -zero |       | <b>75.5</b>          | <b>91.4</b>          | 100          | 0              | 1.44                 | 2.0  | 4.43  |              | <b>40.7</b>          | <b>65.1</b>          | 100.0      | <b>23</b>      | 3.75                 | 2.0  | 5.91 |




Figure 2: Robustness evaluation curves, reporting ASR versus perturbation size, for M2 on MNIST (leftmost plot), C2 on CIFAR10 (middle plot), and I1 on ImageNet (rightmost plot).

### 3.2 EXPERIMENTAL RESULTS

**Attack Performance.** Table 1 reports, for all models and datasets, the median value of  $\|\delta\|_0$  and the attack success rates. The values obtained confirm that our attack can find smaller perturbations in all cases. Specifically, over all the dataset-model configurations,  $\sigma$ -zero drastically improves the state of the art of sparse attacks. For example, on CIFAR10 models,  $\sigma$ -zero outperforms FMN by reducing the median number of manipulated features from 52 to 32 in the best case (C9) and from 7 to 5 in the worst case (C1). On ImageNet models, the median  $\|\delta\|_0$  is reduced from 58 to 23 in the best case (I6) and from 9 to 3 in the worse case (I2). Furthermore, we observe that the ASR of BB, which is the closest attack in terms of performance to  $\sigma$ -zero, drops when used in settings where the input dimensionality increases (e.g., CIFAR10), and it becomes unfeasible in extreme cases (i.e., ImageNet). From Table 1, we can also notice that the median  $\|\delta\|_0$  of BB sometimes is  $\infty$ , since its ASR is lower than 50%. BBadv does not suffer from the same issue but  $\sigma$ -zero continues to outperform that variant too. Lastly, we show in the Appendix that our attack always reaches ASR=100% against all models, even when decreasing the number of iterations. For other attacks, this is not ensured, particularly when reducing the number of iterations.

**Computational Effort.** We report the runtime comparison, the number of queries issued to the model, and the VRAM used by each attack. The results show that our attack is up to 2 (16) times faster than BB when considering MNIST (CIFAR10) models. Therefore, even if BB finds slightly better  $\ell_0$ -adversarial examples in one configuration, its computational effort is much higher than  $\sigma$ -zero. Furthermore, we observed that BB often stops unexpectedly before reaching the specified number of steps because it fails to initialize the attack.

The speed advantage of  $\sigma$ -zero is given because our attack is a simple gradient-based approach that avoids costly inner projections, such as the ones used by BB. On the other hand,  $\sigma$ -zero is slightly slower than FMN and VFGA; however, it compensates by finding better solutions. Notably, similarly to them,  $\sigma$ -zero requires fewer queries than remaining attacks. Furthermore, the speed-competing method VFGA is memory-hungry, forcing us to reduce the batch size when testing its effectiveness on larger models, e.g., C5, C6, and C7. Conversely, running our algorithm also requires reasonable VRAM, as  $\sigma$ -zero implements a lightweight search that includes only the cost of computing gradients and norms for each step. Overall, the practical advantages of our attack make it a promising direction for benchmarking large DNNs in an effective and time-efficient way.

**ImageNet Results.** For ImageNet, we restrict our analysis to EAD, FMN, and VFGA, as they outperform competing attacks on MNIST and CIFAR10 in terms of ASR, perturbation size, and execution time. While all ImageNet models are deemed robust to  $\ell_1$  and  $\ell_\infty$ -norm attacks, they are vulnerable to our  $\ell_0$ -attack. Remarkably, I6 offers higher robustness against  $\ell_0$  attacks, requiring more effort to evade it. The results show that in most configurations, our attack finds adversarial perturbations with a lower median  $\ell_0$ -norm, while being at the same time faster and memory-comparable. The results in the Appendix further confirm that even when decreasing the number of iterations to 100, our attack finds lower  $\ell_0$ -norm solutions and always achieves ASR=100%.

**Robustness Evaluation Curves.** Complementary to the performance results shown in Table 1, we present the robustness evaluation curves in Fig. 2 for each attack on M2, C2, and I1. These curves

<sup>3</sup>The default configurations correspond the hyperparameters used in the original implementation of the attacks.

<sup>3</sup>VRAM is a type of memory designed explicitly for use in Graphics Processing Units (GPUs).






Figure 3: Randomly chosen adversarial examples from MNIST M2 (top-row), CIFAR10 C2 (bottom row) found by adversarial attacks we tested.

go beyond the only median statistic and  $ASR_k$ , providing more evidence that  $\sigma$ -zero achieves higher ASR with smaller  $\ell_0$ -norm perturbations compared to the other attacks. Moreover, the ASR of our attack goes up to 100%, validating the correctness of our gradient-based approach even when considering unbounded perturbations (Carlini & Wagner, 2017). These results reinforce our previous findings that  $\sigma$ -zero is an efficient and effective method for generating adversarial examples with smaller  $\ell_0$  norm. In the Appendix, we include similar curves for all the other experimental configurations, for which results are consistent. In summary, our  $\sigma$ -zero attack consistently outperforms other state-of-the-art methods, suggesting that it can identify smaller and more effective perturbations, making it a highly promising robustness evaluation method.

**Visual Inspection of Adversarial Examples.** In Fig. 3, we show adversarial examples generated with competing  $\ell_0$ -attacks, and our  $\sigma$ -zero. First, we can see that  $\ell_0$  adversarial perturbations are always clearly visually distinguishable. Their goal, indeed, is not to be indistinguishable to the human eye – a common misconception related to adversarial examples (Biggio & Roli, 2018; Gilmer et al., 2018) – but rather to show whether and to what extent models can be fooled by just changing a few input values. For example, note how FMN and VFGA find similar perturbations, as they mostly target overlapping regions of interest. Conversely, EAD finds sparse perturbations scattered throughout the image but with a lower magnitude. This divergence is attributed to EAD’s reliance on an  $\ell_1$  regularizer, which promotes sparsity, thus diminishing perturbation magnitude without necessarily reducing the number of perturbed features. Conversely, our attack does not focus on specific areas or patterns within the images but identifies diverse critical features, whose manipulation is sufficient to mislead the target models. Given the diversity of solutions that the attacks offer, we argue that their combined usage may still improve adversarial robustness evaluation to sparse attacks.

