

---

# GREAT Score: Global Robustness Evaluation of Adversarial Perturbation using Generative Models

---

**ZAITANG LI**

*The Chinese University of Hong Kong  
Sha Tin, Hong Kong  
ztli@cse.cuhk.edu.hk*

Pin-Yu Chen

*IBM Research  
New York, USA*

*pin-yu.chen@ibm.com*

**Tsung-Yi Ho**

*The Chinese University of Hong Kong  
Sha Tin, Hong Kong  
tyho@cse.cuhk.edu.hk*

## Abstract

Current studies on adversarial robustness mainly focus on aggregating *local* robustness results from a set of data samples to evaluate and rank different models. However, the local statistics may not well represent the true *global* robustness of the underlying unknown data distribution. To address this challenge, this paper makes the first attempt to present a new framework, called **GREAT Score**, for global robustness evaluation of adversarial perturbation using generative models. Formally, GREAT Score carries the physical meaning of a global statistic capturing a mean certified attack-proof perturbation level over all samples drawn from a generative model. For finite-sample evaluation, we also derive a probabilistic guarantee on the sample complexity and the difference between the sample mean and the true mean. GREAT Score has several advantages: (1) Robustness evaluations using GREAT Score are efficient and scalable to large models, by sparing the need of running adversarial attacks. In particular, we show high correlation and significantly reduced computation cost of GREAT Score when compared to the attack-based model ranking on RobustBench [12]. (2) The use of generative models facilitates the approximation of the unknown data distribution. In our ablation study with different generative adversarial networks (GANs), we observe consistency between global robustness evaluation and the quality of GANs. (3) GREAT Score can be used for remote auditing of privacy-sensitive black-box models, as demonstrated by our robustness evaluation on several online facial recognition services.

## Project Demo and Code Page:

<https://huggingface.co/spaces/TrustSafeAI/GREAT-Score>  
<https://github.com/IBM/GREAT-Score>

## 1 Introduction

Adversarial robustness is the study of model performance in the worst-case scenario, which is a key element in trustworthy machine learning. Adversarial robustness evaluation refers to the process of assessing a model's resilience against adversarial attacks, which are inputs intentionally designed to deceive the model. Without further remediation, state-of-the-art machine learning models, especially neural networks, are known to be overly sensitive to small human-imperceptible perturbations to data

inputs [19]. Such a property of over-sensitivity could be exploited by bad actors to craft adversarial perturbations leading to prediction-evasive adversarial examples.

Given a threat model specifying the knowledge of the target machine learning model (e.g., white-box or black-box model access) and the setting of plausible adversarial interventions (e.g., norm-bounded input perturbations), the methodology for adversarial robustness evaluation can be divided into two categories: *attack-dependent* and *attack-independent*. Attack-dependent approaches aim to devise the strongest possible attack and use it for performance assessment. A typical example is Auto-Attack [10], a state-of-the-art attack based on an ensemble of advanced white-box and black-box adversarial perturbation methods. On the other hand, attack-independent approaches aim to develop a certified or estimated score for adversarial robustness, reflecting a quantifiable level of attack-proof certificate. Typical examples include neural network verification techniques [60, 66], certified defenses such as randomized smoothing [9], and local Lipschitz constant estimation [59].

Despite a plethora of adversarial robustness evaluation methods, current studies primarily focus on aggregating *local* robustness results from a set of data samples. However, the sampling process of these test samples could be biased and unrepresentative of the true *global* robustness of the underlying data distribution, resulting in the risk of incorrect or biased robustness benchmarks. For instance, we find that when assessing the ranking of Imagenet models through Robustbench [11], using AutoAttack [10] with 10,000 randomly selected samples (the default choice) with 100 independent trials results in an unstable ranking coefficient of  $0.907 \pm 0.0256$  when compared to that of the entire 50,000 test samples. This outcome affirms that AutoAttack's model ranking has notable variations with an undersampled or underrepresented test dataset.

An ideal situation is when the data distribution is transparent and one can draw an unlimited number of samples from the true distribution for reliable robustness evaluation. But in reality, the data distribution is unknown and difficult to characterize. In addition to lacking rigorous global robustness evaluation, many attack-independent methods are limited to the white-box setting, requiring detailed knowledge about the target model (e.g., model parameters and architecture) such as input gradients and internal data representations for robustness evaluation. Moreover, state-of-the-art attack-dependent and attack-independent methods often face the issue of scalability to large models and data volumes due to excessive complexity, such as the computational costs in iterative gradient computation and layer-wise interval bound propagation and relaxation [20, 31].

To address the aforementioned challenges including (i) lack of proper global adversarial robustness evaluation, (ii) limitation to white-box settings, and (iii) computational inefficiency, in this paper we present a novel attack-independent evaluation framework called *GREAT Score*, which is short for *global robustness evaluation of adversarial perturbation using generative models*. We tackle challenge (i) by using a generative model such as a generative adversarial network (GAN) [17, 18] or a diffusion model [27] as a proxy of the true unknown data distribution. Formally, GREAT Score is defined as the mean of a certified lower bound on minimal adversarial perturbation over the data sampling distribution of a generative model, which represents the global distribution-wise adversarial robustness with respect to the generative model in use. It entails a global statistic capturing the mean certified attack-proof perturbation level over all samples from a generative model. For finite-sample evaluation, we also derive a probabilistic guarantee quantifying the sample complexity and the difference between the sample mean and true mean.

For challenge (ii), our derivation of GREAT Score leads to a neat closed-form solution that only requires data forward-passing and accessing the model outputs, which applies to any black-box classifiers giving class prediction confidence scores as model output. Moreover, as a byproduct of using generative models, our adversarial robustness evaluation procedure is executed with only synthetically generated data instead of real data, which is particularly appealing to privacy-aware robustness assessment schemes, e.g., remote robustness evaluation or auditing by a third party with restricted access to data and model. We will present how GREAT Score can be used to assess the robustness of online black-box facial recognition models. Finally, for challenge (iii), GREAT Score is applicable to any off-the-shelf generative models so that we do not take the training cost of generative models into consideration. Furthermore, the computation of GREAT Score is lightweight because it scales linearly with the number of data samples used for evaluation, and each data sample only requires one forward pass through the model to obtain the final predictions.

