A SIMPLE UNSUPERVISED DATA DEPTH-BASED METHOD TO DETECT ADVERSARIAL IMAGES

Anonymous authors
Paper under double-blind review

ABSTRACT

Deep neural networks suffer from critical vulnerabilities regarding robustness, which limits their exploitation in many real-world applications. In particular, a serious concern is their inability to defend against adversarial attacks. Although the research community has developed a large amount of effective attacks, the detection problem has received little attention. Existing detection methods either rely on additional training or on specific heuristics at the risk of overfitting. Moreover, they have mainly focused on ResNet architectures while transformers, which are state-of-the-art for vision tasks, have not been properly investigated. In this paper, we overcome these limitations by introducing `APPROVED`, a simple unsupervised detection method for transformer architectures. It leverages the information available in the logit layer and computes a similarity score with respect to the training distribution. This is accomplished using a *data depth* that is: *(i)* computationally efficient; and *(ii)* non-differentiable, making it harder for gradient-based adversaries to craft malicious samples. Our extensive experiments show that `APPROVED` consistently outperforms previous detectors on CIFAR10, CIFAR100 and Tiny ImageNet.

1 INTRODUCTION

Recent years have seen a rapid development of Deep Neural Networks (DNNs), which have led to a significant improvement over previous state-of-the-art methods (SOTA) in numerous decision-making tasks. However, together with this growth, concerns have been raised about the potential failures of deep learning systems, which limit their large-scale adoption (Alves et al., 2018; Johnson, 2018; Subbaswamy and Saria, 2020). In Computer Vision, a particular source of concern is the existence of *adversarial attacks* (Szegedy et al., 2014), which are samples created by adding to the original (clean) image a well-designed additive perturbation, imperceptible to human eyes, with the goal of fooling a given classifier. The vulnerability of DNNs to such kinds of attacks limits their deployment in safety-critical systems as in aviation safety management (Ali et al., 2020), health monitoring systems (Leibig et al., 2017; Meinke and Hein, 2020)) or in autonomous driving (Bojarski et al., 2016; Guo et al., 2017). Therefore, it is crucial to deploy a proper strategy to defend against adversarial attacks (Amodei et al., 2016).

In this context, the task of distinguishing adversarial samples from clean ones is becoming increasingly challenging as developing new attacks is getting more attention from the community (Gao et al., 2021; Wang et al., 2021a; Naseer et al., 2021; Duan et al., 2020; Zhao et al., 2020; Lin et al., 2019; Deng and Karam, 2020a;b; Wu et al., 2020b; Croce and Hein, 2020; Jia et al., 2020; Dong et al., 2019). Inspired by the concept of *rejection channels* (Chow, 1957), which was proposed over 60 years ago for the character recognition problem, one way to address adversarial attacks is to construct a detector-based rejection strategy. Its objective is to discriminate malicious samples from clean ones, which implies discarding samples detected as adversarial. Research in this area focuses on both *supervised* and *unsupervised* approaches (Aldahdooh et al., 2021c). The supervised approaches rely on features extracted from adversarial examples generated according to one or more attacks (Kherchouche et al., 2020; Feinman et al., 2017; Ma et al., 2018); the unsupervised ones, instead, do not rely on prior knowledge of attacks and, in general, only learn from clean data at the time of training (Xu et al., 2018; Meng and Chen, 2017).

In this work, we focus on the unsupervised scenario, which is often a reasonable approach to real-world use-cases. We model the adversarial detection problem as an *anomaly detection* framework (Breunig et al., 2000; Schölkopf et al., 2001; Liu et al., 2008; Staerman et al., 2019; 2020; Chandola et al., 2009), where the aim is to identify abnormal observations without seeing them during training. Statistical tools called *data depths* are natural similarity score in this context. Data depths have a simple geometric interpretation as they provide center-outward ordering of points with respect to a probability distribution (Tukey, 1975; Zuo and Serfling, 2000). Geometrically speaking, the data depths measure how deep a sample is in a given distribution. Although data depths have received attention from the statistical community, they remain overlooked by the machine learning community.

Contributions. Our contributions can be summarized as follows:

1. **Building on novel tools: data depths.** Our first contribution is to introduce APPROVED, A simPle unsuPeRvised method fOr adVersarial imagE Detection. Given an input, APPROVED considers its embedding in the last layer of the pre-trained classifier and computes the depth of the sample w.r.t the training probability distribution. The deeper it is, the less likely it is to be adversarial. Contrarily to existing methods that involve additional networks training (Meng and Chen, 2017) or heavily rely on opaque feature engineering (Xu et al., 2018), APPROVED is computationally efficient and has a simple geometrical interpretation. Moreover, data depths non-differentiability making it harder for to gradient-based attackers to target APPROVED.

2. **A truly upgraded experimental setting.** Motivated by practical considerations which are different from previous works (Kherchouche et al., 2020; Xu et al., 2018; Meng and Chen, 2017; Ma et al., 2018; Feinman et al., 2017) focusing on ResNets (He et al., 2016), we choose to benchmark APPROVED on vision transformers models (Dosovitskiy et al., 2021; Tolstikhin et al., 2021; Steiner et al., 2021; Chen et al., 2021; Zhai et al., 2022). Indeed, such networks achieve state-of-the-art results on several visual tasks, including object detection (He et al., 2021), image classification (Wang et al., 2021b) and generation (Parmar et al., 2018), largely outperforming ResNets. Interestingly enough, we empirically observe that transformers behave differently from ResNets, which justifies the need to develop detection techniques such as APPROVED, that leverage the specific features of transformers' architectures. Moreover, to avoid overfitting on a specific attack, we test our detection method on a wide range of attack mechanisms.

3. **Ensuring reproducibility.** We provide the open-source code of our method, attacks, and baseline to ensure reproducibility and reduce future research computation and coding overhead.

Organization of the paper. The paper is organized as follows. In Sec. 2, we review detection methods along with attack mechanisms. In Sec. 3, we introduce our detector APPROVED and focus on the description of the data depth on which it relies. In Sec. 4, we study the performance of adversarial attacks on vision transformers and give insights on models' behavior under threat. In Sec. 5, we evaluate APPROVED through numerical experiments and concluding remarks are relegated to Sec. 6.

2 Background and Related Work

Notations. Let us consider the classical supervised learning problem where $x \in \mathcal{X} \subseteq \mathbb{R}^d$ denotes the input sample in the space $\mathcal{X}$, and $y \in \mathcal{Y} = \{1, \dots, C\}$ denotes its associated label. The unknown data distribution is denoted by p_{XY}. The training dataset $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$ is defined as $n \geq 1$ independent identically distributed (i.i.d.) realizations of p_{XY}. Consider $\mathcal{D}_c = \{(x_i, y_i) \in \mathcal{D} : y_i = c\}$, the training data for a given class $c \in \mathcal{Y}$. We define the empirical training distribution for the class c at layer ℓ as $p_c^\ell = \frac{1}{|\mathcal{D}_c|} \sum_{x \in \mathcal{D}_c} \delta_{f_\theta^\ell(x)}$ where δ_u is the dirac mass at point u.

Let $f_\theta^\ell : \mathcal{X} \to \mathbb{R}^{d_\ell}$ with $\ell \in \{1, \dots, L\}$, denote the output of the ℓ-th layer of the DNN, where d_ℓ is the dimension of the latent space induced by the ℓ-th layer. The class prediction is obtained from the L-th layer softmax output as follows:

$$f_\theta^L(x) \triangleq \underset{c \in \mathcal{Y}}{\arg\max}\ q_\theta\left(c|x\right) \text{ with } q_\theta\left(\cdot|x\right) = \text{softmax}(f_\theta^{L-1}(x)).$$

2.1 REVIEW OF ATTACK MECHANISMS

The existence of adversarial examples and their capability to lure a deep neural network have been first introduced in Szegedy et al. (2014). The authors define the adversarial generation problem as:

$$x' = \underset{x' \in \mathbb{R}^d \,:\, \|x'-x\|_p < \varepsilon}{\arg\min} \|x' - x\|_p \;\text{ s.t. }\; f_\theta^L(x') \neq y, \tag{1}$$

where y is the true label associated to a natural sample $x \in \mathcal{X}$ being modified, $\|\cdot\|_p$ is the L_p-norm operator, and ε is the maximal allowed perturbation.

Multiple techniques have since been crafted to solve this problem. They can be divided into two main groups of attack mechanisms depending on the knowledge they have of the DNN model: whitebox and blackbox attacks. The former has full access to the model, its weights, and gradients, while the latter can only rely on queries.

Carlini & Wagner's (CW) (Carlini and Wagner, 2017) attack is among the strongest whitebox attacks developed yet. It attempts to solve the adversarial problem in Eq. 1 by regularizing the minimization of the perturbation norm by a surrogate of the misclassification constraint. *DeepFool (DF)* (Moosavi-Dezfooli et al., 2016) is an iterative method that solves a locally linearized version of the adversarial problem and takes a step in that direction.

The authors of Goodfellow et al. (2014) relaxe the problem as follows:

$$x' = \underset{x' \in \mathbb{R}^d \,:\, \|x'-x\|_p < \varepsilon}{\arg\max} \mathcal{L}(x, x'; \theta), \tag{2}$$

where $\mathcal{L}(x, x'; \theta)$ is the objective of the attacker, which is a surrogate of the misclassification constraint, and propose the *Fast Gradient Sign Method (FGSM)* that approximates the solution of the relaxed problem in Eq. 2 by taking one step in the direction of the sign of the gradient of the attacker's objective w.r.t. the input. *Basic Iterative Method (BIM)* (Kurakin et al., 2018) and *Projected Gradient Descent (PGD)* (Madry et al., 2018) are two iterative extensions of the FGSM algorithm. Their main difference relies on the initialization of the attack algorithm, i.e., while BIM initializes the adversarial examples to the original samples PGD adds a random uniform noise on it. Although created for L_∞-norm constraints, these three methods can be extended to any L_p-norm constraint.

To overcome the absence of knowledge about the model to attack, *Hop Skip Jump (HOP)* (Chen et al., 2020) tries to estimate the model's gradient through queries. *Square Attack (SA)* (Andriushchenko et al., 2020) is based on random searches for a perturbation. If the perturbation doesn't increase the attacker's objective, it is discarded. Finally, *Spatial Transformation Attack (STA)* (Engstrom et al., 2019) rotates and translates the original samples to fool the model.

2.2 REVIEW OF DETECTION METHODS

Defending methods against adversarial attacks have been widely studied for classical CNNs (Madry et al., 2018; Zhang et al., 2019; Alayrac et al., 2019; Wang et al., 2019; Hendrycks et al., 2019; Rice et al., 2020; Atzmon et al., 2019; Huang et al., 2020; Carmon et al., 2019; Wu et al., 2020a). Whereas a few works have focused on studying the robustness of vision transformers to adversarial samples (Aldahdooh et al., 2021a; Benz et al., 2021; Mahmood et al., 2021). Meanwhile, to protect adversarial attacks from disrupting DNNs' functioning, it is possible to craft detectors to ensure that the sample can be trusted.