## 4 RELATED WORK

Due to the inherent complexity of optimizing over non-convex and non-differentiable constraint, classical gradient-based algorithms like ~~Projected Gradient Descent~~ PGD (Madry et al., 2018) cannot be used for computing  $\ell_0$ -norm attacks. We categorize the existing  $\ell_0$ -norm attacks into two main groups: (i) multiple-norm attacks extended to  $\ell_0$ , and (ii) attacks specifically designed to optimize  $\ell_0$  perturbations. Furthermore, we discuss related work that leverages the approximation of  $\ell_0$  for different goals.

**Multiple-norm attacks extended to  $\ell_0$ .** These attacks are developed to work with multiple  $\ell_p$  norms and include the extension of their algorithms to the  $\ell_0$  norm. While they are able to find sparse perturbations, they often require strong use of heuristics to work in this setting. Brendel et al. (2019a) initializes the attack from an adversarial example far away from the clean sample and optimizes the perturbation by walking with small steps on the decision boundary trying to get closest to the original sample. In general, the algorithm can be used for any  $\ell_p$  norm, including  $\ell_0$ , but the individual optimization steps are very costly. Pintor et al. (2021) propose the Fast Minimum-Norm (FMN) attack that does not require an initialization step and converges efficiently with lightweight gradient-descent steps. However, their approach was developed to generalize over  $\ell_p$  norms, but it does not make special adaptations to specifically minimize the  $\ell_0$  norm. Matyasko & Chau (2021) use a two-player approach that optimizes the trade-off between perturbation size and loss of the attack and uses relaxations of the  $\ell_0$  norm (e.g.,  $\ell_{1/2}$ ) to promote sparsity. This scheme however does not strictly minimize the  $\ell_0$  norm, as the relaxation does not set the lowest components exactly to zero.

**$\ell_0$ -specific attacks.** Croce et al. (2022) introduced SparseRS, a random search-based adversarial attack that explores potential perturbation candidates to return the highest confidence solution. Unlike ~~minimum-norm-based minimum-norm~~ attacks, their approach is rooted in a maximum-confidence attack framework with a predefined number of feature manipulations. Césaire et al. (2021) have designed an attack specifically for the  $\ell_0$  norm. This attack is modeled as a stochastic Markov problem. It induces folded Gaussian noise to selected input components, iteratively finding the set that achieves misclassification with minimal perturbation. However, their approach requires a considerable amount of memory to explore the possible combinations and to find an optimal solution. This makes it infeasible to use for larger problems. With  $\sigma$ -zero, we show that the benefits from both groups, efficiency and precision, can be combined to effectively generate sparse  $\ell_0$  attacks. It stands therefore as a promising solution for evaluating DNNs’ robustness within the  $\ell_0$  threat model, which remains relatively underexplored in existing benchmarks (Croce et al., 2021).

**Approximation of the  $\ell_0$  norm.** Given the nonconvex and discontinuous nature of the  $\ell_0$  norm, ~~but also its practicability for sparse optimization,~~ the adoption of surrogate approximation functions ~~have~~ has been extensively studied (Bach et al., 2012; Weston et al., 2003; Zhang, 2008). Chen et al. (2018) use elastic-net regularization to calculate sparse perturbations, however, their attack do not necessarily find minimum  $\ell_0$ -norm perturbations. In our work, we use the formulation proposed by Osborne et al. (2000a), which provides an unbiased estimate of the actual  $\ell_0$ . Furthermore, it has been employed by Cinà et al. (2022) in the context of poisoning attacks ~~(Cinà et al., 2022)~~ to decrease sparsity in the model’s activations, while we use it as a penalty term for crafting minimum  $\ell_0$ -norm adversarial examples.

## 5 CONCLUSIONS, LIMITATIONS, AND FUTURE WORK

Despite numerous proposed attacks for assessing DNN robustness, evaluation methods tend to overlook the significance of  $\ell_0$ -norm attacks (Chen et al., 2018; Croce & Hein, 2021a). However, these attacks can provide valuable insights into identifying the minimum manipulated input values required for successful attacks and reveal crucial information about model limitations. We argue that this literature gap is primarily due to the non-differentiable nature of the  $\ell_0$  norm and its computational complexity, which poses challenges for gradient-based optimization.

In this work, we present  $\sigma$ -zero, a novel approach that leverages a smooth approximation of the  $\ell_0$  norm. By making the objective differentiable, our method becomes amenable to optimization with gradient descent. Through extensive experimentation, we demonstrate the efficacy, precision, and scalability of  $\sigma$ -zero in diverse scenarios, specifically for identifying minimal  $\ell_0$  perturbations. Our approach consistently discovers smaller minimum-norm perturbations across ~~different~~ all models and datasets, while maintaining computational efficiency in execution time and VRAM consumption, and without requiring any computationally-demanding hyperparameter tuning. By identifying the smallest number of input values that can be modified to mislead the target model, our attack provides valuable insights on the vulnerabilities of DNN models and what they learn as salient input characteristics. Additionally, it may also provide meaningful insights on how to mitigate such vulnerabilities to improve robustness.

Although our approach offers promising results for benchmarking DNNs robustness, it relies on the white-box assumption. However, in the absence of such access, attackers may resort to techniques like transferability or gradient estimation to exploit vulnerabilities (Carlini et al., 2019; Tramèr et al., 2020). We acknowledge the significance of this analysis and plan to investigate it further in future research endeavors.

In conclusion,  $\sigma$ -zero emerges as a highly promising candidate for establishing a standardized benchmark to evaluate robustness against sparse  $\ell_0$  perturbations. By facilitating more reliable and scalable assessments, it is poised to drive significant advancements in the development of novel models with improved robustness guarantees against the specific threat model under consideration.