We highlight our main contributions as follows:

- We present GREAT Score as a novel framework for deriving a global statistic representative of the distribution-wise robustness to adversarial perturbation, based on an off-the-shelf generative model for approximating the data generation process.
- Theoretically, we show that GREAT Score corresponds to a mean certified attack-proof level of  $\mathcal{L}_2$ -norm bounded input perturbation over the sampling distribution of a generative model (Theorem 1). We further develop a formal probabilistic guarantee on the quality of using the sample mean as GREAT Score with a finite number of samples from generative models (Theorem 2).
- We evaluate the effectiveness of GREAT Score on all neural network models on RobustBench [11] (the largest adversarial robustness benchmark), with a total of 17 models on CIFAR-10 and 5 models on ImageNet. We show that the model ranking of GREAT Score is highly aligned with that of the original ranking on RobustBench using AutoAttack [10], while GREAT Score significantly reduces the computation time. Specifically, on CIFAR-10 the computation complexity can be reduced by up to 2,000 times. The results suggest that GREAT Score is a competitive and computationally-efficient approach complementary to attack-based robustness evaluations.
- As a demonstration of GREAT Score's capability for remote robustness evaluation of access-limited systems, we show how GREAT Score can audit several online black-box facial recognition APIs.

## 2 Background and Related Works

**Adversarial Attack and Defense.** Adversarial attacks aim to generate examples that can evade classifier predictions in classification tasks. In principle, adversarial examples can be crafted by small perturbations to a native data sample, where the level of perturbation is measured by different  $\mathcal{L}_p$  norms [7, 8, 58]. The procedure of finding adversarial perturbation within a perturbation level is often formulated as a constrained optimization problem, which can be solved by algorithms such as projected gradient descent (PGD) [38]. The state-of-the-art adversarial attack is the Auto-Attack [10], which uses an ensemble of white-box and black-box attacks. There are many methods (defenses) to improve adversarial robustness. A popular approach is adversarial training [38], which generates adversarial perturbation during model training for improved robustness. One common evaluation metric for adversarial robustness is robust accuracy, which is defined as the accuracy of correct classification under adversarial attacks, evaluated on a set of data samples. RobustBench [10] is the largest-scale standardized benchmark that ranks the models using robust accuracy against Auto-Attack on test sets from image classification datasets such as CIFAR-10. In addition to discussed works, several studies evaluate model robustness differently. [43] introduce adversarial sparsity, quantifying the difficulty of finding perturbations, providing insights beyond adversarial accuracy. [48] propose probabilistic robustness, balancing average and worst-case performance by enforcing robustness to most perturbations, better addressing trade-offs. [22] introduce the adversarial hypervolume metric, a comprehensive measure of robustness across varying perturbation intensities.

**Generative Models.** Statistically speaking, let  $X$  denote the observable variable and let  $Y$  denote the corresponding label, the learning objective for a generative model is to model the conditional probability distribution  $P(X | Y)$ . Among all the generative models, GANs have gained a lot of attention in recent years due to their capability to generate realistic high-quality images [18]. The principle of training GANs is based on the formulation of a two-player zero-sum min-max game to learn the high-dimension data distribution. Eventually, these two players reach the Nash equilibrium that  $D$  is unable to further discriminate real data versus generated samples. This adversarial learning methodology aids in obtaining high-quality generative models. In practice, the generator  $G(\cdot)$  takes a random vector  $z$  (i.e., a latent code) as input, which is generated from a zero-mean isotropic Gaussian distribution denoted as  $z \sim \mathcal{N}(0, I)$ , where  $I$  means an identity matrix. Conditional GANs refer to the conditional generator  $G(\cdot|Y)$  given a class label  $Y$ . In addition to GAN, diffusion models (DMs) are also gaining popularity. DMs consist of two stages: the forward diffusion process and the reverse diffusion process. In the forward process, the input data is gradually perturbed by Gaussian Noises and becomes an isotropic Gaussian distribution eventually. In the reverse process, DMs reverse the forward process and implement a sampling process from Gaussian noises to reconstruct the true samples by solving a stochastic differential equation. In our proposed framework, we use off-the-shelf (conditional) GANs and DMs (e.g., DDPM [27]) that are publicly available as our generative models.

**Formal Local Robustness Guarantee and Estimation.** Given a data sample  $x$ , a formal local robustness guarantee refers to a certified range on its perturbation level such that within which the top-1 class prediction of a model will remain unchanged [26]. In  $\mathcal{L}_p$ -norm ( $p \geq 1$ ) bounded perturbations

centered at  $x$ , such a guarantee is often called a certified radius  $r$  such that any perturbation  $\delta$  to  $x$  within this radius (i.e.,  $\|\delta\|_p \leq r$ ) will have the same top-1 class prediction as  $x$ . Therefore, the model is said to be provably locally robust (i.e., attack-proof) to any perturbations within the certified radius  $r$ . By definition, the certified radius of  $x$  is also a lower bound on the minimal perturbation required to flip the model prediction.

Among all the related works on attack-independent local robustness evaluations, the CLEVER framework proposed in [59] is the closest to our study. The authors in [59] derived a closed-form of certified local radius involving the maximum local Lipschitz constant of the model output with respect to the data input around a neighborhood of a data sample  $x$ . They then proposed to use extreme value theory to estimate such a constant and use it to obtain a local robustness score, which is not a certified local radius. Our proposed GREAT Score has major differences from [59] in that our focus is on global robustness evaluation, and our GREAT Score is the mean of a certified radius over the sampling distribution of a generative model. In addition, for every generated sample, our local estimate gives a certified radius.

**Notations.** All the main notations used in the paper are summarized in Appendix A.

### 3 GREAT Score: Methodology and Algorithms

#### 3.1 True Global Robustness and Certified Estimate

Let  $f = [f_1, \dots, f_K] : \mathbb{R}^d \rightarrow \mathbb{R}^K$  denote a fixed  $K$ -way classifier with flattened data input of dimension  $d$ ,  $(x, y)$  denote a pair of data sample  $x$  and its corresponding groundtruth label  $y \in \{1, \dots, K\}$ ,  $P$  denote the true data distribution which in practice is unknown, and  $\Delta_{\min}(x)$  denote the minimal perturbation of a sample-label pair  $(x, y) \sim P$  causing the change of the top-1 class prediction such that  $\arg \max_{k \in \{1, \dots, K\}} f_k(x + \Delta_{\min}(x)) \neq \arg \max_{k \in \{1, \dots, K\}} f_k(x)$ . Note that if the model  $f$  makes an incorrect prediction on  $x$ , i.e.,  $y \neq \arg \max_{k \in \{1, \dots, K\}} f_k(x)$ , then we define  $\Delta_{\min}(x) = 0$ . This means the model is originally subject to prediction evasion on  $x$  even without any perturbation. A higher  $\Delta_{\min}(x)$  means better local robustness of  $f$  on  $x$ .

The following statement defines the true global robustness of a classifier  $f$  based on the probability density function  $p(\cdot)$  of the underlying data distribution  $P$ .