Building a detector falls down to finding a scoring function $s : \mathbb{R}^d \to \mathbb{R}$ and a threshold $\gamma \in \mathbb{R}$ to build a binary rule $g : \mathbb{R}^d \to \{0, 1\}$. For a given test sample $x \in \mathbb{R}^d$,

$$g(x) = \mathbb{I}\{s(x) > \gamma\} = \begin{cases} 1 \text{ if } s(x) > \gamma, \\ 0 \text{ if } s(x) \leq \gamma. \end{cases} \tag{3}$$

If s is an anomaly score, $g(x) = 0$ implies that x is considered as 'natural', i.e., sampled from p_{XY}, and $g(x) = 1$ implies that x is considered as 'adversarial', i.e., perturbed, and if s is a similarity score, the opposite decision is made.

A detection method can act on the model to be protected by modifying its training procedure using tools such as reverse cross-entropy (Pang et al., 2018) or the rejection option (Sotgiu et al., 2020;

Aldahdooh et al., 2021b). In that case, both detector and model are trained jointly. Those methods are usually vulnerable to changes in attack mechanisms, and thus, they need global re-training if a modification of the detector is introduced. On the other hand, it is also possible to craft detectors on top of a fixed trained model. Those methods can be divided into two main categories: supervised methods, where the detector knows the attack that will be perpetrated, and unsupervised methods, where the detector can only rely on clean samples, which is not desired in practice.

Generally, supervised methods use simple machine learning algorithms (e.g., SVM or a logistic regressor) to distinguish between natural and adversarial examples. The effectiveness of such methods heavily relies on natural and adversarial feature extraction. They can be extracted directly from the network's layers (Lu et al., 2017; Carrara et al., 2018; Metzen et al., 2017), or using statistical tools (e.g., maximum mean discrepancy (Grosse et al., 2017), PCA (Li and Li, 2017), kernel density estimation (Feinman et al., 2017), local intrinsic dimensionality (Ma et al., 2018), model uncertainty (Feinman et al., 2017) or natural scene statistics (Kherchouche et al., 2020)). Supervised methods, which heavily depend on the knowledge about the perpetrated attack, tend to overfit to that attack mechanism and usually generalize poorly.

Unsupervised methods do not assume any knowledge of the attacker. Indeed, new attack mechanisms are crafted every year, and it is realistic to assume knowledge about the attacker. To overcome that absence of prior knowledge about the attacker, unsupervised methods can only rely on natural samples. The features extraction rely on different techniques, such feature squeezing (Xu et al., 2018), adaptive noise, (Liang et al., 2021), using denoising autoencoders (Meng and Chen, 2017), network invariant (Ma et al., 2019) or training an auxiliary model (Sotgiu et al., 2020; Aldahdooh et al., 2021b; Zheng and Hong, 2018). Raghuram et al. (2021) uses dimension reduction, kNN and layer aggregation to distinguish between natural and adversarial samples. In this paper, we only focus on unsupervised methods that cannot act on the model to be protected.

3 Our Proposed Detector

3.1 Background on data depth

The notion of data depth goes back to John Tukey in 1975, who introduced the halfspace depth (Tukey, 1975). Data depth functions are useful nonparametric statistics allowing to rank elements of a multivariate space $\mathbb{R}^d$ w.r.t. a probability distribution (or a dataset). Given a random variable Z which follows the distribution p_Z, a data depth can be defined as:

$$\begin{array}{rccc} D: & \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) & \longrightarrow & [0,1], \\ & (z, p_Z) & \longmapsto & D(z, p_Z). \end{array} \tag{4}$$

The higher the value of the depth function, the deeper the element is in the reference distribution. Various data depths have been introduced over the year (see Chapter 2 of Staerman (2022) for a survey), including halfspace depth (Tukey, 1975), the simplicial depth (Liu, 1990), the projection depth (Liu, 1992) or the zonoid depth (Koshevoy and Mosler, 1997). Despite their applications in statistics and machine learning (e.g., regression (Rousseeuw and Hubert, 1999; Hallin et al., 2010), classification (Mozharovskyi et al., 2015), automatic text evaluation (Staerman et al., 2021b) or anomaly detection (Serfling, 2006; Rousseeuw and Hubert, 2018; Staerman et al., 2020; 2022)) the use of data depth with representation models, and more generally to deep learning, remains overlooked by the community. The halfspace depth is the first and the most studied depth in the literature probably due to its appealing properties (Donoho and Gasko, 1992; Zuo and Serfling, 2000) as well as its connections with univariate quantiles. However, it suffers from computational burden in practice (Rousseeuw and Struyf, 1998; Dyckerhoff and Mozharovskyi, 2016). Indeed, it requires solving an optimization problem over the unit hypersphere of a non-differentiable quantity. To remedy this drawback, (Ramsay et al., 2019) introduced the Integrated Rank-Weighted (IRW) depth (see also Chen et al. (2015); Staerman et al. (2021a)), which involves an expectation as an alternative to the infimum over the unit hypersphere of the halfspace depth, making it easier to compute. The IRW depth is scale and translation invariant and has been successfully applied to anomaly detection (Chen et al., 2015; Staerman et al., 2021a) making it a good candidate for our purposes. Formally, the IRW depth is defined as:

$$D_{\mathrm{IRW}}(z, p_Z) = \int_{\mathbb{S}^{d-1}} \min \left\{ F_u\left(\langle u, z\rangle\right), 1 - F_u\left(\langle u, z\rangle\right) \right\} \mathrm{d}u, \tag{5}$$

where the unit hypersphere is denoted as $\mathbb{S}^{d-1}$ and $F_u(t) = \mathbb{P}(\langle u, Z\rangle \leq t)$. A Monte-Carlo scheme is used to approximate the expectation by an empirical means. Given a training dataset $\mathcal{S}_n = \{z_1, \ldots, z_n\}$ following p_Z, denotes $u_k \in \mathbb{S}^{d-1}$, the empirical version of the IRW depth, which can be computed in $\mathcal{O}(n_{\text{proj}}nd)$ and is then linear in all of its parameters, is defined as:

$$\widetilde{D}^{\text{MC}}_{\text{IRW}}(z, \mathcal{S}_n) = \frac{1}{n_{\text{proj}}} \sum_{k=1}^{n_{\text{proj}}} \min \left\{ \frac{1}{n} \sum_{i=1}^{n} \mathbb{I}\left\{\langle u_k, z_i - z\rangle \leq 0\right\}, \frac{1}{n} \sum_{i=1}^{n} \mathbb{I}\left\{\langle u_k, z_i - z\rangle > 0\right\} \right\}, \quad (6)$$

3.2 APPROVED: OUR DEPTH-BASED DETECTOR

Intuition. Our detector tries to answer this simple question: can we find a metric that will be able to distinguish between natural and arbitrary adversarial samples? At the logit layer, we want to compare the new input to the training samples of its predicted class to measure whether the new sample is behaving as expected. Data depths, particularly the IRW depth, are serious candidates as they measure the 'distance' between a given new input to the training probability distribution.

APPROVED in a nutshell. To detect whether a given model f_θ can trust a new input x, APPROVED will perform three steps:

1. **Logits computation**. For an new input x, APPROVED first require to extract the logits (i.e., $f_\theta^{L-1}(x)$) from the pretrained classifier.
2. **Similarity score computation**. APPROVED relies on the IRW depth score $D_{\text{IRW}}(f_\theta^{L-1}(x), p_{\hat{y}}^{L-1})$, between $p_{\hat{y}}^{L-1}$, the training distribution of the predicted class $\hat{y} = f_\theta^L(x)$ at the logit layer, and $f_\theta^{L-1}(x)$, using Algo.1 in App. B to evaluate equation 6.
3. **Thresholding.** For a given threshold γ, the test input sample x is detected as clean if $D_{\text{IRW}}(f_\theta^{L-1}(x), p_{\hat{y}}^{L-1}) > \gamma$, otherwise, it is considered as adversarial. A classical way to select γ it by selecting an amount of training samples the detector can wrongfully detect.

3.3 COMPARISON WITH EXISTING DETECTORS

We benchmark our approach with two unsupervised detection methods: FS and MagNet. We chose these baselines because they are unsupervised and do not modify the model to protect. We could consider NIC (Ma and Liu, 2019) but extracting features at each layer is computationally expensive.

The Feature Squeezing method (FS; (Xu et al., 2018)). It computes the feature squeezing of the input, extracts its prediction, and compares it to the original prediction. The further away they are, the more likely the input is adversarial. In practice, four versions of the input are needed: the original input, a low-precision version, a median-filtered version, and a denoising-filtered version. One inference on the model is required for each of the four inputs. Later, the maximal L_1 difference between the original prediction and each of the other three is picked. FS is, therefore, parameter-free and does not require training. However, the necessary time to extract the essential features and the memory needed to store all the input modifications and their prediction are quite high.

MagNet (Meng and Chen, 2017). It is based on the training of two components: first, a detector that detects if a sample x is clean or not, then, a reformer that tries to find an approximation of the input closer to the training manifold. In practice, each of those components is an autoencoder that must be trained on clean samples before testing new inputs. MagNet requires three inferences at testing time, one on the detector, one on the reformer, and one on the original model to protect. Therefore, even though, at inference time, little time is necessary to output the prediction, MagNet requires careful training, which is time-consuming, and storing two autoencoders, which is highly memory-consuming.

APPROVED, similarly to FS and contrary to MagNet, does not require training time and is parameter-free. Contrary to FS, it only requires one inference on the model to extract the logits of the input. It is, therefore, *less computationally and time-consuming*. The summary of computational time and resources needed to deploy each detection method is provided in App. C. Finally, since data depths are non-differentiable, it is not straightforward for gradient-based attacks that have full access to the detection method to attack APPROVED.

4 ADVERSARIAL ATTACKS ON VISION TRANSFORMERS (VIT)

In the following, we provide insights on the behavior of vision transformers under the threat of adversarial attacks, along with a comparison to the classically used ResNets models.

4.1 SET-UP

Datasets and classifiers. We conducted our study on pretrained Vision Transformers. We rely on three widely used vision datasets: CIFAR10 (Krizhevsky, 2009), CIFAR100 and TinyImageNet (Tiny) Jiao et al. (2019). Training details can be found in App. A.

Model	Dataset	Acc (%)
ViT	CIFAR10	98.7
	CIFAR100	92.4
	Tiny ImageNet	86.4

Table 1: ViT accuracy for each dataset

Performance measures. We use two different metrics to compare the different detection methods:

AUROC↑: Area Under the Receiver Operating Characteristic curve (Davis and Goadrich, 2006). It represents the relation between True Positive Rates (TPR), i.e., the percentage of perturbed samples detected as adversarial, and False Positive Rates (FPR), i.e., the percentage of clean samples detected as adversarial. The higher the AUROC↑ is, the better the detector's performances are.