**Ethics Statement.** Based on our comprehensive analysis, we assert that there are no identifiable ethical considerations or foreseeable negative societal consequences that warrant specific attention within the confines of this study. Rather this study will help improve the understanding of adversarial robustness properties of DNNs, and identify potential ways in which robustness can be improved.

## REFERENCES

- Sravanti Addepalli, Samyak Jain, and Venkatesh Babu R. Efficient and effective augmentation strategy for adversarial training. In *NeurIPS*, 2022.
- Maximilian Augustin, Alexander Meinke, and Matthias Hein. Adversarial robustness on in- and out-distribution improves explainability. In *Computer Vision - ECCV 2020 - 16th European Conference*, volume 12371 of *Lecture Notes in Computer Science*, pp. 228–245. Springer, 2020.
- Francis R. Bach, Rodolphe Jenatton, Julien Mairal, and Guillaume Obozinski. Optimization with sparsity-inducing penalties. *Found. Trends Mach. Learn.*, 4(1):1–106, 2012.
- Battista Biggio and Fabio Roli. Wild patterns: Ten years after the rise of adversarial machine learning. *Pattern Recognition*, 84:317–331, 2018.
- Battista Biggio, Iginio Corona, Davide Maiorca, Blaine Nelson, Nedim Srndic, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. Evasion attacks against machine learning at test time. In *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD*, volume 8190 of *Lecture Notes in Computer Science*, pp. 387–402. Springer, 2013.
- Wieland Brendel, Jonas Rauber, Matthias Kümmeler, Ivan Ustyuzhaninov, and Matthias Bethge. Accurate, reliable and fast robustness evaluation. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems, NeurIPS*, 2019a.
- Wieland Brendel, Jonas Rauber, Matthias Kümmeler, Ivan Ustyuzhaninov, and Matthias Bethge. Accurate, reliable and fast robustness evaluation. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019b.
- Nicholas Carlini and David A. Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy SP*, pp. 39–57. IEEE Computer Society, 2017.
- Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian J. Goodfellow, Aleksander Madry, and Alexey Kurakin. On evaluating adversarial robustness. *CoRR*, abs/1902.06705, 2019.
- Yair Carmon, Aditi Raghunathan, Ludwig Schmidt, John C Duchi, and Percy S Liang. Unlabeled data improves adversarial robustness. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- Manon Césaire, Lucas Schott, Hatem Hajri, Sylvain Lamprier, and Patrick Gallinari. Stochastic sparse adversarial attacks. In *33rd IEEE International Conference on Tools with Artificial Intelligence, ICTAI*, pp. 1247–1254. IEEE, 2021.
- Pin-Yu Chen, Yash Sharma, Huan Zhang, Jinfeng Yi, and Cho-Jui Hsieh. EAD: elastic-net attacks to deep neural networks via adversarial examples. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18)*, pp. 10–17. AAAI Press, 2018.
- Tianlong Chen, Sijia Liu, Shiyu Chang, Yu Cheng, Lisa Amini, and Zhangyang Wang. Adversarial robustness: From self-supervised pre-training to fine-tuning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR*, pp. 696–705. Computer Vision Foundation / IEEE, 2020.
- Antonio Emanuele Cinà, Ambra Demontis, Battista Biggio, Fabio Roli, and Marcello Pelillo. Energy-latency attacks via sponge poisoning. *CoRR*, abs/2203.08147, 2022.
- Antonio Emanuele Cinà, Kathrin Grosse, Ambra Demontis, Sebastiano Vascon, Werner Zellinger, Bernhard Alois Moser, Alina Oprea, Battista Biggio, Marcello Pelillo, and Fabio Roli. Wild patterns reloaded: A survey of machine learning security against training data poisoning. *ACM Computing Surveys*, 2022.
- Francesco Croce and Matthias Hein. Sparse and imperceivable adversarial attacks. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4723–4731, 2019.

- Francesco Croce and Matthias Hein. Mind the box:  $l_1$ -apgd for sparse adversarial attacks on image classifiers. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML*, volume 139 of *Proceedings of Machine Learning Research*, pp. 2201–2211. PMLR, 2021a.
- Francesco Croce and Matthias Hein. Mind the box:  $l_1$ -apgd for sparse adversarial attacks on image classifiers. In *International Conference on Machine Learning (ICML)*, 2021b.
- Francesco Croce, Maksym Andriushchenko, Vikash Sehwal, Edoardo Debenedetti, Nicolas Flammarion, Mung Chiang, Prateek Mittal, and Matthias Hein. Robustbench: a standardized adversarial robustness benchmark. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks*, 2021.
- Francesco Croce, Maksym Andriushchenko, Naman D. Singh, Nicolas Flammarion, and Matthias Hein. Sparse-rs: A versatile framework for query-efficient sparse black-box adversarial attacks. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI*, pp. 6437–6445. AAAI Press, 2022.
- Edoardo Debenedetti, Vikash Sehwal, and Prateek Mittal. A light recipe to train robust vision transformers. In *First IEEE Conference on Secure and Trustworthy Machine Learning*, 2023. URL <https://openreview.net/forum?id=IztT98ky0cKs>.
- Logan Engstrom, Andrew Ilyas, Hadi Salman, Shibani Santurkar, and Dimitris Tsipras. Robustness (python library), 2019. URL <https://github.com/MadryLab/robustness>.
- Justin Gilmer, Ryan P. Adams, Ian J. Goodfellow, David Andersen, and George E. Dahl. Motivating the rules of the game for adversarial example research. *CoRR*, abs/1807.06732, 2018.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations (ICLR)*, 2015.
- Sven Gowal, Sylvestre-Alvise Rebuffi, Olivia Wiles, Florian Stimberg, Dan Andrei Calian, and Timothy A. Mann. Improving robustness using generated data. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS*, pp. 4218–4233, 2021.
- Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2015.
- Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV*, pp. 8320–8329. IEEE, 2021.
- Yulun Jiang, Chen Liu, Zhichao Huang, Mathieu Salzmann, and Sabine Süsstrunk. Towards stable and efficient adversarial training against  $l_1$  bounded adversarial attacks. In *International Conference on Machine Learning*, 2023.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60:84 – 90, 2012.
- Yann LeCun and Corinna Cortes. The mnist database of handwritten digits. 2005.
- A Madry, A Makelov, L Schmidt, D Tsipras, and A Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations (ICLR)*, 2018.
- Alexander Matyasko and Lap-Pui Chau. PDPGD: primal-dual proximal gradient descent adversarial attack. *CoRR*, abs/2106.01538, 2021. URL <https://arxiv.org/abs/2106.01538>.
- Apostolos Modas, Seyed-Mohsen Moosavi-Dezfooli, and Pascal Frossard. Sparsefool: a few pixels make a big difference. In *Conference on computer vision and pattern recognition (CVPR)*, 2019.
- Michael R. Osborne, Brett Presnell, and Berwin A. Turlach. On the lasso and its dual. *Journal of Computational and Graphical Statistics*, 9:319 – 337, 2000a.