**Definition 1** (True global robustness w.r.t.  $P$ ). The true global robustness of a classifier  $f$  with respect to a data distribution  $P$  is defined as:

$$\Omega(f) = \mathbb{E}_{x \sim P}[\Delta_{\min}(x)] = \int_{x \sim P} \Delta_{\min}(x) p(x) dx \quad (1)$$

Unless the probability density function of  $P$  and every local minimal perturbation are known, the exact value of the true global robustness cannot be computed. An alternative is to estimate such a quantity. Extending Definition 1, let  $g(x)$  be a local robustness statistic. Then the corresponding global robustness estimate is defined as

$$\hat{\Omega}(f) = \mathbb{E}_{x \sim P}[g(x)] = \int_{x \sim P} g(x) p(x) dx \quad (2)$$

Furthermore, if one can prove that  $g(x)$  is a valid lower bound on  $\Delta_{\min}(x)$  such that  $g(x) \leq \Delta_{\min}(x)$ ,  $\forall x$ , then the estimate  $\hat{\Omega}(f)$  is said to be a certified lower bound on the true global robustness with respect to  $P$ , and larger  $\hat{\Omega}(f)$  will imply better true global robustness. In what follows, we will formally introduce our proposed GREAT Score and show that it is a certified estimate of the lower bound on the true robustness with respect to the data-generating distribution learned by a generative model.

#### 3.2 Using GMs to Evaluate Global Robustness

Recall that a generative model (GM) takes a random vector  $z \sim \mathcal{N}(0, I)$  sampled from a zero-mean isotropic Gaussian distribution as input to generate a data sample  $G(z)$ . In what follows, we present our first main theorem that establishes a certified lower bound  $\hat{\Omega}(f)$  on the true global robustness of a classifier  $f$  measured by the data distribution given by  $G(\cdot)$ .

Without loss of generality, we assume that all data inputs are confined in the scaled data range  $[0, 1]^d$ , where  $d$  is the size of any flattened data input. The  $K$ -way classifier  $f : [0, 1]^d \mapsto \mathbb{R}^K$  takes a data sample  $x$  as input and outputs a  $K$ -dimensional vector  $f(x) = [f_1(x), \dots, f_K(x)]$  indicating the likelihood of its prediction on  $x$  over  $K$  classes, where the top-1 class prediction is defined as  $\hat{y} = \arg \max_{k \in \{1, \dots, K\}} f_k(x)$ . We further denote  $c$  as the groundtruth class of  $x$ . Therefore, if  $\hat{y} \neq c$ , then the classifier is said to make a wrong top-1 prediction. When considering the adversarial robustness on a wrongly classified sample  $x$ , we define the minimal perturbation for altering model prediction as  $\Delta_{\min}(x) = 0$ . The intuition is that an attacker does not need to take any action to make the sample  $x$  evade the correct prediction by  $f$ , and therefore the required minimal adversarial perturbation level is 0 (i.e., zero robustness).

Given a generated data sample  $G(z)$ , we now formally define a local robustness score function as

$$g(G(z)) = \sqrt{\frac{\pi}{2}} \cdot \max\{f_c(G(z)) - \max_{k \in \{1, \dots, K\}, k \neq c} f_k(G(z)), 0\} \quad (3)$$

The scalar  $\sqrt{\pi/2}$  is a constant associated with the sampling Gaussian distribution of  $G$ , which will be apparent in later analysis. We further offer several insights into understanding the physical meaning of the considered local robustness score in (3): (i) The inner term  $f_c(G(z)) - \max_{k \in \{1, \dots, K\}, k \neq c} f_k(G(z))$  represents the gap in the likelihood of model prediction between the correct class  $c$  and the most likely class other than  $c$ . A positive and larger value of this gap reflects higher confidence of the correct prediction and thus better robustness. (ii) Following (i), a negative gap means the model is making an incorrect prediction, and thus the outer term  $\max\{\text{gap}, 0\} = 0$ , which corresponds to zero robustness.

Next, we use the local robustness score  $g$  defined in (3) to formally state our theorem on establishing a certified lower bound on the true global robustness and the proof sketch.

**Theorem 1** (certified global robustness estimate). *Let  $f : [0, 1]^d \mapsto [0, 1]^K$  be a  $K$ -way classifier and let  $f_k(\cdot)$  be the predicted likelihood of class  $k$ , with  $c$  denoting the groundtruth class. Given a generator  $G$  such that it generates a sample  $G(z)$  with  $z \sim \mathcal{N}(0, I)$ . Define  $g(G(z)) = \sqrt{\frac{\pi}{2}} \cdot \max\{f_c(G(z)) - \max_{k \in \{1, \dots, K\}, k \neq c} f_k(G(z)), 0\}$ . Then the global robustness estimate of  $f$  evaluated with  $\mathcal{L}_2$ -norm bounded perturbations, defined as  $\hat{\Omega}(f) = \mathbb{E}_{z \sim \mathcal{N}(0, I)}[g(G(z))]$ , is a certified lower bound of the true global robustness  $\Omega(f)$  with respect to  $G$ .*

The complete proof is given in Appendix C.

### 3.3 Probabilistic Guarantee on Sample Mean

As defined in Theorem 1, the global robustness estimate  $\hat{\Omega}(f) = \mathbb{E}_{z \sim \mathcal{N}(0, I)}[g(G(z))]$  is the mean of the local robustness score function introduced in (3) evaluated through a generator  $G$  and its sampling distribution. In practice, one can use a finite number of samples  $\{G(z_i|y_i)\}_{i=1}^n$  generated from a conditional generator  $G(\cdot|y)$  to estimate  $\hat{\Omega}(f)$ , where  $y$  denotes a class label and it is also an input parameter to the conditional generator. The simplest estimator of  $\hat{\Omega}(f)$  is the sample mean, defined as

$$\hat{\Omega}_S(f) = \frac{1}{n} \sum_{i=1}^n g(G(z_i|y_i)) \quad (4)$$

In what follows, we present our second main theorem to deliver a probabilistic guarantee on the sample complexity to achieve  $\epsilon$  difference between the sample mean  $\hat{\Omega}_S(f)$  and the true mean  $\hat{\Omega}(f)$ .

**Theorem 2** (probabilistic guarantee on sample mean). *Let  $f$  be a  $K$ -way classifier with its outputs bounded by  $[0, 1]^K$  and let  $e$  denote the natural base. For any  $\epsilon, \delta > 0$ , if the sample size  $n \geq \frac{32e \cdot \log(2/\delta)}{\epsilon^2}$ , then with probability at least  $1 - \delta$ , the sample mean  $\hat{\Omega}_S(f)$  is  $\epsilon$ -close to the true mean  $\hat{\Omega}(f)$ . That is,  $|\hat{\Omega}_S(f) - \hat{\Omega}(f)| \leq \epsilon$ .*

The complete proof is given in Appendix D. The proof is built on a concentration inequality in [40]. It is worth noting that the bounded output assumption of the classifier  $f$  in Theorem 2 can be easily satisfied by applying a normalization layer at the final model output, such as the softmax function or the element-wise sigmoid function.