FPR$\downarrow_{90\%}$: False Positive Rate at 90% True Positive Rate. It represents the number of natural samples detected as adversarial when 90% of the attacked samples are detected. Lower is better.

Remark. We discard the perturbed samples that do not fool the underlying classifier. Indeed, detecting a sample that does not perturb the classifier's functioning as natural or adversarial is a valid answer.

Attacks. For the experiments, we will evaluate the different detection methods on the attacks presented in Sec. 2.1. Under L_1-norm constraint, we craft attacks following PGD^1 scheme. For the L_2-norm constraint, we consider PGD^2, DF and HOP. Under L_∞-norm constraint, we study PGD$^\infty$, BIM and FGSM attacks, CW$^\infty$ and SA. Finally, we create STA attacks, which are not subject to a norm constraint. The values of the maximal allowed perturbation are discussed in the next section.

4.2 ADVERSARIAL ATTACK CALIBRATION

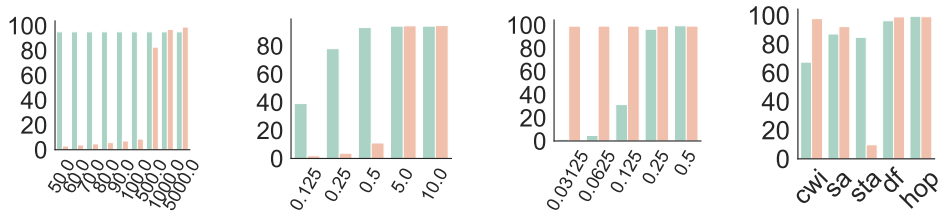

Figure 1: Percentage of successful attacks depending on the L_p-norm constraint, the maximal perturbation ε and the attack algorithm on ResNet18 (green) and ViT (orange).

Given that the variety of attacks comes from choosing the L_p-norm constraint and the maximal allowed L_p-norm perturbation ε, it is crucial to select them carefully. Adversarial attacks and defense mechanisms have been widely studied for classical convolutional networks, particularly for ResNets (Goodfellow et al., 2014; Moosavi-Dezfooli et al., 2016; Zhang et al., 2019; Madry et al., 2018; Xu et al., 2018; Meng and Chen, 2017). Hence, choosing the maximally allowed perturbation ε for ViT comes naturally from comparing the success attack rates between attacks on ResNets and ViTs.

In Fig. 1, we present the success rates for attacks on Resnet18 (resp. on ViT) in blue (resp. in orange), for different attack mechanisms, different L_p-norms and different maximal perturbation ε (the results for FGSM and BIM are relegated to App. D). Attacks behave differently on ResNets and on ViTs: on L_∞-norm constraints, at equal ε, the attacks are more potent on the ViT than on ResNet18. Indeed, the input of a ViT has more pixels than the input of a ResNet ($32 \times 32 \times 3$ for ResNet and $224 \times 224 \times 3$ for ViT). Limiting the perturbation by an L_∞-norm constraint, i.e., controlling the maximal perturbation pixel-wise without controlling the number of modified pixels, will create samples further away from the original sample if it has more pixels. On the contrary, under L_1 and L_2-norm constraints, the opposite behavior is observable: at fixed ε, the attack are more potent on ResNets than on ViTs. This can be explained by the fact that limiting L_1 or L_2-norm perturbations controls the average perturbations on the whole input sample. The modifications are therefore smaller

Table 2: Averaged results over the different attacks for each considered L_p-Norm constraints for APPROVED, FS and MagNet, along with the Averaged results over the norms. The results are presented as mean $\pm standard\_deviation$. The best results are presented in **bold**.

	APPROVED						FS						MagNet					
	CIFAR10		CIFAR100		Tiny		CIFAR10		CIFAR100		Tiny		CIFAR10		CIFAR100		Tiny	
	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$
L_1	**94.0** ±5.2	**13.2** ±13.5	**78.3** ±7.6	**46.4** ±13.1	**75.2** ±1.3	**59.2** ±2.7	79.5 ±3.3	34.9 ±8.3	71.1 ±5.1	55.5 ±8.0	54.2 ±14.0	75.1 ±11.0	51.3 ±1.1	91.1 ±0.9	50.1 ±0.2	90.2 ±0.2	49.4 ±0.9	90.0 ±1.4
L_2	**94.1** ±3.7	**14.6** ±13.5	**80.5** ±4.9	**44.0** ±11.3	**76.8** ±4.6	**53.8** ±10.1	77.3 ±1.8	37.2 ±8.6	68.2 ±5.1	58.9 ±10.5	61.8 ±12.0	72.4 ±10.6	51.0 ±1.2	89.8 ±2.7	50.6 ±0.7	89.3 ±2.0	49.9 ±1.4	89.0 ±2.5
L_∞	**95.3** ±6.5	**13.4** ±20.9	**86.7** ±0.4	**29.9** ±19.3	**91.8** ±8.8	**19.1** ±19.6	73.0 ±3.5	53.4 ±18.3	62.6 ±6.8	67.3 ±11.6	74.6 ±17.8	61.2 ±23.9	56.2 ±9.7	80.0 ±17.4	55.0 ±8.7	81.3 ±15.8	50.9 ±2.6	88.3 ±4.3
no Norm	**94.9** ±0.0	**10.5** ±0.0	**87.4** ±0.0	**32.1** ±0.0	**80.2** ±0.0	**42.5** ±0.0	78.8 ±0.0	37.5 ±0.0	65.4 ±0.0	50.0 ±0.0	53.0 ±0.0	77.5 ±0.0	39.4 ±0.0	93.5 ±0.0	38.3 ±0.0	92.8 ±0.0	34.9 ±0.0	95.6 ±0.0
Average	**94.7** ±5.6	**13.5** ±17.5	**83.2** ±8.9	**37.2** ±17.8	**83.9** ±10.3	**37.7** ±23.8	75.8 ±4.2	44.2 ±16.4	66.1 ±7.0	62.0 ±11.8	65.8 ±18.0	67.7 ±19.6	53.3 ±7.7	85.3 ±13.5	52.3 ±7.0	85.7 ±12.0	49.8 ±3.4	89.2 ±3.6

pixel-wise if the image is bigger. While on ResNets, the classical values of ε are lower than 40 on L_1-norm constraints and 2 on L_2-norm-constraints, we had to increase the maximum ε studied for those L_p-norm constraint to have successful enough attacks. Finally, Spatial Transformation Attacks (STA) disturb ResNets' functioning more easily than ViTs'.

Summary. To sum up, in the remaining of the paper, under L_1-norm constraint, we craft PGD[1] attacks with maximum norm constraint $\varepsilon \in \{50, 60, 70, 80, 90, 100, 500, 1000, 5000\}$. For the L_2-norm, we consider PGD[2] with $\varepsilon \in \{0.125, 0.25, 0.5, 5, 10\}$, DF with no ε, and HOP attacks with 3 restarts and $\varepsilon = 0.1$. Under L_∞-norm constraint, we consider PGD$^\infty$, BIM and FGSM attacks with $\varepsilon \in \{0.03125, 0.0625, 0.125, 0.25, 0.5\}$, CW$^\infty$ with $\varepsilon = 0.3125$ and SA with $\varepsilon = 0.125$. Finally, STA attacks, which are not subject to a norm constraint, can rotate the input up to 60°, and translate it up to 10 pixels.

4.3 Locating the relevant information

In the previous section, we saw that the attacks behave differently w.r.t. the classifier on which they are perpetrated. We now continue this investigation by looking at the differences between the two models from the depth scores' perspective. In this framework, we define the layer to have relevant information when the difference between the depth score on the naturals and the depth score on the adversarial is significant. Indeed, the higher the difference, the more evident the shift between the distributions of the natural and the adversarial induced by the depth score will be, and hence the easier it will be to find a threshold that distinguishes natural from adversarial samples. We start by computing layer per layer the differences between the IRW depth on the natural samples ($D_{\text{IRW}}(f_\theta^\ell(x); p_{\hat{y}}^\ell)$) and on the adversarial samples ($D_{\text{IRW}}(f_\theta^\ell(x'); p_{\hat{y}}^\ell)$) both for ViT and for ResNet18. In Fig. 2, we plot the mean and standard deviation for each layer and each network.The diamond points represent the outliers. Fig. 2 shows that the information about whether a sample is natural or adversarial, based on the study of the IRW depth, is significantly spread across the ResNet18 model: in each layer, the values range between 0 and 0.06. On the contrary, on ViT, this information is concentrated in the logit layer, where the values range between 0.05 and 0.2 while the values range from 0 to 0.05 for the other layers. To summarize, while relevant information to distinguish between

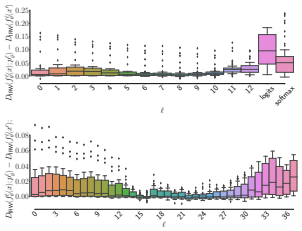

Figure 2: Difference between natural and adversarial IRW depth values as a function of the layer on ViT (top) and on ResNet18 (bottom), averaged over the attacks.

Table 3: Averaged results over the different types of attack mechanism for APPROVED, FS, and MagNet, along with the averaged results over the norms. The results are presented as mean $\pm standard\_deviation$. The best results are presented in **bold**. Dashed values (–) corresponds to attacks that take more than 48 hours to run on V100 GPUs.