- Michael R. Osborne, Brett Presnell, and Berwin A. Turlach. On the lasso and its dual. *Journal of Computational and Graphical Statistics*, 9:319–337, 2000b.
- Maura Pintor, Fabio Roli, Wieland Brendel, and Battista Biggio. Fast minimum-norm adversarial attacks through adaptive norm constraints. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems, NeurIPS*, pp. 20052–20062, 2021.
- Maura Pintor, Luca Demetrio, Angelo Sotgiu, Ambra Demontis, Nicholas Carlini, Battista Biggio, and Fabio Roli. Indicators of attack failure: Debugging and improving optimization of adversarial examples. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 23063–23076. Curran Associates, Inc., 2022. URL [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/91ffdc5e2f12436d99914418e38d0a09-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/91ffdc5e2f12436d99914418e38d0a09-Paper-Conference.pdf).
- Jonas Rauber, Wieland Brendel, and Matthias Bethge. Foolbox: A python toolbox to benchmark the robustness of machine learning models, 2017. URL <https://github.com/bethgelab/foolbox>.
- Jérôme Rony, Luiz G. Hafemann, Luiz Oliveira, Ismail Ben Ayed, Robert Sabourin, and Eric Granger. Decoupling direction and norm for efficient gradient-based l2 adversarial attacks and defenses. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4317–4325, 2018.
- Jérôme Rony, Eric Granger, Marco Pedersoli, and Ismail Ben Ayed. Augmented lagrangian adversarial attacks. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV*, pp. 7718–7727. IEEE, 2021a.
- Jérôme Rony, Eric Granger, Marco Pedersoli, and Ismail Ben Ayed. Augmented lagrangian adversarial attacks. In *Conference on computer vision and pattern recognition (CVPR)*, 2021b.
- Jérôme Rony and Ismail Ben Ayed. Adversarial Library. URL <https://github.com/jeromerony/adversarial-library>.
- Hadi Salman, Andrew Ilyas, Logan Engstrom, Ashish Kapoor, and Aleksander Madry. Do adversarially robust imagenet models transfer better? In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS*, 2020.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations (ICLR)*, 2014a.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Yann LeCun. Intriguing properties of neural networks. In *2nd International Conference on Learning Representations, ICLR*, 2014b.
- Florian Tramèr, Nicholas Carlini, Wieland Brendel, and Aleksander Madry. On adaptive attacks to adversarial example defenses. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS*, 2020.
- Jason Weston, André Elisseeff, Bernhard Schölkopf, and Michael E. Tipping. Use of the zero-norm with linear models and kernel methods. *J. Mach. Learn. Res.*, 3:1439–1461, 2003.
- Eric Wong, Leslie Rice, and J. Zico Kolter. Fast is better than free: Revisiting adversarial training. In *8th International Conference on Learning Representations, ICLR*. OpenReview.net, 2020.
- Yuancheng Xu, Yanchao Sun, Micah Goldblum, Tom Goldstein, and Furong Huang. Exploring and exploiting decision boundary dynamics for adversarial robustness. In *International Conference on Learning Representations (ICLR)*, 2023.
- Tong Zhang. Multi-stage convex relaxation for learning with sparse regularization. In *NIPS*, 2008.



## A APPENDIX

### A.1 ROBUSTBENCH-ROBUST MODELS

The experimental setup described in this paper (Sect. 3.1) utilizes pre-trained baseline and robust models obtained from RobustBench Croce et al. (2021). The goal of RobustBench is to track the progress in adversarial robustness for  $\ell_\infty$  and  $\ell_2$ -norm attacks since these are the most studied settings in the literature. We summarize in Table 2 the models we employed in our paper. Each entry in the table includes the label reference from RobustBench, the short name we assigned to the model, and the corresponding clean and robust accuracy under the specific threat model. The robustness of these models is evaluated against an ensemble of white-box and black-box attacks, specifically AutoAttack. [Complementary, we also include models trained to be robust against  \$\ell\_1\$  sparse attacks, i.e.,](#) (Croce & Hein, 2021b)[and](#) (Jiang et al., 2023). Our experimental setup is designed to encompass a wide range of model architectures and defensive techniques, ensuring a comprehensive and thorough performance evaluation of the considered attacks.