### 3.4 Algorithm and Computational Complexity

Algorithm 1 summarizes the procedure of computing GREAT Score using the sample mean estimator. It can be seen that the computation complexity of GREAT Score is linear in the number of generated samples  $N_S$ , and for each sample, the computation of the statistic  $g$  defined in (3) only requires drawing a sample from the generator  $G$  and taking a forward pass to the classifier  $f$  to obtain the model predictions on each class. As a byproduct, GREAT Score applies to the setting when the classifier  $f$  is a black-box model, meaning only the model outputs are observable by an evaluator.

---

#### Algorithm 1: GREAT Score Computation

---

**Input:**  $K$ -way classifier  $f(\cdot)$ , conditional generator  $G(\cdot)$ , local score function  $g(\cdot)$  defined in (3), number of generated samples  $N_S$

**Output:** GREAT Score  $\hat{\Omega}_S(f)$

**for**  $i \leftarrow 1$  **to**  $N_S$  **do**

    Randomly select a class label  $y \in \{1, 2, \dots, K\}$

    Sample  $z \sim \mathcal{N}(0, I)$  from a Gaussian distribution and generate a sample  $G(z|y)$  with class  $y$

    Pass  $G(z|y)$  into the model  $f$  and get the prediction for each class  $\{f_k(G(z|y))\}_{k=1}^K$

    Record the statistic

$$g^{(i)}(G(z|y)) = \sqrt{\frac{\pi}{2}} \cdot \max\{f_y(G(z|y)) - \max_{k \in \{1, \dots, K\}, k \neq y} f_k(G(z|y)), 0\}$$

**end**

$\hat{\Omega}_S(f) \leftarrow$  Compute the sample mean of  $\{g^{(i)}\}_{i=1}^{N_S}$

---

### 3.5 Calibrated GREAT Score

In cases when one has additional knowledge of adversarial examples on a set of images from a generative model, e.g., successful adversarial perturbations (an upper bound on the minimal perturbation of each sample) returned by any norm-minimization adversarial attack method such as the CW attack [7], the CW attack employs two loss terms, classification loss and distance metric, to generate adversarial examples. See Appendix E for details. We can further “calibrate” the GREAT Score with respect to the available perturbations. Moreover, since Theorem 1 informs some design choices on the model output layer, as long as the model output is a non-negative  $K$ -dimensional vector  $f \in [0, 1]^K$  reflecting the prediction confidence over  $K$  classes, we will incorporate such flexibility in the calibration process.

Specifically, we use calibration in the model ranking setup where there are  $M$  models  $\{f^{(j)}\}_{j=1}^M$  for evaluation, and each model (indexed by  $j$ ) has a set of known perturbations  $\{\delta_i^{(j)}\}_{i=1}^N$  on a common set of  $N$  image-label pairs  $\{x_i, y_i\}_{i=1}^N$  from the same generative model. We further consider four different model output layer designs (that are attached to the model logits): (i)  $\text{sigmoid}(\cdot|T_1)$ : sigmoid with temperature  $T_1$ , (ii)  $\text{softmax}(\cdot|T_2)$ : softmax with temperature  $T_2$ , (iii)  $\text{sigmoid}(\text{softmax}(\cdot|T_2 = 1)|T_1)$ : sigmoid with temperature after softmax, and (iv)  $\text{softmax}(\text{sigmoid}(\cdot|T_1 = 1)|T_2)$ : softmax with temperature after sigmoid. Finally, let  $\{\hat{\Omega}_S(f^{(j)})\}_{j=1}^M$  denote the GREAT Score computed based on  $\{x_i, y_i\}_{i=1}^N$  for each model. We calibrate GREAT Score by optimizing some rank statistics (e.g., the Spearman’s rank correlation coefficient) over the temperature parameter by comparing the ranking consistency between  $\{\hat{\Omega}_S(f^{(j)})\}_{j=1}^M$  and  $\{\delta_i^{(j)}\}_{i=1}^N$ . In our experiments, we find that setting (iv) gives the best result and use it as the default setup for calibration, as detailed in Appendix F.

## 4 Experimental Results

### 4.1 Experiment Setup

**Datasets and Models.** We conduct our experiment on several datasets including CIFAR-10 [32], ImageNet-1K [13] and CelebA-HQ [29]/CelebA [36]. For neural network models, we use the available models on RobustBench [11] (see more details in the next paragraph), which includes 17/5 models on CIFAR-10/ImageNet, correspondingly. We also use several off-the-shelf GANs and




Figure 1: Comparison of local GREAT Score and CW attack in  $\mathcal{L}_2$  perturbation on CIFAR-10 with Rebuffi\_extra model [46]. The x-axis is the image id. The result shows the local GREAT Score is indeed a lower bound of the perturbation level found by CW attack.




Figure 2: Cumulative robust accuracy (RA) with varying  $\mathcal{L}_2$  perturbation level using 500 samples. Note that GREAT Score gives a certified RA for attack-proof robustness, whereas Auto-Attack is an empirical robustness evaluation.

diffusion models (DMs) trained on CIFAR-10 and ImageNet for computing GREAT Score in an ablation study (we defer the model details to later paragraphs).

**Summary of Classifiers on RobustBench.** The RobustBench [49] is to-date the largest benchmark for robustness evaluation with publicly accessible neural network models submitted by contributors. RobustBench uses the default test dataset from several standard image classification tasks, such as CIFAR-10 and ImageNet-1K, to run Auto-Attack [10] and report the resulting accuracy with  $\mathcal{L}_2$ -norm and  $\mathcal{L}_\infty$ -norm perturbations (i.e., the robust accuracy – RA) as a metric for adversarial robustness. Even under one perturbation type, it is not easy to make a direct and fair comparison among all submitted models on RobustBench because they often differ by the training scheme, network architecture, as well as the usage of additional real and/or synthetic data. To make a meaningful comparison with GREAT Score, we select all non-trivial models (having non-zero RA) submitted to the CIFAR-10 and ImageNet-1K benchmarks and evaluated with  $\mathcal{L}_2$ -norm perturbation with a fixed perturbation level of 0.5 using Auto-Attack. We list the model names in Table 1 and provide their descriptions in Appendix G.

**GANs and DMs.** We used off-the-shelf GAN models provided by StudioGAN [41], a library containing released GAN models. StudioGAN also reports the Inception Score (IS) to rank the model quality. We use the GAN model with the highest IS value as our default GAN for GREAT Score, which are StyleGAN2 [30]/ BigGAN [6] for CIFAR-10/ImageNet with IS = 10.477/99.705, respectively. For the ablation study of using different generative models in GREAT Score (Section 4.4), we also use the following GAN/DM models: LSGAN [39], GGAN [35], SAGAN [65], SNGAN [42], DDPM [27] and StyleGAN2 [30].