	APPROVED						FS						MagNet					
	CIFAR10		CIFAR100		Tiny		CIFAR10		CIFAR100		Tiny		CIFAR10		CIFAR100		Tiny	
	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$	AUROC↑	FPR$_{↓90\%}$
PGD	**95.5** ±4.3	**9.6** ±10.8	**81.3** ±7.8	**41.2** ±14.6	**81.0** ±10.2	**45.0** ±24.1	77.2 ±3.8	44.4 ±14.2	70.1 ±4.9	62.5 ±11.1	65.6 ±19.0	66.2 ±22.5	51.5 ±1.4	89.7 ±2.7	62.4 ±13.3	68.1 ±24.1	49.5 ±0.7	90.0 ±1.1
BIM	**96.8** ±4.4	**7.1** ±10.1	**82.1** ±13.1	**37.9** ±26.2	**95.0** ±7.4	**11.8** ±16.6	71.2 ±1.5	69.6 ±2.9	64.3 ±1.9	77.8 ±3.4	86.5 ±2.2	60.4 ±19.2	52.6 ±2.2	86.0 ±3.3	51.8 ±1.6	86.9 ±2.7	49.9 ±0.2	89.7 ±0.3
FGSM	**90.5** ±8.8	**29.7** ±29.4	**90.4** ±6.4	**23.9** ±15.8	**85.6** ±7.2	**33.5** ±14.6	73.7 ±4.3	32.7 ±5.0	54.8 ±6.2	56.0 ±5.4	52.8 ±3.3	75.1 ±2.3	62.6 ±11.1	65.2 ±26.4	62.4 ±13.3	68.1 ±24.1	52.9 ±4.0	85.4 ±6.8
HOP	**98.3** ±0.0	**3.3** ±0.0	**89.1** ±0.0	**24.8** ±0.0	**87.1** ±0.0	**31.8** ±0.0	74.5 ±0.0	25.0 ±0.0	62.7 ±0.0	50.0 ±0.0	59.1 ±0.0	76.3 ±0.0	53.4 ±0.0	83.6 ±0.0	50.0 ±0.0	89.9 ±0.0	52.7 ±0.0	83.8 ±0.0
DeepFool	**86.5** ±0.0	45.4 ±0.0	**75.5** ±0.0	59.9 ±0.0	–	–	79.7 ±0.0	**31.2** ±0.0	62.2 ±0.0	**50.0** ±0.0	–	–	50.3 ±0.0	89.7 ±0.0	50.0 ±0.0	89.9 ±0.0	–	–
SA	**98.2** ±0.0	**3.3** ±0.0	**89.6** ±0.0	**26.0** ±0.0	**77.0** ±0.0	**49.1** ±0.0	72.0 ±0.0	25.0 ±0.0	63.3 ±0.0	50.0 ±0.0	48.7 ±0.0	78.5 ±0.0	55.1 ±0.0	82.4 ±0.0	54.9 ±0.0	82.6 ±0.0	50.6 ±0.0	89.4 ±0.0
CW	**90.4** ±0.0	**30.6** ±0.0	**81.7** ±0.0	**42.2** ±0.0	–	–	78.8 ±0.0	37.5 ±0.0	67.0 ±0.0	50.0 ±0.0	–	–	50.6 ±0.0	89.3 ±0.0	50.0 ±0.0	89.8 ±0.0	–	–
STA	**94.9** ±0.0	**10.5** ±0.0	**87.4** ±0.0	**32.1** ±0.0	**80.2** ±0.0	**42.5** ±0.0	78.8 ±0.0	37.5 ±0.0	65.4 ±0.0	50.0 ±0.0	53.0 ±0.0	77.5 ±0.0	39.4 ±0.0	93.5 ±0.0	38.3 ±0.0	92.8 ±0.0	34.9 ±0.0	95.6 ±0.0

natural and adversary samples is diffused in the ResNet18 model, which has small and similar values for all the layers, the most valuable information is instead concentrated at the logit layer for the ViT network, which experiences larger values only for that particular layer. It seems, therefore, relevant to build a detector *specific* for vision transformers based *only* on the output of the logit layer.

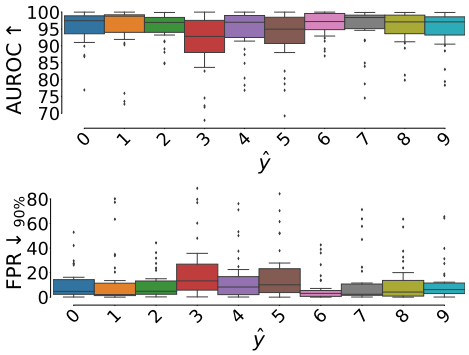

Figure 3: APPROVED's AUROC↑ and FPR↓$_{90\%}$ per class, averaged over CIFAR10.

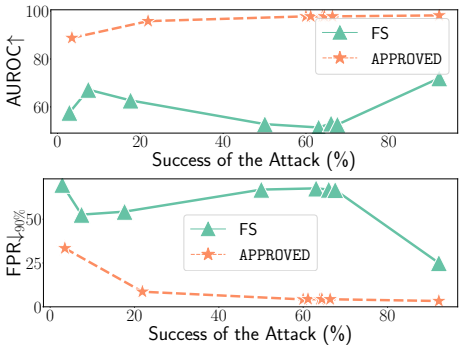

Figure 4: Detector Performances under blackbox Adaptive Attack.

5 Experiments

5.1 Results

Performances of APPROVED compared to other unsupervised detection methods. In Tab. 6, Tab. 7, and Tab. 8 relegated to App. E, we report the detailed results for each considered detection method under the threat of each attack mechanism, L_p-norm constraint and maximum perturbation ε. In Tab. 2, we report the averaged AUROC↑ and FPR↓$_{90\%}$ on each of the considered L_p-norm, along with the global average for each detector, on CIFAR10, CIFAR100, and Tiny ImageNet. Overall, APPROVED shows better results than the SOTA detection methods. On CIFAR10, APPROVED creates an increase of AUROC↑ of 18.9% and a decrease of FPR↓$_{90\%}$ of 30.7% compared to the best performing state-of-the-art detector, i.e., FS. On CIFAR100, the improvements are 17.1% and 24.8%, respectively, while they are 18.0% and 29.9% on Tiny ImageNet. In addition, FS and APPROVED have similar dispersions. Moreover, under specific L_p-norm constraints, our method consistently outperforms SOTA methods, especially under the L_∞-norm constraint where APPROVED outperforms FS (resp. MagNet) by 22.3% (resp. 39.1%) in terms of AUROC↑ and 40.0% (resp. 66.6%) in terms of FPR↓$_{90\%}$ on CIFAR10. Finally, by looking at the detailed results presented in Tab. 6 and Tab. 7, we can deduce that FS and APPROVED have opposite behaviors: when the performances of FS decrease, APPROVED's performances tend to improve. For example, under the L_∞-norm constraint, APPROVED has more trouble detecting attacks with small perturbations, while FS has more difficulty detecting attacks with large perturbations. Indeed, since APPROVED measures the depth of a sample within a distribution, it will be able to recognize the strongest attacks well.

Performances per attack. In Tab. 3 we give the overall idea of the results on all three datasets per attack mechanism by showing them in terms of *mean* and *standard deviation* (std) on the AUROC↑ and on the FPR↓$_{90\%}$. APPROVED turns out to consistently outperform the state-of-the-art detectors forall datasets. In particular, we notice that the FGSM attacks that are the easiest to generate are the ones that present the highest diversity among the results in the methods examined. Indeed, by looking at Tab. 3, we can find larger values of the standard deviation in correspondence to that attack. Moreover, APPROVED is capable of recognizing attacks that are more difficult for the competitors (e.g., STA for MagNet or FGSM for FS). We also observe that for APPROVED the most challenging task is to distinguish natural and adversarial samples when they are crafted with DeepFool. However, it is the best choice even in this case as it reaches better performances than the other detectors.

5.2 Adaptive Attacks

In this experiment, we evaluate APPROVED against adaptive attacks, which has knowledge about the defense (Athalye et al., 2018; Tramer et al., 2020; Carlini and Wagner, 2017). Two scenarii can be considered with adaptive attacks: whitebox and blackbox. Whitebox attacks (*e.g.* BPDA (Athalye et al., 2018)) are not straightforward to adapt in our case since finding a differentiable surrogate of IRW remains a very challenging open research question in the statistical community, which has never been tackled. As a matter of fact, the only attempts to approximate a non-differentiable depth was performed not on the IRW depth but on the Tuckey depth in Dyckerhoff et al. (2021), with very poor results as pointed out in She et al. (2021). Thus, in this experiment, we rely on blackbox attacks and

present the results in Fig. 4. We attacked both APPROVED and FS using a modified version of SA (Engstrom et al., 2019), for which the attack objective has been modified to allow the attacker to fool both the detection method as well as the classifier. We rely on an hyperparameter α that weights the relative importance of the two parts of the objective. It is straightforward (cf. Fig. 4) that APPROVED is less sensitive to adaptive attacks than FS. This results further validates the use of the IRW depth to craft detection method and further assesses the superiority of APPROVED.

5.3 Finer Analysis

Per class analysis. As explained in Sec. 3.2, APPROVED is based on the IRW depth, which computes the depth score of the sample w.r.t. the original distribution by class. Fig. 3 shows the per-class performances averaged over the different attacks on CIFAR10, while Fig. 7, relegated to App. F due to space constraints, shows the performances on CIFAR100. It is clear from Fig. 3 that APPROVED does not have equal performances on every class. In particular, some classes present extremely high mean average AUROC↑ (i.e., class 7), others exhibit very low FPR↓$_{90\%}$ (i.e., class 6), while some others have their adversarial and clean samples tough to distinguish (i.e., class 3 and 5). The same behavior is observable of CIFAR100 (see Fig. 7).

AUROC↑ vs FPR↓$_{90\%}$. We conclude our analysis by looking at the trade-off between AUROC↑and FPR↓$_{90\%}$(see Fig. 5). The ideal method would concentrate the results on the upper left of the figure, which corresponds to high AUROC↑ and low FPR↓$_{90\%}$, while a poor detector would concentrate them in the bottom right corner of the figure, which corresponds to low AUROC↑ and high FPR↓$_{90\%}$. We observe that on CIFAR10, the APPROVED points are more concentrated in the upper left corner, while the FS points are concentrated in the center of the figure and MagNet's in the lower right corner. On CIFAR100 and Tiny ImageNet, the results for our method are slightly more spread out in the top left and center of the figure, while for FS and MagNet, they are still in the center and bottom right, respectively. Note that FS has a different behavior than expected, i.e., the line connecting the top left corner with the bottom right corner. This behavior change can be observed for FPR↓$_{90\%}$ between 25%-35% on CIFAR10 and between 50%-75% on CIFAR100 and Tiny ImageNet. On CIFAR10, FS presents a lower AUROC↑ for a fixed FPR↓$_{90\%}$ than expected, whereas, on CIFAR100, it presents a lower AUROC↑ (for FPR↓$_{90\%}$ values between 50%-60%) or higher (for FPR↓$_{90\%}$ values around 75%) than expected.

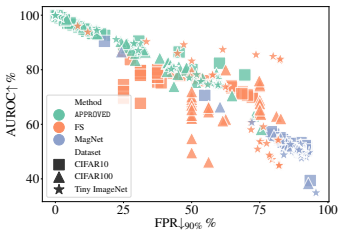

Figure 5: AUROC↑ as a function of FPR↓$_{90\%}$ for APPROVED, FS, and MagNet on all considered datasets.

6 Conclusions and Limitations

We introduced APPROVED, an efficient unsupervised detection method designed to defend against adversarial attacks. In contrast with previous detection methods, which were built for ResNet architectures, APPROVED is well suited for vision transformers which nowadays represent the state-of-the-art. While the relevant information about the discrepancy between clean and adversarial samples is distributed across all layers of ResNets, for the transformers, it was empirically shown to be concentrated in the logit layer. This motivated us to build APPROVED on top of this logit layer by computing a similarity score between an input sample and the training distribution based on the statistical notion of *data depth*. We chose to use the Integrated Rank-Weighted depth, which lends itself to fast inference computations and is non-differentiable, making it harder for gradient-based adversarial methods to craft malicious samples.We conduct extensive numerical experiments and prove that APPROVED outperforms the other state-of-the-art methods significantly.