Table 2: Summary of Robustbench Croce et al. (2021) models used in our experiments. For each model, we report its reference label in Robustbench Croce et al. (2021), its threat model, and corresponding clean and robust accuracy.

| Dataset  | Reference                     | med                           | Threat model             | Clean accuracy % | Robust accuracy % |
|----------|-------------------------------|-------------------------------|--------------------------|------------------|-------------------|
| CIFAR10  | Standard                      | C1 (Croce et al., 2021)       | -                        | 94.78            | 0                 |
|          | Carmon2019Unlabeled           | C2 (Carmon et al., 2019)      | $\ell_\infty$            | 89.69            | 59.53             |
|          | Augustin2020Adversarial       | C3 (Augustin et al., 2020)    | $\ell_2$                 | 91.08            | 72.91             |
|          | Engstrom2019Robustness        | C4 (Engstrom et al., 2019)    | $\ell_\infty$ - $\ell_2$ | 87.03 - 90.83    | 49.25 - 69.24     |
|          | Gowal2020Uncovering           | C5 (Gowal et al., 2021)       | $\ell_2$                 | 90.90            | 74.50             |
|          | Chen2020Adversarial           | C6 (Chen et al., 2020)        | $\ell_\infty$            | 86.04            | 51.56             |
|          | Xu2023Exploring_WRN-28-10     | C7 (Xu et al., 2023)          | $\ell_\infty$            | 93.69            | 63.89             |
|          | Addepalli2022Efficient_RN18   | C8 (Addepalli et al., 2022)   | $\ell_\infty$            | 85.71            | 52.48             |
| Imagenet | Standard_R18                  | I1 (He et al., 2015)          | -                        | 76.52            | 0                 |
|          | Engstrom2019Robustness        | I2 (Engstrom et al., 2019)    | $\ell_\infty$            | 62.56            | 29.22             |
|          | Wong2020Fast                  | I3 (Wong et al., 2020)        | $\ell_\infty$            | 55.62            | 26.24             |
|          | Salman2020Do_R18              | I4 (Salman et al., 2020)      | $\ell_\infty$            | 64.02            | 34.96             |
|          | Hendrycks2020Many             | I5 (Hendrycks et al., 2021)   | $\ell_\infty$            | 76.86            | 52.90             |
|          | Debenedetti2022Light_XCiT-S12 | I6 (Debenedetti et al., 2023) | $\ell_\infty$            | 72.34            | 41.78             |

### A.2 $\sigma$ -ZERO OBJECTIVE FUNCTION VISUALIZATION

In Figure 4, we depict the behavior of the loss terms of  $\sigma$ -zero when applied to the Imagenet data sample, specifically, the frog in Figure 1. When the sample is not adversarial, the attack algorithm increases the  $\ell_0$  norm, highlighted by the bumps in the orange curve, to find a valid adversarial  $\delta$ . Conversely, when an adversarial example is found, the loss term is cropped to zero, and the algorithm focus solely on minimizing the  $\ell_0$  in  $\delta$ .




Figure 4:  $\sigma$ -zero loss terms during the optimization procedure.




Figure 5: Ablation study on  $\sigma$  (y-axis) and  $\tau$  (x-axis) for ~~MNIST-CIFAR10 C1~~, CIFAR10 C8, ~~(middle-row)~~ and Imagenet (bottom-row). For each combination, we report the attack success rate (leftmost matrix) and the  $\ell_0$  norm on the output adversarial perturbation (rightmost matrix).

### A.3 ABLATION STUDY

To assess the strength and potential limitations of our proposed attack, we conducted an ablation study on its key hyperparameters. Specifically, we investigated the impact of varying two critical parameters,  $\tau$  and  $\sigma$ . The parameter  $\tau$  governs the tolerance threshold in Algorithm 1, which induces sparsity within the adversarial noise. Conversely,  $\sigma$  defines the approximation quality of  $\hat{\ell}_0$  in Equation 4 compared to the actual  $\ell_0$  function. Our ablation study, depicted in Figure 5, involved three distinct models: M1 for MNIST (top row), C8 for CIFAR10 (middle row), and I1 for Imagenet (bottom row). We executed the attack on 1000 randomly selected samples from each dataset and recorded the Attack Success Rate (ASR) and the median  $\ell_0$  norm of the resulting adversarial perturbations. Remarkably, we observe a significant robustness of  $\sigma$ -zero with respect to these two hyperparameters. Except for the extreme case of  $\tau = 0$ , where we exclude our sparsity-enforcement step in Algorithm 1 (Line 6), the adversarial noise perturbed more features. However, even with a small value of  $\tau$ , the  $\ell_0$  norm decreased while maintaining an ASR close to or equal to 100%. This suggests that a substantial portion of the adversarial noise  $\delta$  consists of minimal and negligible contributions, yet, remains effective for a successful attack. Concerning  $\sigma$ , as also noted in Cinà et al. (2022), its selection is not particularly challenging, especially when incorporating the sparsity projection operator. Overall, the ablation study revealed consistent trends across various models. In all cases, we identified a broad parameter configuration range where our attack maintained robustness, making hyperparameter optimization for the attacker a swift task. This robustness is further evidenced by the results presented in Table 1, where our attack consistently outperforms state-of-the-art attacks even with a shared hyperparameter configuration across all models.

### A.4 PRECISION AND QUERY COST ANALYSIS

In Fig. ??, we present a comprehensive comparison of different attacks based on the average number of queries utilized for crafting the attack and the median norm of the resulting adversarial noise. Our analysis demonstrates that while VFGA is fast, its precision in producing minimum-norm adversarial examples may not be optimal. Compared to BB Brendel et al. (2019a), our findings indicate that BB can be more accurate but at the cost of increased query complexity and computational demands. Furthermore, when compared against PDPGD Matyasko & Chau (2021) and FMN Pintor et al. (2021), our attack consistently offers the same number of queries on average.

These findings collectively support the argument that provides a way for assessing the robustness of DNNs against  $\sigma$ -zero attackers in an accurate and computationally feasible manner.

Precision vs. Query Cost for  $\sigma$ -zero attacks against CIFAR10 models (leftmost plot) and , , and in Imagenet (rightmost plot).

#### A.4 ATTACK COMPARISON WITH 100 STEPS

In our experimental setup, we also explore scenarios where the attacker’s access to queries is limited, thus reducing the number of iterations for the attack. To simulate this perspective, we replicate our experimental comparison involving  $\sigma$ -zero and state-of-the-art sparse attacks while restricting the number of steps to 100. The results are summarized in Table 3. Notably, compared to the results presented in Table 1, most competitive attacks undergo a decrease in their ASR, while  $\sigma$ -zero consistently maintains a 100% success rate. In conclusion,  $\sigma$ -zero remains a promising choice for crafting minimum  $\ell_0$ -norm attacks against DNNs, even when the attacker has limited query resources.