**GREAT Score implementation.** The implementation follows Algorithm 1 in Appendix ?? with a sigmoid/softmax function on the logits of the CIFAR-10/ImageNet classifier to ensure the model output of each dimension is within  $[0, 1]$ , as implied by Theorem 1. As ImageNet-1K has 1000 classes, applying sigmoid will make the robustness score function in (3) degenerate. We use softmax instead. 500 samples drawn from a generative model were used for computing GREAT Score.

**Comparative methods.** We compare the effectiveness of GREAT Score in two objectives: robustness ranking (global robustness) and per-sample perturbation. For the former, we compare the RA reported in RobustBench on the test dataset (named RobustBench Accuracy) as well as the RA of Auto-Attack on the generated data samples (named AutoAttack Accuracy). For the latter, we report the RA of Auto-Attack in  $\mathcal{L}_2$ -norm with a fixed perturbation level of 0.5.

**Evaluation metrics.** For robustness ranking, we report Spearman’s rank correlation coefficient between two sets of model rankings (e.g., GREAT Score v.s. RobustBench Accuracy). A value closer to 1 means higher consistency. Robust accuracy refers to the fraction of correctly classified samples against adversarial perturbations.

**Calibration Method.** We run  $\mathcal{L}_2$ -norm CW attack [7] (with learning rate 0.005 and 200 iterations) on each generated data sample to find the minimal adversarial perturbation. Then, we use grid search in the range  $[0, 2]$  with an interval of 0.00001 to find temperature value maximizing the Spearman’s rank correlation coefficient between GREAT Score and CW attack distortion.

**Compute Resources.** All our experiments were run on a GTX 2080 Ti GPU with 12GB RAM.

Table 1: Comparison of (Calibrated) GREAT Score v.s. minimal distortion found by CW attack [7] on CIFAR-10. The results are averaged over 500 samples from StyleGAN2.

| Model Name                  | RobustBench Accuracy(%) | AutoAttack Accuracy(%) | GREAT Score | Calibrated GREAT Score | CW Distortion |
|-----------------------------|-------------------------|------------------------|-------------|------------------------|---------------|
| Rebuffi_extra [46]          | 82.32                   | 87.20                  | 0.507       | 1.216                  | 1.859         |
| Gowal_extra [21]            | 80.53                   | 85.60                  | 0.534       | 1.213                  | 1.324         |
| Rebuffi_70_ddpm [46]        | 80.42                   | 90.60                  | 0.451       | 1.208                  | 1.943         |
| Rebuffi_28_ddpm [46]        | 78.80                   | 90.00                  | 0.424       | 1.214                  | 1.796         |
| Augustin_WRN_extra [3]      | 78.79                   | 86.20                  | 0.525       | 1.206                  | 1.340         |
| Schwag [54]                 | 77.24                   | 89.20                  | 0.227       | 1.143                  | 1.392         |
| Augustin_WRN [3]            | 76.25                   | 86.40                  | 0.583       | 1.206                  | 1.332         |
| Rade [45]                   | 76.15                   | 86.60                  | 0.413       | 1.200                  | 1.486         |
| Rebuffi_R18[46]             | 75.86                   | 87.60                  | 0.369       | 1.210                  | 1.413         |
| Gowal [21]                  | 74.50                   | 86.40                  | 0.124       | 1.116                  | 1.253         |
| Schwag_R18 [54]             | 74.41                   | 88.60                  | 0.236       | 1.135                  | 1.343         |
| Wu2020Adversarial [62]      | 73.66                   | 84.60                  | 0.128       | 1.110                  | 1.369         |
| Augustin2020Adversarial [3] | 72.91                   | 85.20                  | 0.569       | 1.199                  | 1.285         |
| Engstrom2019Robustness [15] | 69.24                   | 82.20                  | 0.160       | 1.020                  | 1.084         |
| Rice2020Overfitting [47]    | 67.68                   | 81.80                  | 0.152       | 1.040                  | 1.097         |
| Rony2019Decoupling [50]     | 66.44                   | 79.20                  | 0.275       | 1.101                  | 1.165         |
| Ding2020MMA [14]            | 66.09                   | 77.60                  | 0.112       | 0.909                  | 1.095         |

Table 2: Spearman’s rank correlation coefficient on CIFAR-10 using GREAT Score, RobustBench (with test set), and Auto-Attack (with generated samples).

|                                         | Uncalibrated | Calibrated |
|-----------------------------------------|--------------|------------|
| GREAT Score vs. RobustBench Correlation | 0.6618       | 0.8971     |
| GREAT Score vs. AutoAttack Correlation  | 0.3690       | 0.6941     |
| RobustBench vs. AutoAttack Correlation  | 0.7296       | 0.7296     |

## 4.2 Local and Global Robustness Analysis

Recall from Theorem 1 that the local robustness score proposed in (3) gives a certified perturbation level for generated samples from a generative model. To verify this claim, we randomly select 20 generated images on CIFAR-10 and compare their local certified perturbation level to the perturbation found by the CW attack [7] using the Rebuffi\_extra model [46]. Figure 1 shows the perturbation level of local GREAT Score in (3) and that of the corresponding CW attack per sample. We can see that the local GREAT Score is a lower bound of CW attack, as the CW attack finds a successful adversarial perturbation that is no smaller than the minimal perturbation  $\Delta_{\min}$  (i.e., an over-estimation). The true  $\Delta_{\min}$  value lies between these lower and upper bounds.

In Figure 2, we compare the cumulative robust accuracy (RA) of GREAT Score and Auto-Attack over 500 samples by sweeping the  $\mathcal{L}_2$  perturbation level from 0 to 1 with a 0.05 increment for Auto-Attack. The cumulative RA of GREAT Score at a perturbation level  $r$  represents the fraction of samples with local GREAT Scores greater than  $r$ , providing an attack-proof guarantee that no attacks can achieve a lower RA at the same perturbation level. For Auto-Attack, the RA at each perturbation level is calculated as the fraction of correctly classified samples under that specific perturbation. The blue curve in the figure represents the RA from empirical Auto-Attack, while the orange curve shows the RA derived from GREAT Score, offering a certified robustness guarantee. We observe that the trend of attack-independent certified robustness (GREAT Score) closely mirrors that of empirical attacks (Auto-Attack), suggesting that GREAT Score effectively reflects empirical robustness. It is important to note that the gap between our certified curve and the empirical curve of AutoAttack does not necessarily indicate inferiority of GREAT Score. Instead, this discrepancy could point to the existence of undiscovered adversarial examples at higher perturbation radii. This gap illustrates the fundamental difference between certified and empirical robustness measures, highlighting the potential for GREAT Score to provide a more conservative, yet guaranteed, estimate of model robustness.