Future Research. We think our method paves the way for future research efforts. Indeed, there is still room for improvement: even if the AUROC↑ performances are good, the FPR↓$_{90\%}$ are also fairly high. We believe the idea of leveraging information contained in layers of transformers through data depths can be fruitful in improving defense mechanisms against adversarial attacks. Our research is expected to have a positive societal impact by protecting the integrity of AI systems, especially necessary in critical systems such as autonomous cars (Morgulis et al., 2019) or stock predictions (Xie et al., 2022).

REFERENCES

Jean-Baptiste Alayrac, Jonathan Uesato, Po-Sen Huang, Alhussein Fawzi, Robert Stanforth, and Pushmeet Kohli. Are labels required for improving adversarial robustness? In *Advances in Neural Information Processing Systems*, pages 12214–12223, 2019.

Ahmed Aldahdooh, Wassim Hamidouche, and Olivier Deforges. Reveal of vision transformers robustness against adversarial attacks. *arXiv preprint arXiv:2106.03734*, 2021a.

Ahmed Aldahdooh, Wassim Hamidouche, and Olivier Déforges. Revisiting model's uncertainty and confidences for adversarial example detection. *arXiv preprint arXiv: 2103.05354*, 2021b.

Ahmed Aldahdooh, Wassim Hamidouche, Sid Ahmed Fezza, and Olivier Déforges. Adversarial example detection for DNN models: A review. *arXiv preprint arXiv:2105.00203*, 2021c.

Muhammad Ali, Yim-Fun Hu, Doanh Kim Luong, George Oguntala, Jian-Ping Li, and Kanaan Abdo. Adversarial attacks on ai based intrusion detection system for heterogeneous wireless communications networks. In *2020 AIAA/IEEE 39th Digital Avionics Systems Conference (DASC)*, pages 1–6. IEEE, 2020.

Erin Alves, Devesh Bhatt, Brendan Hall, Kevin Driscoll, Anitha Murugesan, and John Rushby. Considerations in assuring safety of increasingly autonomous systems. *NASA*, 2018.

Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.

Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion, and Matthias Hein. Square attack: a query-efficient black-box adversarial attack via random search. In *European Conference on Computer Vision*, pages 484–501. Springer, 2020.

Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *International conference on machine learning*, pages 274–283. PMLR, 2018.

Matan Atzmon, Niv Haim, Lior Yariv, Ofer Israelov, Haggai Maron, and Yaron Lipman. Controlling neural level sets. In *Advances in Neural Information Processing Systems*, pages 2034–2043, 2019.

Philipp Benz, Soomin Ham, Chaoning Zhang, Adil Karjauv, and In So Kweon. Adversarial robustness comparison of vision transformer and mlp-mixer to cnns. *arXiv preprint arXiv:2110.02797*, 2021.

Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.

M.M. Breunig, H.-P. Kriegel, R.T. Ng, and J. Sander. Lof: Identifying density-based local outliers. In *ACM SIGMOD*, volume 29, pages 93–104. ACM, 2000.

Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017.

Yair Carmon, Aditi Raghunathan, Ludwig Schmidt, John C Duchi, and Percy S Liang. Unlabeled data improves adversarial robustness. In *Advances in Neural Information Processing Systems*, pages 11192–11203, 2019.

Fabio Carrara, Rudy Becarelli, Roberto Caldelli, Fabrizio Falchi, and Giuseppe Amato. Adversarial examples detection in features distance spaces. In *Computer Vision - ECCV 2018 Workshops - Munich, Germany, Proceedings, Part II*, volume 11130, pages 313–327. Springer, 2018.

V. Chandola, A. Banerjee, and V. Kumar. Anomaly detection: A survey. *ACM Comput. Surv.*, 41(3): 15:1–15:58, 2009. ISSN 0360-0300.

Bo Chen, Kai Ming Ting, Takashi Washio, and Gholamreza Haffari. Half-space mass: a maximally robust and efficient data depth method. *Machine Learning*, 100(2):677–699, 2015.

Jianbo Chen, Michael I Jordan, and Martin J Wainwright. Hopskipjumpattack: A query-efficient decision-based attack. In *2020 ieee symposium on security and privacy (sp)*, pages 1277–1294. IEEE, 2020.

Xiangning Chen, Cho-Jui Hsieh, and Boqing Gong. When vision transformers outperform resnets without pretraining or strong data augmentations. *arXiv preprint arXiv:2106.01548*, 2021.

Chi-Keung Chow. An optimum character recognition system using decision functions. *IRE Transactions on Electronic Computers*, (4):247–254, 1957.

Francesco Croce and Matthias Hein. Minimally distorted adversarial examples with a fast adaptive boundary attack. In *International Conference on Machine Learning*, pages 2196–2205. PMLR, 2020.

Jesse Davis and Mark Goadrich. The relationship between precision-recall and roc curves. In *Proceedings of the 23rd international conference on Machine learning*, pages 233–240, 2006.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.

Yingpeng Deng and Lina J Karam. Frequency-tuned universal adversarial attacks. *arXiv preprint arXiv:2003.05549*, 2020a.

Yingpeng Deng and Lina J Karam. Towards imperceptible universal attacks on texture recognition. *arXiv preprint arXiv:2011.11957*, 2020b.

Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Evading defenses to transferable adversarial examples by translation-invariant attacks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4312–4321, 2019.

David L. Donoho and Miriam Gasko. Breakdown properties of location estimates based on half space depth and projected outlyingness. *The Annals of Statistics*, 20:1803–1827, 1992.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021.

Ranjie Duan, Xingjun Ma, Yisen Wang, James Bailey, A Kai Qin, and Yun Yang. Adversarial camouflage: Hiding physical-world attacks with natural styles. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1000–1008, 2020.

Rainer Dyckerhoff and Pavlo Mozharovskyi. Exact computation of the halfspace depth. *Computational Statistics & Data Analysis*, 98:19–30, 2016.

Rainer Dyckerhoff, Pavlo Mozharovskyi, and Stanislav Nagy. Approximate computation of projection depths. *Computational Statistics & Data Analysis*, 157:107166, 2021.

Logan Engstrom, Brandon Tran, Dimitris Tsipras, Ludwig Schmidt, and Aleksander Madry. Exploring the landscape of spatial robustness. In *International Conference on Machine Learning*, pages 1802–1811. PMLR, 2019.

Reuben Feinman, Ryan R Curtin, Saurabh Shintre, and Andrew B Gardner. Detecting adversarial samples from artifacts. *arXiv preprint arXiv:1703.00410*, 2017.

Ruijun Gao, Qing Guo, Felix Juefei-Xu, Hongkai Yu, and Wei Feng. Advhaze: Adversarial haze attack. *arXiv preprint arXiv:2104.13673*, 2021.

Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.

Kathrin Grosse, Praveen Manoharan, Nicolas Papernot, Michael Backes, and Patrick D. McDaniel. On the (statistical) detection of adversarial examples. *arXiv preprint arXiv:1702.06280*, 2017.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML*, volume 70, pages 1321–1330, 2017.

Marc Hallin, Davy Paindaveine, and Miroslav Šiman. Multivariate quantiles and multiple-output regression quantiles: From l1 optimization to halfspace depth. *The Annals of Statistics*, 38(2): 635–669, 04 2010.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

Lu He, Qianyu Zhou, Xiangtai Li, Li Niu, Guangliang Cheng, Xiao Li, Wenxuan Liu, Yunhai Tong, Lizhuang Ma, and Liqing Zhang. End-to-end video object detection with spatial-temporal transformers. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 1507–1516, 2021.

Dan Hendrycks, Kimin Lee, and Mantas Mazeika. Using pre-training can improve model robustness and uncertainty. In *International Conference on Machine Learning*, pages 2712–2721. PMLR, 2019.

Lang Huang, Chao Zhang, and Hongyang Zhang. Self-adaptive training: beyond empirical risk minimization. *Advances in Neural Information Processing Systems*, 33, 2020.

Ethan Huynh. Vision transformers in 2022: An update on tiny imagenet. *arXiv preprint arXiv:2205.10660*, 2022.

Yunhan Jia Jia, Yantao Lu, Junjie Shen, Qi Alfred Chen, Hao Chen, Zhenyu Zhong, and Tao Wei Wei. Fooling detection alone is not enough: Adversarial attack against multiple object tracking. In *International Conference on Learning Representations (ICLR'20)*, 2020.

Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*, 2019.

CW Johnson. The increasing risks of risk assessment: On the rise of artificial intelligence and non-determinism in safety-critical systems. In *the 26th Safety-Critical Systems Symposium*, page 15. Safety-Critical Systems Club York, UK., 2018.

Anouar Kherchouche, Sid Ahmed Fezza, Wassim Hamidouche, and Olivier Déforges. Natural scene statistics for detecting adversarial examples in deep neural networks. In *22nd IEEE International Workshop on Multimedia Signal Processing*, pages 1–6. IEEE, 2020.

Gleb Koshevoy and Karl Mosler. Zonoid trimming for multivariate distributions. *The Annals of Statistics*, 25(5):1998–2017, 10 1997.

Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.

Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-100 (canadian institute for advanced research). URL http://www.cs.toronto.edu/~kriz/cifar.html.

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

Christian Leibig, Vaneeda Allken, Murat Seçkin Ayhan, Philipp Berens, and Siegfried Wahl. Leveraging uncertainty information from deep neural networks for disease detection. *Scientific reports*, 7(1):1–14, 2017.

Xin Li and Fuxin Li. Adversarial examples detection in deep networks with convolutional filter statistics. In *IEEE International Conference on Computer Vision, ICCV*, pages 5775–5783. IEEE Computer Society, 2017.

Bin Liang, Hongcheng Li, Miaoqiang Su, Xirong Li, Wenchang Shi, and Xiaofeng Wang. Detecting adversarial image examples in deep neural networks with adaptive noise reduction. *IEEE Trans. Dependable Secur. Comput.*, 18(1):72–85, 2021.

Jiadong Lin, Chuanbiao Song, Kun He, Liwei Wang, and John E Hopcroft. Nesterov accelerated gradient and scale invariance for adversarial attacks. *arXiv preprint arXiv:1908.06281*, 2019.

Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *2008 eighth ieee international conference on data mining*, pages 413–422. IEEE, 2008.

Regina Y. Liu. On a notion of data depth based on random simplices. *The Annals of Statistics*, 18(1): 405–414, 1990.

Regina Y. Liu. *Data Depth and Multivariate Rank Tests*, page 279–294. North-Holland, Amsterdam, 1992.