#### A.5 ROBUSTNESS EVALUATION CURVES COMPARISON WITH MAXIMUM-CONFIDENCE ATTACKS

We present in Figs. ??-?? the robustness evaluation curves depicting the performance of attacks against all the models analyzed in our paper. These findings reinforce our experimental analysis, explicitly demonstrating that the attack consistently achieves higher values of ASR while employing smaller norm perturbations compared to alternative attacks. While it is essential to acknowledge that in a few instances, may exhibit slight inferiority or comparability to BB Brendel et al. (2019b) , we emphasize that these cases are limited. Furthermore, as depicted in Fig. ??, the success rate of BB drops for different models in CIFAR10 and its computational demands render it unfeasible for application on larger datasets such as ImageNet.

#### A.6 COMPARISON WITH SPARSE-RS

In our experimental comparisons presented in both Table 1 and Table 3, we include the Sparse-RS attack introduced by Croce et al. (2022). This attack is and PGD- $\ell_0$  attacks introduced by (Croce et al., 2022) and (Croce & Hein, 2019). These attacks have been designed to generate sparse adversarial perturbations given a fixed budget  $k$ . Specifically, in their threat model, the attacker imposes a maximum limit on the number of perturbed features. Sparse-RS, and the attack then outputs the adversarial example that minimizes the model’s confidence in predicting the true label of the sample. Notably, Croce et al. (2022) has demonstrated promising results when comparing Sparse-RS against different state-of-the-art attacks like FMN or VFGA. However, since the fixed-budget threat model differs from the minimum-norm scenario we consider in this paper, as we do not assume a maximum budget, we have developed a wrapper around Sparse-RS and PGD- $\ell_0$  to ensure a fair comparison. Our process employs a binary search strategy to determine the smallest budget, denoted as  $k$ , that must be provided as input to Sparse-RS them to achieve an Attack Success Rate (ASR) equal to or greater than  $\rho$ . During 100% on the data. At each iteration of the binary search, if the ASR exceeds or equals  $\rho$  attack is successful, we halve the value of  $k$ , asking Sparse-RS and PGD- $\ell_0$  to perturb fewer features. Conversely, if the ASR falls below  $\rho$  sample is not adversarial, we double its value. The comparison tables showcase the best result with the minimum norm perturbation in the median across the iterations. Thus, our binary search approach ensures that the attacker finds the value of  $k$ . Thus, we enable the attacker to find the value of  $k$  that results in the minimum  $\ell_0$  norm for each sample. Finally, the comparison tables showcase the best result with the minimum norm perturbation in the median. In our comparison, we consider two different values for  $\rho$ , namely  $\rho = 85\%$  and  $\rho = 100\%$ , denoted as Sparse-RS<sup>85</sup> and Sparse-RS<sup>100</sup>, respectively. across the iterations.

Robustness evaluation curves for and in MNIST (two leftmost plots) , and in CIFAR10 (two rightmost plot). The dashed line represents the model prediction error on pristine samples. Ours referees to. Robustness evaluation curves for and (two leftmost plots), and (two rightmost plots) in CIFAR10. Ours referees to. Robustness evaluation curves for and (two leftmost plots) in CIFAR10,

Table 3: For each attack, we report the corresponding ASR, median  $\|\delta\|_0$ , sample-level average execution time and the number of queries, and the maximum VRAM consumed during the execution. \* in VFGA indicates that the usage of a smaller batch size, to fit its execution in memory, may have led to a slight overestimation of the execution time.