Table 1 compares the global robustness statistics of the 17 grouped CIFAR-10 models on RobustBench for uncalibrated and calibrated versions respectively, in terms of the GREAT Score and the average distortion of CW attack, which again verifies GREAT Score is a certified lower bound on the true global robustness (see its definition in Section 3.1), while any attack with 100% attack success rate only gives an upper bound on the true global robustness. We also observe that calibration can indeed enlarge the GREAT Score and tighten its gap to the distortion of CW attack.

## 4.3 Model Ranking on CIFAR-10 and ImageNet

Following the experiment setup in Section 4.1, we compare the model ranking on CIFAR-10 using GREAT Score (evaluated with generated samples), RobustBench (evaluated with Auto-Attack on the test set), and Auto-Attack (evaluated with Auto-Attack on generated samples). Table 2 presents their mutual rank correlation (higher value means more aligned ranking) with calibrated and uncalibrated versions. We note that there is an innate discrepancy between Spearman’s rank correlation coefficient (way below 1) of RobustBench v.s. Auto-Attack, which means Auto-Attack will give inconsistent model rankings when evaluated on different data samples. In addition, GREAT Score measures *classification margin*, while AutoAttack measures *accuracy* under a fixed perturbation budget  $\epsilon$ .

AutoAttack’s ranking will change if we use different  $\epsilon$  values. E.g., comparing the ranking of  $\epsilon = 0.3$  and  $\epsilon = 0.7$  on 10000 CIFAR-10 test images for AutoAttack, the Spearman’s correlation is only 0.9485. Therefore, we argue that GREAT Score and AutoAttack are *complementary* evaluation metrics and they don’t need to match perfectly. Despite their discrepancy, before calibration, the correlation between GREAT Score and RobustBench yields a similar value. With calibration, there is a significant improvement in rank correlation between GREAT Score to RobustBench and Auto-Attack, respectively.

Table 3 presents the global robustness statistics of these three methods on ImageNet. We observe almost perfect ranking alignment between GREAT Score and RobustBench, with their Spearman’s rank correlation coefficient being 0.8, which is higher than that of Auto-Attack and RobustBench (0.6). These results suggest that GREAT Score is a useful metric for *margin-based* robustness evaluation.

#### 4.4 Ablation Study and Run-time Analysis

**Ablation study on GANs and DMs.** Evaluating on CIFAR-10, Figure 3 compares the inception score (IS) and the Spearman’s rank correlation coefficient between GREAT Score and RobustBench on five GANs and DDPM. One can observe that models with higher IS attain better ranking consistency.

**Limitations and Further Analysis for generation models.** While our experiments demonstrate the effectiveness of GREAT Score, it’s important to acknowledge certain limitations and provide further analysis. The performance of GREAT Score relies on the generative model’s ability to produce valid samples belonging to the conditioned class. Recent studies [34, 53] have shown GANs’ convergence to true data distributions under specific conditions, and our experiments further demonstrate high-quality instances produced by the generative models, as evidenced by the inception score and the strong Spearman’s rank correlation between GREAT Score and RobustBench. We recognize that in some cases, class ambiguity may exist. However, given our focus on evaluating classifier robustness, we typically deal with well-defined and distinctive labels, considering the issue of label ambiguity is beyond the scope of our method. Furthermore, the assumption that the generative model provides a good approximation of the true data-generating distribution is crucial. Recent work [34, 53] has also demonstrated the convergence rate of approaching the true data distribution for a family of GANs under certain conditions. These considerations highlight areas for potential future work and underscore the importance of careful generative model selection when applying GREAT Score.

**Run-time analysis.** Figure 4 compares the run-time efficiency of GREAT Score over Auto-Attack on the same 500 generated CIFAR-10 images. We show the ratio of their average per-sample run-time (wall clock time of GREAT Score/Auto-Attack is reported in Appendix I) and observe around 800-2000 times improvement, validating the computational efficiency of GREAT Score. Furthermore, our framework demonstrates excellent scalability with increasing dataset sizes and model complexity, as detailed in Appendix N, showing linear scaling behavior that makes it suitable for large-scale applications.

**Sample Complexity and GREAT Score.** In Appendix J, we report the mean and variance of GREAT Score with a varying number of generated data samples. The results show that the statistics of GREAT Score are quite stable even with a small number of data samples (i.e.,  $\geq 500$ ).

Table 3: Robustness evaluation on ImageNet using GREAT Score, RobustBench (with test set), and Auto Attack (with generated samples). The Spearman’s rank correlation coefficient for GREAT Score v.s. RobustBench and Auto-Attack v.s. RobustBench is 0.9 and 0.872, respectively.

| Model Name   | RobustBench Accuracy (%) | AutoAttack Accuracy (%) | GREAT Score |
|--------------|--------------------------|-------------------------|-------------|
| Trans1 [52]  | 38.14                    | 30.4                    | 0.504       |
| Trans2 [52]  | 34.96                    | 25.8                    | 0.443       |
| LIBRARY [15] | 29.22                    | 30.6                    | 0.449       |
| Fast [61]    | 26.24                    | 19.2                    | 0.273       |
| Trans3 [52]  | 25.32                    | 19.6                    | 0.275       |




Figure 3: Comparison of Spearman’s rank correlation coefficient between GREAT Score and RobustBench using 500 generated CIFAR-10 images. Figure 4: Run-time improvement (GREAT Score over Auto-Attack) on 500 generated CIFAR-10 images. GREAT Score with different GANs.

Table 4: Group-wise and overall robustness evaluation for online gender classification APIs over 500 generated samples (per group).

| Online API Name | Old   | Young | With Eyeglasses | Without Eyeglasses | Total |
|-----------------|-------|-------|-----------------|--------------------|-------|
| BetaFace        | 0.950 | 0.662 | 0.547           | 0.973              | 0.783 |
| Inferdo         | 0.707 | 0.487 | 0.458           | 0.669              | 0.580 |
| ARSA-Technology | 1.033 | 0.958 | 0.739           | 1.082              | 0.953 |
| DEEPFACE        | 0.979 | 0.774 | 0.763           | 0.969              | 0.872 |
| Baidu           | 1.097 | 1.029 | 0.931           | 1.134              | 1.048 |
| Luxand          | 1.091 | 0.912 | 0.673           | 1.010              | 0.944 |

Table 5: GREAT Score v.s. robust accuracy under square attack [1].

| DEEPFACE      | Old    | Young  | With Eyeglasses | Without Eyeglasses |
|---------------|--------|--------|-----------------|--------------------|
| Square Attack | 84.40% | 72.60% | 65.80%          | 89.00%             |
| GREAT Score   | 0.979  | 0.774  | 0.763           | 0.969              |

#### 4.5 Evaluation on Online Facial Recognition APIs

To demonstrate GREAT Score enables robustness evaluation of black-box models that only provide model inference outcomes based on data inputs, we use synthetically generated face images with hidden attributes to evaluate six online face recognition APIs for gender classification. It is worth noting that GREAT Score is suited for privacy-sensitive assessment because it only uses synthetic face images for evaluation and does not require using real face images.