Jiajun Lu, Theerasit Issaranon, and David A. Forsyth. Safetynet: Detecting and rejecting adversarial examples robustly. In *IEEE International Conference on Computer Vision*, pages 446–454. IEEE Computer Society, 2017.

Shiqing Ma and Yingqi Liu. Nic: Detecting adversarial samples with neural network invariant checking. In *Proceedings of the 26th network and distributed system security symposium (NDSS 2019)*, 2019.

Shiqing Ma, Yingqi Liu, Guanhong Tao, Wen-Chuan Lee, and Xiangyu Zhang. NIC: detecting adversarial samples with neural network invariant checking. In *26th Annual Network and Distributed System Security Symposium*. The Internet Society, 2019.

Xingjun Ma, Bo Li, Yisen Wang, Sarah M. Erfani, Sudanthi N. R. Wijewickrema, Grant Schoenebeck, Dawn Song, Michael E. Houle, and James Bailey. Characterizing adversarial subspaces using local intrinsic dimensionality. In *6th International Conference on Learning Representations*, 2018.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.

Kaleel Mahmood, Rigel Mahmood, and Marten Van Dijk. On the robustness of vision transformers to adversarial examples. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7838–7847, 2021.

Alexander Meinke and Matthias Hein. Neural networks that provably know when they don't know. In *8th International Conference on Learning Representations, ICLR*, 2020.

Dongyu Meng and Hao Chen. Magnet: A two-pronged defense against adversarial examples. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 135–147. ACM, 2017.

Jan Hendrik Metzen, Tim Genewein, Volker Fischer, and Bastian Bischoff. On detecting adversarial perturbations. In *5th International Conference on Learning Representations*, 2017.

Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2574–2582, 2016.

Nir Morgulis, Alexander Kreines, Shachar Mendelowitz, and Yuval Weisglass. Fooling a real car with adversarial traffic signs. *arXiv preprint arXiv:1907.00374*, 2019.

Pavlo Mozharovskyi, Karl Mosler, and Tatjana Lange. Classifying real-world data with the $DD\alpha$-procedure. *Advances in Data Analysis and Classification*, 9(3):287–314, 2015.

Muzammal Naseer, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Fatih Porikli. On generating transferable targeted perturbations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7708–7717, 2021.

Maria-Irina Nicolae, Mathieu Sinn, Minh Ngoc Tran, Beat Buesser, Ambrish Rawat, Martin Wistuba, Valentina Zantedeschi, Nathalie Baracaldo, Bryant Chen, Heiko Ludwig, Ian Molloy, and Ben Edwards. Adversarial robustness toolbox v1.2.0. *CoRR*, 1807.01069, 2018. URL https://arxiv.org/pdf/1807.01069.

Tianyu Pang, Chao Du, Yinpeng Dong, and Jun Zhu. Towards robust detection of adversarial examples. In *Advances in Neural Information Processing Systems 31*, pages 4584–4594, 2018.

Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *International Conference on Machine Learning*, pages 4055–4064. PMLR, 2018.

Jayaram Raghuram, Varun Chandrasekaran, Somesh Jha, and Suman Banerjee. A general framework for detecting anomalous inputs to dnn classifiers. In *International Conference on Machine Learning*, pages 8764–8775. PMLR, 2021.

Kelly Ramsay, Stéphane Durocher, and Alexandre Leblanc. Integrated rank-weighted depth. *Journal of Multivariate Analysis*, 173:51–69, 2019.

Leslie Rice, Eric Wong, and Zico Kolter. Overfitting in adversarially robust deep learning. In *International Conference on Machine Learning*, pages 8093–8104. PMLR, 2020.

Peter J. Rousseeuw and Mia Hubert. Regression depth. *Journal of the American Statistical Association*, 94(446):388–402, 1999.

Peter J. Rousseeuw and Mia Hubert. Anomaly detection by robust statistics. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(2):e1236, 2018.

Peter J. Rousseeuw and Anja Struyf. Computing location depth and regression depth in higher dimensions. *Statistics and Computing*, 8(3):193–203, 1998.

Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2016.

B. Schölkopf, J.C. Platt, J. Shawe-Taylor, A. Smola, and R. Williamson. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7):1443–1471, 2001.

Robert Serfling. Depth functions in nonparametric multivariate inference. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 72, 2006.

Yiyuan She, Shao Tang, and Jingze Liu. On generalization and computation of tukey's depth: Part i. *arXiv preprint arXiv:2112.08475*, 2021.

Angelo Sotgiu, Ambra Demontis, Marco Melis, Battista Biggio, Giorgio Fumera, Xiaoyi Feng, and Fabio Roli. Deep neural rejection against adversarial examples. *EURASIP J. Inf. Secur.*, 2020:5, 2020.

Guillaume Staerman. *Functional anomaly detection and robust estimation*. PhD thesis, Institut polytechnique de Paris, 2022.

Guillaume Staerman, Pavlo Mozharovskyi, Stephan Clémençon, and Florence d'Alché Buc. Functional isolation forest. In *Proceedings of The Eleventh Asian Conference on Machine Learning*, volume 101, pages 332–347, 2019.

Guillaume Staerman, Pavlo Mozharovskyi, and Stéphan Clémençon. The area of the convex hull of sampled curves: a robust functional statistical depth measure. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108, pages 570–579. PMLR, 2020.

Guillaume Staerman, Pavlo Mozharovskyi, and Stéphan Clémençon. Affine-invariant integrated rank-weighted depth: Definition, properties and finite sample analysis. *arXiv preprint arXiv:2106.11068*, 2021a.

Guillaume Staerman, Pavlo Mozharovskyi, Pierre Colombo, Stéphan Clémençon, and Florence d'Alché Buc. A pseudo-metric between probability distributions based on depth-trimmed regions. *arXiv preprint arXiv:2103.12711*, 2021b.

Guillaume Staerman, Eric Adjakossa, Pavlo Mozharovskyi, Vera Hofer, Jayant Sen Gupta, and Stephan Clémençon. Functional anomaly detection: a benchmark study. *arXiv preprint arXiv:2201.05115*, 2022.

Andreas Steiner, Alexander Kolesnikov, , Xiaohua Zhai, Ross Wightman, Jakob Uszkoreit, and Lucas Beyer. How to train your vit? data, augmentation, and regularization in vision transformers. *arXiv preprint arXiv:2106.10270*, 2021.

Adarsh Subbaswamy and Suchi Saria. From development to deployment: dataset shift, causality, and shift-stable models in health ai. *Biostatistics*, 21(2):345–352, April 2020. ISSN 1465-4644. doi: 10.1093/biostatistics/kxz041.

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *2nd International Conference on Learning Representations*, 2014.

Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. Mlp-mixer: An all-mlp architecture for vision. *arXiv preprint arXiv:2105.01601*, 2021.

Florian Tramer, Nicholas Carlini, Wieland Brendel, and Aleksander Madry. On adaptive attacks to adversarial example defenses. *Advances in Neural Information Processing Systems*, 33:1633–1645, 2020.

John W. Tukey. Mathematics and the picturing of data. In *Proceedings of the International Congress of Mathematicians*, volume 2, pages 523–531, 1975.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

Xiaosen Wang, Xuanran He, Jingdong Wang, and Kun He. Admix: Enhancing the transferability of adversarial attacks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16158–16167, 2021a.

Yisen Wang, Difan Zou, Jinfeng Yi, James Bailey, Xingjun Ma, and Quanquan Gu. Improving adversarial robustness requires revisiting misclassified examples. In *International Conference on Learning Representations*, 2019.

Yulin Wang, Rui Huang, Shiji Song, Zeyi Huang, and Gao Huang. Not all images are worth 16x16 words: Dynamic transformers for efficient image recognition. *Advances in Neural Information Processing Systems*, 34, 2021b.

Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. *Advances in Neural Information Processing Systems*, 33, 2020a.

Kaiwen Wu, Allen Wang, and Yaoliang Yu. Stronger and faster wasserstein adversarial attacks. In *International Conference on Machine Learning*, pages 10377–10387. PMLR, 2020b.

Yong Xie, Dakuo Wang, Pin-Yu Chen, Jinjun Xiong, Sijia Liu, and Oluwasanmi Koyejo. A word is worth a thousand dollars: Adversarial attack on tweets fools stock prediction. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 587–599, Seattle, United States, July 2022. Association for Computational Linguistics. URL https://aclanthology.org/2022.naacl-main.43.

Weilin Xu, David Evans, and Yanjun Qi. Feature squeezing: Detecting adversarial examples in deep neural networks. In *25th Annual Network and Distributed System Security Symposium*. The Internet Society, 2018.

Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. *CVPR*, 2022.

Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P Xing, Laurent El Ghaoui, and Michael I Jordan. Theoretically principled trade-off between robustness and accuracy. In *International Conference on Machine Learning*, pages 1–11, 2019.

Zhengyu Zhao, Zhuoran Liu, and Martha Larson. Adversarial color enhancement: Generating unrestricted adversarial images by optimizing a color filter. *arXiv preprint arXiv:2002.01008*, 2020.

Zhihao Zheng and Pengyu Hong. Robust detection of adversarial attacks by modeling the intrinsic properties of deep neural networks. In *Advances in Neural Information Processing Systems 31*, pages 7924–7933, 2018.

Yijun Zuo and Robert Serfling. General notions of statistical depth function. *The Annals of Statistics*, 28(2):461–482, 2000.

A TRAINING DETAILS

We compare the different detection methods on three vision datasets: CIFAR10, CIFAR100 (Krizhevsky et al.) and Tiny ImageNet (Jiao et al., 2019) for which we use the ViT models presented in Sec. 4.1 to build a classifier.

We trained two different models: a ViT, and a ResNet18. The ResNet18 has been trained on 100 epochs, with a Stochastic Gradient Descent (SGD) optimizer, with a learning rate of 0.1, a momentum of 0.9, and a weight decay of 10^{-5}. We use the base model with 16 layers (85.8 million of parameters) from https://github.com/jeonsworld/ViT-pytorch trained on ImageNet (Deng et al., 2009) as our ViT classifier for CIFAR10 and CIFAR100. To train it we set the batch size to 512. The learning rate of SGD (Ruder, 2016) is set to 3×10^{-2} and we use 500 warming steps with no gradient accumulation (Vaswani et al., 2017). For Tiny ImageNet, we used as the underlying classifier a ViT with 16 layers, trained by Huynh (2022) and available at https://github.com/ehuynh1106/TinyImageNet-Transformers. Note that we only use the class token to output the layer-wise input's representations.

Remark. We compare our proposed APPROVED method with FS and MagNet, recalled in Sec. 2.2. We train MagNet according to its original training procedure, while FS and our APPROVED, presented in Sec. 3.2, do not require any training.

B APPROXIMATION ALGORITHM

In this appendix, we display the algorithm used to compute the IRW depth (see Algorithm 1).