| Attack         |    | Model | Performance |                      |                      |                | Computational Effort |           |       | Model | Performance |                      |                      |                | Computational Effort |           |      |
|----------------|----|-------|-------------|----------------------|----------------------|----------------|----------------------|-----------|-------|-------|-------------|----------------------|----------------------|----------------|----------------------|-----------|------|
|                |    |       | ASR(%)      | ASR(%) <sub>10</sub> | ASR(%) <sub>50</sub> | $\ \delta\ _0$ | t(s)                 | q (x1000) | VRAM  |       | ASR(%)      | ASR(%) <sub>10</sub> | ASR(%) <sub>50</sub> | $\ \delta\ _0$ | t(s)                 | q (x1000) | VRAM |
| MNIST          |    |       |             |                      |                      |                |                      |           |       |       |             |                      |                      |                |                      |           |      |
| EAD            | M1 |       | 100.0       | 1.11                 | 46.65                | 52.0           | 0.09                 | 1.14      | 0.05  | M2    | 100.0       | 1.2                  | 35.57                | 61.0           | 0.07                 | 0.99      | 0.05 |
| VFGA           |    |       | 95.71       | 9.57                 | 82.56                | 27.0           | 0.04                 | 0.74      | 0.21  |       | 92.32       | 1.81                 | 39.28                | 57.0           | 0.04                 | 1.3       | 0.21 |
| PDPGD          |    |       | 100.0       | 0.98                 | 0.98                 | 359.0          | 0.02                 | 0.2       | 0.04  |       | 95.02       | 0.52                 | 0.52                 | 254.0          | 0.02                 | 0.2       | 0.04 |
| BB             |    |       | 100.0       | 12.8                 | 98.0                 | 20.0           | 0.13                 | 1.19      | 0.05  |       | 87.87       | 26.53                | 83.0                 | 18.0           | 0.12                 | 1.69      | 0.05 |
| FMN            |    |       | 88.93       | 7.22                 | 83.09                | 30.0           | 0.01                 | 0.2       | 0.04  |       | 14.81       | 4.02                 | 14.03                | $\infty$       | 0.01                 | 0.2       | 0.04 |
| $\sigma$ -zero |    |       | 100.0       | <b>12.46</b>         | <b>98.55</b>         | <b>21.0</b>    | 0.03                 | 0.2       | 0.05  |       | 100.0       | <b>38.3</b>          | <b>99.88</b>         | <b>13.0</b>    | 0.03                 | 0.2       | 0.05 |
| CIFAR10        |    |       |             |                      |                      |                |                      |           |       |       |             |                      |                      |                |                      |           |      |
| EAD            | C1 |       | 100.0       | 6.82                 | 19.09                | 146.0          | 0.26                 | 0.77      | 1.47  | C5    | 100.0       | 14.35                | 32.94                | 83.0           | 1.59                 | 0.82      | 9.92 |
| VFGA           |    |       | 98.99       | 49.14                | 93.46                | <b>11.0</b>    | 0.16                 | 0.38      | 11.96 |       | 87.75       | 27.64                | 67.1                 | 29.0           | 3.6*                 | 0.86      | >40  |
| PDPGD          |    |       | 100.0       | 5.23                 | 5.23                 | 3057.0         | 0.06                 | 0.2       | 1.31  |       | 99.75       | 11.26                | 11.26                | 2814.0         | 0.32                 | 0.2       | 8.86 |
| BB             |    |       | 100.0       | <b>53.48</b>         | 97.55                | <b>10.0</b>    | 0.59                 | 0.95      | 1.47  |       | 17.92       | 13.18                | 16.88                | $\infty$       | 2.67                 | 1.95      | 9.93 |
| FMN            |    |       | 98.86       | <b>62.85</b>         | 97.72                | <b>8.0</b>     | 0.05                 | 0.2       | 1.31  |       | 72.34       | 27.16                | 61.49                | 33.0           | 0.31                 | 0.2       | 8.86 |
| $\sigma$ -zero |    |       | 100.0       | 42.54                | <b>99.11</b>         | 12.0           | 0.08                 | 0.2       | 1.47  |       | 100.0       | <b>35.09</b>         | <b>85.95</b>         | <b>18.0</b>    | 0.44                 | 0.2       | 9.92 |
| EAD            | C2 |       | 100.0       | 12.74                | 28.74                | 100.0          | 0.27                 | 0.8       | 1.47  | C6    | 100.0       | 16.95                | 29.59                | 128.0          | 0.7                  | 0.66      | 5.4  |
| VFGA           |    |       | 93.69       | 28.99                | 75.38                | 24.0           | 0.22                 | 0.72      | 11.71 |       | 97.08       | 34.27                | 82.07                | 20.0           | 4.24*                | 0.61      | >40  |
| PDPGD          |    |       | 99.39       | 10.31                | 10.31                | 2421.0         | 0.06                 | 0.2       | 1.32  |       | 54.16       | 13.96                | 13.96                | 3072.0         | 0.2                  | 0.2       | 5.12 |
| BB             |    |       | 14.97       | 11.58                | 14.29                | $\infty$       | 0.44                 | 1.95      | 1.47  |       | 84.46       | 35.89                | 78.14                | 17.0           | 1.68                 | 1.67      | 5.39 |
| FMN            |    |       | 80.68       | 28.06                | 69.36                | 27.0           | 0.05                 | 0.2       | 1.31  |       | 87.1        | 32.98                | 76.59                | 22.0           | 0.19                 | 0.2       | 5.12 |
| $\sigma$ -zero |    |       | 100.0       | <b>32.58</b>         | <b>86.2</b>          | <b>18.0</b>    | 0.07                 | 0.2       | 1.47  |       | 100.0       | <b>37.99</b>         | <b>90.93</b>         | <b>16.0</b>    | 0.24                 | 0.2       | 5.39 |
| EAD            | C3 |       | 100.0       | 9.14                 | 10.67                | 451.0          | 0.31                 | 0.71      | 1.89  | C7    | 100.0       | 9.23                 | 21.61                | 162.0          | 0.31                 | 0.8       | 2.15 |
| VFGA           |    |       | 91.64       | 21.7                 | 66.55                | 33.0           | 0.34                 | 0.87      | 16.53 |       | 75.79       | 22.76                | 56.58                | 39.0           | 1.34*                | 1.06      | >40  |
| PDPGD          |    |       | 75.31       | 8.92                 | 8.92                 | 3052.0         | 0.09                 | 0.2       | 1.8   |       | 96.2        | 6.31                 | 6.31                 | 2773.0         | 0.07                 | 0.2       | 2.0  |
| BB             |    |       | 57.05       | 19.46                | 53.53                | 40.0           | 0.59                 | 1.9       | 1.89  |       | 100.0       | <b>35.43</b>         | <b>90.74</b>         | <b>16.0</b>    | 0.64                 | 1.07      | 2.16 |
| FMN            |    |       | 71.3        | 20.37                | 62.41                | 36.0           | 0.09                 | 0.2       | 1.8   |       | 68.88       | 23.21                | 59.69                | 35.0           | 0.06                 | 0.2       | 2.0  |
| $\sigma$ -zero |    |       | 100.0       | <b>23.24</b>         | <b>85.46</b>         | <b>24.0</b>    | 0.11                 | 0.2       | 1.89  |       | 100.0       | <b>32.42</b>         | <b>89.12</b>         | <b>18.0</b>    | 0.08                 | 0.2       | 2.15 |
| EAD            | C4 |       | 100.0       | 9.38                 | 10.56                | 434.0          | 0.4                  | 0.9       | 1.89  | C8    | 100.0       | 15.76                | 26.17                | 144.5          | 0.14                 | 0.79      | 0.41 |
| VFGA           |    |       | 99.16       | 30.44                | 90.13                | 19.0           | 0.27                 | 0.52      | 16.53 |       | 94.16       | 29.57                | 74.22                | 26.0           | 0.13                 | 0.73      | 3.07 |
| PDPGD          |    |       | 99.9        | 9.17                 | 9.17                 | 2709.0         | 0.09                 | 0.2       | 1.8   |       | 90.95       | 14.29                | 14.29                | 3057.0         | 0.04                 | 0.2       | 0.36 |
| BB             |    |       | 32.83       | 16.28                | 32.41                | $\infty$       | 0.49                 | 1.93      | 1.89  |       | 100.0       | <b>35.98</b>         | <b>89.39</b>         | <b>17.0</b>    | 0.43                 | 1.13      | 0.41 |
| FMN            |    |       | 87.2        | 26.63                | 79.48                | 24.0           | 0.09                 | 0.2       | 1.8   |       | 80.68       | 29.47                | 69.77                | 27.0           | 0.03                 | 0.2       | 0.36 |
| $\sigma$ -zero |    |       | 100.0       | <b>29.45</b>         | <b>92.77</b>         | <b>19.0</b>    | 0.11                 | 0.2       | 1.89  |       | 100.0       | <b>32.49</b>         | <b>85.19</b>         | <b>19.0</b>    | 0.04                 | 0.2       | 0.41 |
| Imagenet       |    |       |             |                      |                      |                |                      |           |       |       |             |                      |                      |                |                      |           |      |
| EAD            | I1 |       | 100.0       | 34.4                 | 35.7                 | 501.5          | 0.2                  | 0.71      | 1.21  | I4    | 100.0       | 56.1                 | 60.7                 | 0.0            | 0.2                  | 0.77      | 1.21 |
| VFGA           |    |       | 85.3        | 47.0                 | <b>72.4</b>          | <b>14.0</b>    | 1.74*                | 0.7       | >40   |       | 83.9        | 61.9                 | 75.4                 | 1.0            | 1.49*                | 0.59      | >40  |
| FMN            |    |       | 68.1        | <b>47.6</b>          | 64.9                 | 14.0           | 0.06                 | 0.2       | 1.14  |       | 77.2        | 63.2                 | 75.1                 | 0.0            | 0.06                 | 0.2       | 1.14 |
| $\sigma$ -zero |    |       | 100.0       | 42.9                 | 70.4                 | 20.0           | 0.08                 | 0.2       | 1.21  |       | 100.0       | <b>69.4</b>          | <b>88.0</b>          | 0.0            | 0.07                 | 0.2       | 1.2  |
| EAD            | I2 |       | 100.0       | 44.5                 | 50.2                 | 48.0           | 0.48                 | 0.75      | 4.36  | I5    | 100.0       | 26.6                 | 27.4                 | 1105.0         | 0.4                  | 0.62      | 4.35 |
| VFGA           |    |       | 72.9        | 49.2                 | 63.3                 | 12.0           | 4.34*                | 0.86      | >40   |       | 74.0        | <b>36.5</b>          | <b>58.9</b>          | <b>30.0</b>    | 4.5*                 | 0.96      | >40  |
| FMN            |    |       | 62.4        | 50.4                 | 60.0                 | 10.0           | 0.13                 | 0.2       | 4.25  |       | 53.0        | 36.0                 | 50.3                 | 46.0           | 0.13                 | 0.2       | 4.25 |
| $\sigma$ -zero |    |       | 100.0       | <b>54.8</b>          | <b>79.9</b>          | <b>6.0</b>     | 0.15                 | 0.2       | 4.35  |       | 100.0       | <b>32.3</b>          | <b>57.1</b>          | <b>40.0</b>    | 0.15                 | 0.2       | 4.35 |
| EAD            | I3 |       | 100.0       | 54.8                 | 60.5                 | 0.0            | 0.47                 | 0.74      | 4.35  | I6    | 99.9        | 32.2                 | 32.9                 | 818.5          | 1.14                 | 0.72      | 5.67 |
| VFGA           |    |       | 83.0        | 62.2                 | 76.0                 | 1.0            | 3.47*                | 0.59      | >40   |       | 57.4        | 35.4                 | <b>46.9</b>          | 67.0           | 13.35*               | 1.2       | >40  |
| FMN            |    |       | 74.2        | 63.3                 | 71.8                 | 0.0            | 0.13                 | 0.2       | 4.25  |       | 47.5        | <b>35.7</b>          | 45.2                 | $\infty$       | 0.34                 | 0.2       | 5.54 |
| $\sigma$ -zero |    |       | 100.0       | <b>69.9</b>          | <b>86.8</b>          | 0.0            | 0.15                 | 0.2       | 4.35  |       | 100.0       | <b>36.9</b>          | <b>53.3</b>          | <b>40.0</b>    | 0.41                 | 0.2       | 5.68 |