We use an off-the-shelf face image generator InterFaceGAN [56] trained on CelebA-HQ dataset [29], which can generate controllable high-quality face images with the choice of attributions such as eyeglasses, age, and expression. We generate four different groups (attributes) of face images for evaluation: Old, Young, With Eyeglasses, and Without Eyeglasses. For annotating the ground truth gender labels of the generated images, we use the gender predictions from the FAN classifier [25]. In total, 500 gender-labeled face images are generated for each group. Appendix L shows some examples of the generated images for each group.

We evaluate the GREAT Score on six online APIs for gender classification: BetaFace [5], Inferdo [28], Arsa-Technology [2], DeepFace [55], Baidu [4] and Luxand [37]. These APIs are “black-box” models to end users or an external model auditor because the model details are not revealed and only the model inference results returned by APIs (prediction probabilities on Male/Female) are provided.

Finally, we upload these images to the aforementioned online APIs and calculate the GREAT Score using the returned prediction results. Table 4 displays the group-level and overall GREAT Score results. Our evaluation reveals interesting observations. For instance, APIs such as BetaFace, Inferdo, and DEEPFACE exhibit a large discrepancy for Old v.s. Young, while other APIs have comparable scores. For all APIs, the score of With Eyeglasses is consistently and significantly lower than that of Without Eyeglasses, which suggests that eyeglasses could be a common spurious feature that affects the group-level robustness in gender classification. The analysis demonstrates how GREAT Score can be used to study the group-level robustness of an access-limited model in a privacy-enhanced manner.

To verify our evaluation, in Table 5 we compare GREAT Score to the black-box square attack [1] with  $\epsilon = 2$  and # queries = 100 on DEEPFACE. For both Age and Eyeglasses groups (Old v.s. Young and W/O eyeglasses), we see consistently that a higher GREAT Score (second row) indicates better robust accuracy (% , first row) against square attack.

## 5 Conclusion

In this paper, we presented GREAT Score, a novel and computation-efficient attack-independent metric for global robustness evaluation against adversarial perturbations. GREAT Score uses an off-the-shelf generative model such as GANs for evaluation and enjoys theoretical guarantees on its estimation of the true global robustness. Its computation is lightweight and scalable because it only requires accessing the model predictions on the generated data samples. Our extensive experimental results on CIFAR-10 and ImageNet also verified high consistency between GREAT Score and the attack-based model ranking on RobustBench, demonstrating that GREAT Score can be used as an efficient measure complementary to existing robustness benchmarks. We also demonstrated the novel use of GREAT Score for the robustness evaluation of online facial recognition APIs.

**Limitations.** One limitation could be that our framework of global adversarial robustness evaluation using generative models is centered on  $\mathcal{L}_2$ -norm based perturbations. This limitation could be addressed if the Stein’s Lemma can be extended for other  $\mathcal{L}_p$  norms.

## Acknowledgments and Disclosure of Funding

This work was supported by the JC STEM Lab of Intelligent Design Automation funded by The Hong Kong Jockey Club Charities Trust for Zaitang Li and Tsung-Yi Ho. Also, this material is based upon work supported by the Chief Digital and Artificial Intelligence Office under Contract No. W519TC-23-9-2037 for Pin-Yu Chen.

## References

- [1] 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. 10
- [2] Arsa-Technology. Arsa API. <https://rapidapi.com/arsa-technology-arsa-technology-default/api/face-recognition18>. 10
- [3] Maximilian Augustin, Alexander Meinke, and Matthias Hein. Adversarial robustness on in-and out-distribution improves explainability. In *European Conference on Computer Vision*, pages 228–245. Springer, 2020. 8, 20, 21, 24
- [4] Baidu. Baidu API. <https://console.bce.baidu.com/>. 10
- [5] BetaFace. BetaFace API. <https://rapidapi.com/betaface/api/face-recognition>. 10
- [6] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Blxsqj09Fm>. 7
- [7] 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. 3, 6, 7, 8
- [8] 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 AAAI Conference on Artificial Intelligence*, pages 10–17, 2018. 3
- [9] Jeremy M Cohen, Elan Rosenfeld, and J Zico Kolter. Certified adversarial robustness via randomized smoothing. *International Conference on Machine Learning*, 2019. 2
- [10] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *International conference on machine learning*, pages 2206–2216. PMLR, 2020. 2, 3, 7
- [11] Francesco Croce, Maksym Andriushchenko, Vikash Sehwal, Edoardo Debenedetti, Nicolas Flammarion, Mung Chiang, Prateek Mittal, and Matthias Hein. Robustbench: a standardized adversarial robustness benchmark. *arXiv preprint arXiv:2010.09670*, 2020. 2, 3, 6
- [12] Francesco Croce, Maksym Andriushchenko, Vikash Sehwal, Edoardo Debenedetti, Nicolas Flammarion, Mung Chiang, Prateek Mittal, and Matthias Hein. Robustbench: a standardized adversarial robustness benchmark. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. URL <https://openreview.net/forum?id=SSKZPJct7B>. 1
- [13] 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. 6
- [14] Gavin Weiguang Ding, Yash Sharma, Kry Yik Chau Lui, and Ruitong Huang. Mma training: Direct input space margin maximization through adversarial training. *arXiv preprint arXiv:1812.02637*, 2018. 8, 20, 21, 24