Algorithm 1 Approximation of the IRW depth

Initialization: test sample x, n_{proj}, $\mathbf{X} = [x_1, \ldots, x_n]^\top$.

1: Construct $\mathbf{U} \in \mathbb{R}^{d \times n_{\text{proj}}}$ by sampling uniformly n_{proj} vectors $U_1, \ldots, U_{n_{\text{proj}}}$ in $\mathbb{S}^{d-1}$

2: Compute $\mathbf{M} = \mathbf{X}\mathbf{U}$ and $x^\top \mathbf{U}$

3: Compute the rank value $\sigma(j)$, the rank of $x^\top \mathbf{U}$ in $\mathbf{M}_{:,j}$ for every $j \leq n_{\text{proj}}$

4: Set $D = \frac{1}{n_{\text{proj}}} \sum_{j=1}^{n_{\text{proj}}} \sigma(j)$

Output: $\widetilde{D}^{\text{MC}}_{\text{IRW}}(x, \mathbf{X}) = D$

Complexity. The complexity of the algorithm is detailed as follows. Line 1 requires sampling n_{proj} Gaussian samples and normalizing them in order to define unit sphere directions and can be computed in $\text{O}(n_{\text{proj}}d)$. Line 2 requires $\text{O}(n_{\text{proj}}dn)$ to project data on the n_{proj} unit sphere Monte-Carlo directions. Line 3 requires computing the sorting operation on n_{proj} columns of the matrix M and then leads to a complexity of $\text{O}(n_{\text{proj}}n)$. Line 4 requires the computation of the mean and can be done in n_{proj} operations. Finally, the total complexity of the algorithm is then in $\text{O}(n_{\text{proj}}dn)$ which is linear in all of its parameters.

Remarks. Given that the algorithm is linear in all its parameters, computing the IRW depth can be scaled to any datasets. Note that the IRW data depth makes no assumption on the training distribution. In line 3 of Algorithm 1, "rank values" consists in ranking the elements of the projection of each input on U. This is achieved by a sorting algorithm. This step allows us to define an ordering of the projected inputs, which is used to compute the final depth score.

C TIME AND COMPUTATIONAL REQUIREMENTS

C.1 TO GENERATE ATTACKS

We here present the computational requirements to generate the attacks on the transformer, along with the required time to generate them. We use the Adversarial-Robustness Toolbox (ART) (Nicolae et al., 2018) to generate the attacks.

Table 4: Resources and time needed to generate different types of attack on CIFAR10

Attack	GPUs	CPUs	Time
FGSM	V100-32G	20G	0h25
BIM	V100-32G	20G	3h13
PGD	V100-32G	20G	4h30
DF	V100-32G	20G	1h54
HOP	V100-32G	20G	47h39
CW^{∞}	V100-32G	30G	2h48
SA	V100-32G	20G	5h04
STA	V100-32G	20G	1h25

C.2 To deploy detectors

This section presents the computational requirements, along with the time needed to deploy each of the studied detection methods on CIFAR10. For FS and MagNet, we use the codes available at https://github.com/aldahdooh/detectors_review.

Table 5: Resources and time needed to train and test each detection method

Method	GPUs	CPUs	Training Time	Testing Time
APPROVED	V100-32G	40G	N/A	0h11
FS	V100-32G	80G	N/A	0h53
MagNet	V100-32G	180G	3h01	0h13

D Success Rates of Attacks on CIFAR10

We here report the success rate per attack for all the different threat mechanisms (i.e., PGD^1, PGD^2, PGD^{∞}, BIM, FGSM, CW^{∞}, SA, STA, DF and HOP). In orange are the attack performances on ViT while the ones on ResNet are in green (see Sec. 4 for a detailed analysis).

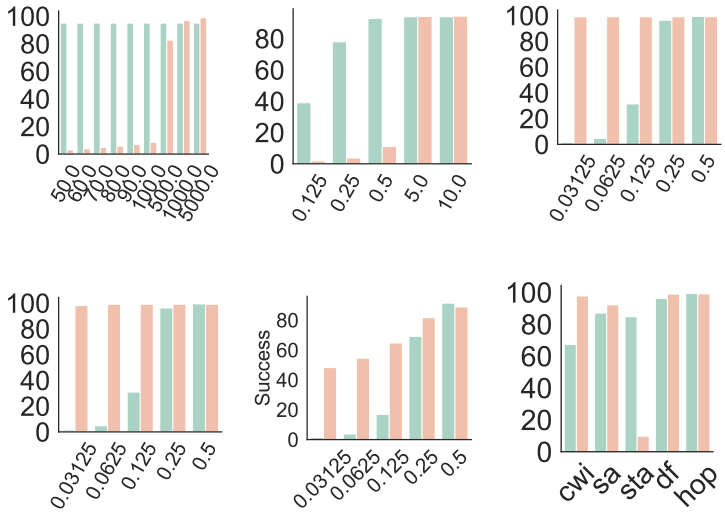

Figure 6: Percentage of successful attacks depending on the L_p-norm constraint, the maximal perturbation ε and the attack algorithm on ResNet18 (orange) and ViT (blue).

E DETAILED RESULTS FOR CIFAR10, CIFAR100, AND TINY IMAGENET

E.1 DETAILED TABLES

In Tab. 6, Tab. 7 and Tab. 8, we present the detailed results for CIFAR10, CIFAR100 and Tiny ImageNet under multiple threats. From Tab. 6, Tab. 7 and Tab. 8, it is straightforward to conclude

Table 6: AUROC↑ and FPR$\downarrow_{90\%}$ for each considered attack mechanisms, L_p-norm constraint and ε on CIFAR10 for APPROVED, FS, and MagNet. The best result for each attack is shown in **bold**.

CIFAR10						
Norm L1	APPROVED		FS		MagNet	
	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$
PGD[1]						
$\varepsilon = 50$	**97.2**	**5.0**	77.6	37.5	53.3	90.1
$\varepsilon = 60$	**97.0**	**5.7**	77.4	37.5	51.6	92.1
$\varepsilon = 70$	**96.4**	**6.8**	78.0	31.2	51.9	92.0
$\varepsilon = 80$	**95.7**	**8.6**	78.1	31.2	51.3	91.9
$\varepsilon = 90$	**94.8**	**11.1**	78.7	31.2	52.0	91.6
$\varepsilon = 100$	**93.9**	**13.9**	79.0	37.5	51.6	91.6
$\varepsilon = 500$	80.1	50.1	**86.8**	**25.0**	49.6	90.5
$\varepsilon = 1000$	**93.0**	**14.2**	83.7	37.5	49.9	90.0
$\varepsilon = 5000$	**98.0**	**3.6**	76.0	55.2	50.1	89.9
Norm L2	APPROVED		FS		MagNet	
	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$
PGD[2]						
$\varepsilon = 0.125$	**97.1**	**4.5**	75.5	37.5	50.6	92.1
$\varepsilon = 0.25$	**97.1**	**5.5**	77.2	37.5	52.2	91.7
$\varepsilon = 0.5$	**92.6**	**18.1**	79.8	31.2	50.6	91.6
$\varepsilon = 5$	**93.3**	**13.6**	77.0	45.9	50.0	89.8
$\varepsilon = 10$	**94.1**	**11.5**	76.8	52.1	50.1	89.8
HOP						
$\varepsilon = 0.1$	**98.3**	**3.3**	74.5	25.0	53.4	83.6
DeepFool						
No ε	**86.5**	45.4	79.7	**31.2**	50.3	89.7
Norm L_∞	APPROVED		FS		MagNet	
	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$
PGD$^\infty$						
$\varepsilon = 0.03125$	**96.5**	**6.4**	78.7	42.9	50.3	89.6
$\varepsilon = 0.0625$	**99.1**	**2.1**	73.4	64.7	51.0	88.4
$\varepsilon = 0.125$	**99.7**	**0.8**	71.8	68.6	52.9	85.5
$\varepsilon = 0.25$	**99.8**	**0.5**	70.9	70.0	54.3	83.4
$\varepsilon = 0.5$	**99.8**	**0.5**	70.8	70.1	54.4	83.3
BIM						
$\varepsilon = 0.03125$	**88.3**	**27.0**	74.0	64.5	50.3	89.6
$\varepsilon = 0.0625$	**97.1**	**5.4**	70.2	72.3	50.7	88.9
$\varepsilon = 0.125$	**99.0**	**2.2**	70.0	72.2	51.8	87.2
$\varepsilon = 0.25$	**99.7**	**0.7**	70.7	70.5	53.6	84.4
$\varepsilon = 0.5$	**99.9**	**0.2**	71.2	68.4	56.4	80.1
FGSM						
$\varepsilon = 0.03125$	**78.1**	69.5	75.2	**38.8**	51.9	88.1
$\varepsilon = 0.0625$	**82.4**	60.2	77.2	**37.5**	53.0	86.1
$\varepsilon = 0.125$	**93.1**	**16.6**	78.9	31.2	57.3	79.2
$\varepsilon = 0.25$	**99.1**	**1.6**	69.6	25.0	70.6	54.8
$\varepsilon = 0.5$	**99.7**	**0.6**	67.7	31.2	80.4	18.0
SA						
$\varepsilon = 0.125$	**98.2**	**3.3**	72.0	25.0	55.1	82.4
CW$^\infty$						
$\varepsilon = 0.3125$	**90.4**	**30.6**	78.8	37.5	50.6	89.3
No Norm	APPROVED		FS		MagNet	
	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$
STA						
No ε	**94.9**	**10.5**	78.8	37.5	39.4	93.5

that APPROVED significantly outperforms FS, and MagNet.

Table 7: AUROC↑ and FPR$\downarrow_{90\%}$ for each considered attack mechanisms, L_p-norm constraint and ε on CIFAR100 for APPROVED, FS and MagNet. The best result for each attack is shown in **bold**.