Figure 6: Randomly chosen adversarial examples from MNIST M2.

and (two rightmost plots) in Imagenet. Ours denotes  $\sigma$ -zero. Robustness evaluation curves for and (two leftmost plots), and (two rightmost plots) in Imagenet. Ours denotes  $\sigma$ -zero.




Figure 7: Randomly chosen adversarial examples from CIFAR10 C2 .

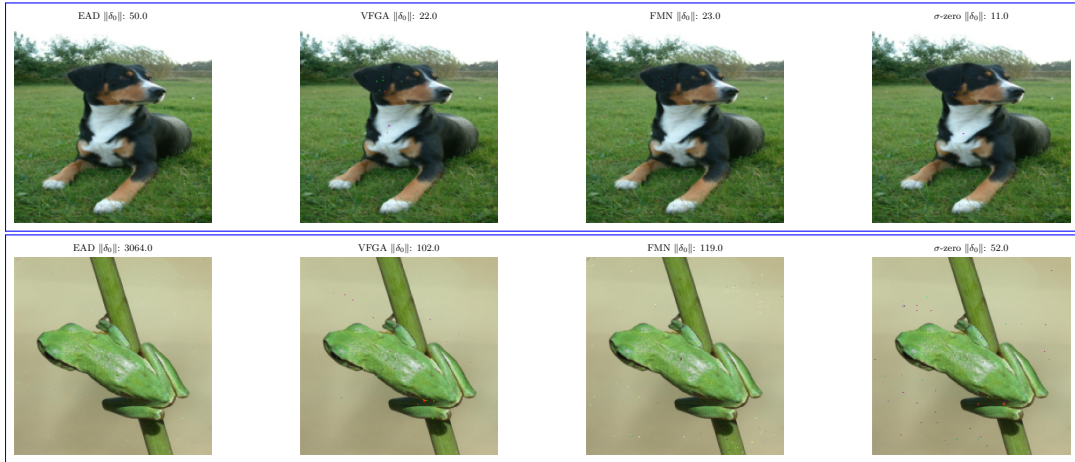


Figure 8: Randomly chosen adversarial examples from Imagenet I1 .