CIFAR100						
Norm L1	APPROVED		FS		MagNet	
	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$
PGD[1]						
$\varepsilon = 50$	**83.5**	**39.3**	65.5	56.2	50.5	90.5
$\varepsilon = 60$	**82.4**	**41.0**	66.6	56.2	50.5	90.3
$\varepsilon = 70$	**81.2**	**45.3**	67.4	50.0	50.0	90.4
$\varepsilon = 80$	**79.8**	**47.8**	68.3	50.0	50.0	90.4
$\varepsilon = 90$	**78.4**	**50.0**	69.2	**50.0**	50.2	90.3
$\varepsilon = 100$	**77.0**	54.0	70.1	**50.0**	50.1	90.4
$\varepsilon = 500$	58.1	75.5	**79.3**	**50.0**	50.0	90.0
$\varepsilon = 1000$	78.3	**44.9**	**80.0**	62.5	50.0	89.9
$\varepsilon = 5000$	**86.1**	**29.4**	74.0	75.0	50.0	89.8
Norm L2	APPROVED		FS		MagNet	
	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$
PGD[2]						
$\varepsilon = 0.125$	**84.3**	**38.3**	64.6	56.2	50.8	90.8
$\varepsilon = 0.25$	**82.7**	**41.4**	66.2	56.2	50.8	90.1
$\varepsilon = 0.5$	**73.9**	59.1	72.0	**50.0**	50.3	90.0
$\varepsilon = 5$	**78.6**	**43.5**	75.1	75.0	50.0	89.9
$\varepsilon = 10$	**79.4**	**41.0**	74.4	75.0	50.0	89.9
HOP						
$\varepsilon = 0.1$	**89.1**	**24.8**	62.7	50.0	52.1	84.5
DeepFool						
No ε	**75.5**	59.9	62.2	**50.0**	50.0	89.9
Norm L_∞	APPROVED		FS		MagNet	
	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$
PGD$^\infty$						
$\varepsilon = 0.03125$	75.4	**51.5**	**76.0**	74.8	50.2	89.7
$\varepsilon = 0.0625$	**88.1**	**26.0**	68.9	75.0	50.6	88.9
$\varepsilon = 0.125$	**93.3**	**14.9**	65.5	75.0	52.1	86.5
$\varepsilon = 0.25$	**94.4**	**12.8**	64.3	75.0	53.0	84.9
$\varepsilon = 0.5$	**89.7**	**26.4**	64.2	75.0	53.1	84.8
BIM						
$\varepsilon = 0.03125$	63.1	**72.9**	**67.6**	75.0	50.2	89.7
$\varepsilon = 0.0625$	**70.5**	**64.8**	63.0	81.1	50.5	89.2
$\varepsilon = 0.125$	**87.2**	**28.1**	62.1	82.7	51.3	87.8
$\varepsilon = 0.25$	**93.2**	**15.4**	63.7	75.4	52.5	85.7
$\varepsilon = 0.5$	**96.5**	**8.3**	65.3	75.0	54.6	82.2
FGSM						
$\varepsilon = 0.03125$	**80.8**	**48.1**	61.9	62.5	51.0	88.8
$\varepsilon = 0.0625$	**86.5**	**33.0**	61.3	61.4	52.1	86.8
$\varepsilon = 0.125$	**90.4**	**24.0**	54.8	50.0	55.8	80.2
$\varepsilon = 0.25$	**95.7**	**10.3**	49.6	50.0	66.4	60.4
$\varepsilon = 0.5$	**98.6**	**4.1**	46.2	56.2	86.6	24.2
SA						
$\varepsilon = 0.125$	**89.6**	**26.0**	63.3	50.0	54.9	82.6
CW$^\infty$						
$\varepsilon = 0.3125$	**81.7**	**42.2**	67.0	50.0	50.0	89.8
No Norm	APPROVED		FS		MagNet	
	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$
STA						
No ε	**87.4**	**32.1**	65.4	50.0	38.3	92.8

Table 8: AUROC↑ and FPR$\downarrow_{90\%}$ for each considered attack mechanisms, L_p-norm constraint and ε on Tiny ImageNet for APPROVEDand FS. The best result for each attack is shown in **bold**.

	Tiny ImageNet					
Norm L1	APPROVED		FS		MagNet	
	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$
PGD[1]						
ε = 50	**74.2**	**61.1**	44.8	81.6	50.4	88.9
ε = 60	**74.3**	**60.7**	45.0	81.8	50.3	88.9
ε = 70	**74.8**	**60.7**	45.1	82.0	50.0	89.0
ε = 80	**74.7**	**60.5**	45.1	82.3	49.6	88.9
ε = 90	**74.9**	**59.8**	45.0	82.2	49.7	89.3
ε = 100	**74.6**	**59.4**	44.9	82.0	49.6	89.0
ε = 500	**76.5**	**59.7**	60.7	71.7	48.0	93.1
ε = 1000	**74.2**	**59.4**	73.7	62.4	47.6	92.0
ε = 5000	78.2	51.8	**83.2**	**50.0**	49.1	90.3
Norm L2	APPROVED		FS		MagNet	
	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$	AUROC↑	FPR$\downarrow_{90\%}$
PGD[2]						
ε = 0.125	**74.2**	**60.2**	45.2	81.4	50.2	88.7
ε = 0.25	**75.0**	**57.2**	45.2	81.8	49.3	89.7
ε = 0.5	**75.7**	**53.4**	47.1	79.5	49.6	91.0
ε = 5	74.3	60.6	**77.9**	**57.5**	48.7	91.0
ε = 10	74.4	59.7	**78.1**	**57.7**	48.8	90.9
HOP						
ε = 0.1	**87.1**	**31.8**	59.1	76.3	52.7	83.8
Norm L_∞	APPROVED		FS		MagNet	
	AUROC	FPR	AUROC	FPR	AUROC	FPR
PGD$^\infty$						
ε = 0.03125	89.6	28.8	**96.0**	**8.2**	49.7	90.0
ε = 0.0625	**99.1**	**1.9**	93.8	11.9	49.8	89.9
ε = 0.125	**99.9**	**0.0**	89.2	47.1	49.9	89.6
ε = 0.25	**99.9**	**0.0**	85.5	73.6	50.0	89.5
ε = 0.5	**99.9**	**0.0**	83.6	82.2	50.1	89.4
BIM						
ε = 0.03125	80.7	**43.1**	**86.0**	44.8	49.5	90.1
ε = 0.0625	**95.1**	**15.1**	90.3	33.4	49.9	89.9
ε = 0.125	**99.6**	**1.0**	87.4	61.4	49.9	89.8
ε = 0.25	**99.9**	**0.0**	84.9	79.9	50.0	89.5
ε = 0.5	**99.9**	**0.0**	83.9	82.5	50.2	89.1
FGSM						
ε = 0.03125	**74.5**	**55.9**	56.3	75.5	49.7	90.2
ε = 0.0625	**80.8**	**43.5**	58.0	71.8	50.4	89.6
ε = 0.125	**87.1**	**30.4**	53.6	75.1	50.9	88.7
ε = 0.25	**91.1**	**22.3**	48.1	78.8	52.6	86.2
ε = 0.5	**94.4**	**15.2**	50.9	74.2	60.7	72.1
SA						
ε = 0.125	**77.0**	**49.1**	48.7	78.5	50.6	89.4
No Norm	APPROVED		FS		MagNet	
	AUROC	FPR	AUROC	FPR	AUROC	FPR
STA						
No ε	**80.2**	**42.5**	53.0	77.5	34.9	95.6

F PER CLASS ANALYSIS

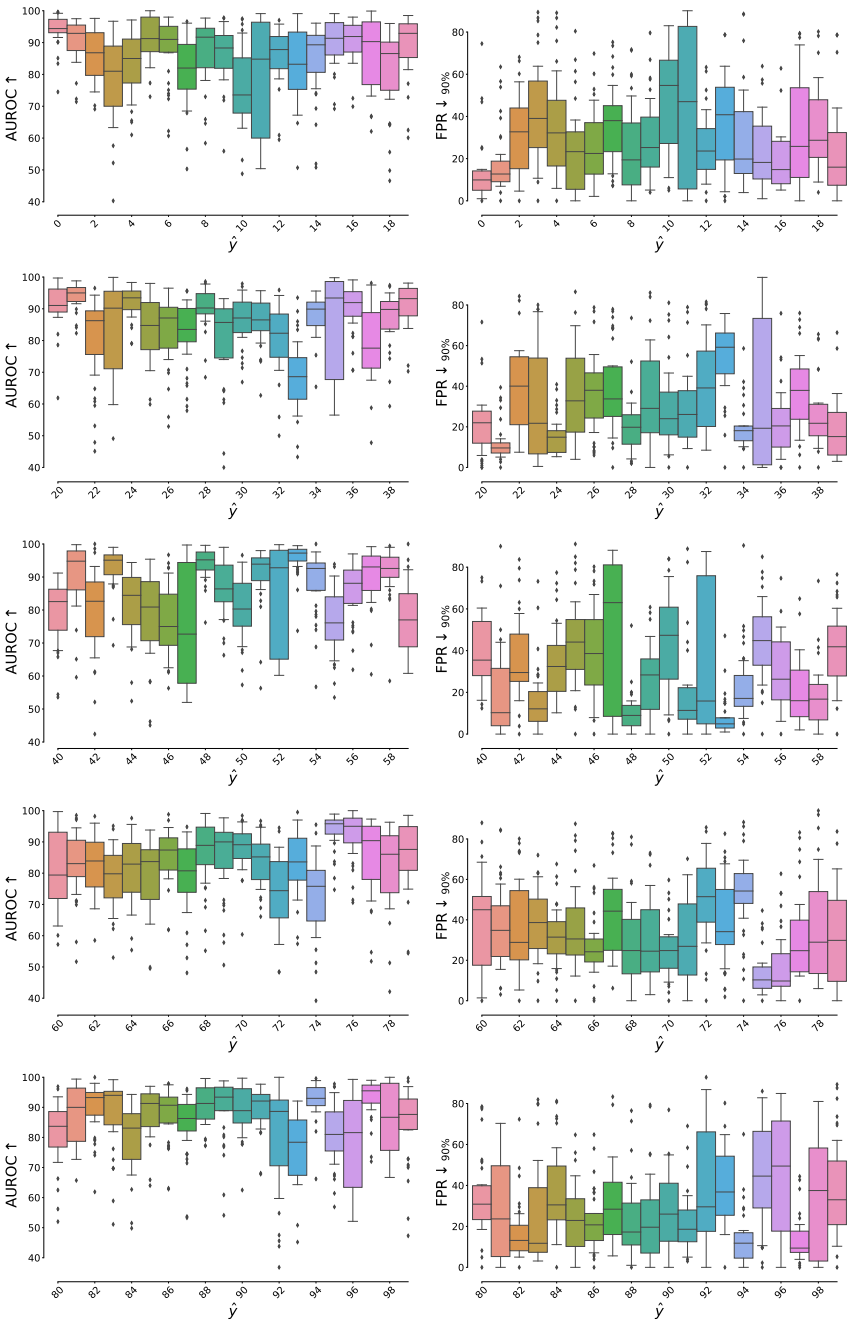

Figure 7: APPROVED's AUROC↑ and FPR$\downarrow_{90\%}$ per class, averaged over the attacks on CIFAR100.

As for CIFAR10 (see Sec. 5), the detector performances depend on the predicted class. Some classes are easy to detect (i.e., classes 0, 21, 53, 75, and 94), others are more difficult (i.e., classes 3, 10, 33, 47, 60, 74, and 93). Some have low variance (i.e., 0, 1, 24, 34, 75, 82 and, 94) while others have an extremely large dispersion (i.e., 11, 35,47, 52, 96, and 98).