- [15] Logan Engstrom, Andrew Ilyas, Hadi Salman, Shibani Santurkar, and Dimitris Tsipras. Robustness (python library), 2019. URL <https://github.com/MadryLab/robustness>. 8, 9, 21, 24
- [16] Mahyar Fazlyab, Alexander Robey, Hamed Hassani, Manfred Morari, and George Pappas. Efficient and accurate estimation of lipschitz constants for deep neural networks. *Advances in Neural Information Processing Systems*, 32, 2019. 15
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 2
- [18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 2, 3
- [19] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 2
- [20] Sven Gowal, Krishnamurthy Dj Dvijotham, Robert Stanforth, Rudy Bunel, Chongli Qin, Jonathan Uesato, Relja Arandjelovic, Timothy Mann, and Pushmeet Kohli. Scalable verified training for provably robust image classification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4842–4851, 2019. 2
- [21] Sven Gowal, Chongli Qin, Jonathan Uesato, Timothy Mann, and Pushmeet Kohli. Uncovering the limits of adversarial training against norm-bounded adversarial examples. *arXiv preprint arXiv:2010.03593*, 2020. 8, 20, 21, 24
- [22] Ping Guo, Cheng Gong, Xi Lin, Zhiyuan Yang, and Qingfu Zhang. Exploring the adversarial frontier: Quantifying robustness via adversarial hypervolume. *arXiv preprint arXiv:2403.05100*, 2024. 3
- [23] 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. 20
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. 20
- [25] Keke He, Yanwei Fu, Wuhao Zhang, Chengjie Wang, Yu-Gang Jiang, Feiyue Huang, and Xiangyang Xue. Harnessing synthesized abstraction images to improve facial attribute recognition. 10
- [26] Matthias Hein and Maksym Andriushchenko. Formal guarantees on the robustness of a classifier against adversarial manipulation. *arXiv preprint arXiv:1705.08475*, 2017. 3
- [27] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 2, 3, 7
- [28] Inferdo. Inferdo API. <https://rapidapi.com/inferdo/api/face-detection6>. 10
- [29] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018. 6, 10
- [30] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. 7
- [31] Guy Katz, Clark Barrett, David L Dill, Kyle Julian, and Mykel J Kochenderfer. Reluplex: An efficient smt solver for verifying deep neural networks. In *International Conference on Computer Aided Verification*, pages 97–117. Springer, 2017. 2

- [32] Alex Krizhevsky et al. Learning multiple layers of features from tiny images. 2009. 6
- [33] Klas Leino, Zifan Wang, and Matt Fredrikson. Globally-robust neural networks. In *International Conference on Machine Learning*, pages 6212–6222. PMLR, 2021. 15
- [34] Tengyuan Liang. How well generative adversarial networks learn distributions. *The Journal of Machine Learning Research*, 22(1):10366–10406, 2021. 9, 15
- [35] Jae Hyun Lim and Jong Chul Ye. Geometric gan. *arXiv preprint arXiv:1705.02894*, 2017. 7
- [36] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. 6
- [37] Luxand. Luxand API. <https://dashboard.luxand.cloud/>. 10
- [38] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *International Conference on Learning Representations*, 2018. 3
- [39] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2794–2802, 2017. 7
- [40] Andreas Maurer and Massimiliano Pontil. Some hoeffding-and bernstein-type concentration inequalities. *arXiv preprint arXiv:2102.06304*, 2021. 5, 17
- [41] Minsu Cho Minguk Kang, Woohyeon Shim and Jaesik Park. Rebooting ACGAN: Auxiliary Classifier GANs with Stable Training. 2021. 7
- [42] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018. 7
- [43] Raphael Olivier and Bhiksha Raj. How many perturbations break this model? evaluating robustness beyond adversarial accuracy. In *International Conference on Machine Learning*, pages 26583–26598. PMLR, 2023. 3
- [44] Remigijus Paulavičius and Julius Žilinskas. Analysis of different norms and corresponding lipschitz constants for global optimization. *Technological and Economic Development of Economy*, 12(4):301–306, 2006. 16
- [45] Rahul Rade and Seyed-Mohsen Moosavi-Dezfooli. Helper-based adversarial training: Reducing excessive margin to achieve a better accuracy vs. robustness trade-off. In *ICML 2021 Workshop on Adversarial Machine Learning*, 2021. 8, 20, 21, 24
- [46] Sylvestre-Alvise Rebuffi, Sven Gowal, Dan A Calian, Florian Stimberg, Olivia Wiles, and Timothy Mann. Fixing data augmentation to improve adversarial robustness. *arXiv preprint arXiv:2103.01946*, 2021. 7, 8, 20, 21, 24
- [47] Leslie Rice, Eric Wong, and Zico Kolter. Overfitting in adversarially robust deep learning. In *International Conference on Machine Learning*, pages 8093–8104. PMLR, 2020. 8, 20, 21, 24
- [48] Alexander Robey, Luiz Chamon, George J Pappas, and Hamed Hassani. Probabilistically robust learning: Balancing average and worst-case performance. In *International Conference on Machine Learning*, pages 18667–18686. PMLR, 2022. 3
- [49] RobustBench. Robustbench. <https://robustbench.github.io/>. 7
- [50] Jérôme Rony, Luiz G Hafemann, Luiz S Oliveira, Ismail Ben Ayed, Robert Sabourin, and Eric Granger. Decoupling direction and norm for efficient gradient-based l2 adversarial attacks and defenses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4322–4330, 2019. 8, 20, 21, 24

- [51] Wenjie Ruan, Min Wu, Youcheng Sun, Xiaowei Huang, Daniel Kroening, and Marta Kwiatkowska. Global robustness evaluation of deep neural networks with provable guarantees for the  $l_0$  norm. *arXiv preprint arXiv:1804.05805*, 2018. 15
- [52] Hadi Salman, Andrew Ilyas, Logan Engstrom, Ashish Kapoor, and Aleksander Madry. Do adversarially robust imagenet models transfer better? *Advances in Neural Information Processing Systems*, 33:3533–3545, 2020. 9, 20
- [53] Nicolas Schreuder, Victor-Emmanuel Brunel, and Arnak Dalalyan. Statistical guarantees for generative models without domination. In *Algorithmic Learning Theory*, pages 1051–1071. PMLR, 2021. 9, 15
- [54] Vikash Sehwal, Saeed Mahloujifar, Tinashe Handina, Sihui Dai, Chong Xiang, Mung Chiang, and Prateek Mittal. Robust learning meets generative models: Can proxy distributions improve adversarial robustness? *arXiv preprint arXiv:2104.09425*, 2021. 8, 20, 21, 24
- [55] Sefik Ilkin Serengil and Alper Ozpinar. Hyperextended lightface: A facial attribute analysis framework. In *2021 International Conference on Engineering and Emerging Technologies (ICEET)*, pages 1–4. IEEE, 2021. doi: 10.1109/ICEET53442.2021.9659697. URL <https://doi.org/10.1109/ICEET53442.2021.9659697>. 10
- [56] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *CVPR*, 2020. 10
- [57] Charles M Stein. Estimation of the mean of a multivariate normal distribution. *The annals of Statistics*, pages 1135–1151, 1981. 16
- [58] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *International Conference on Learning Representations*, 2014. 3
- [59] Tsui-Wei Weng, Huan Zhang, Pin-Yu Chen, Jinfeng Yi, Dong Su, Yupeng Gao, Cho-Jui Hsieh, and Luca Daniel. Evaluating the robustness of neural networks: An extreme value theory approach. *arXiv preprint arXiv:1801.10578*, 2018. 2, 4, 15, 16
- [60] Eric Wong and Zico Kolter. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *International Conference on Machine Learning*, pages 5286–5295, 2018. 2
- [61] Eric Wong, Leslie Rice, and J Zico Kolter. Fast is better than free: Revisiting adversarial training. *arXiv preprint arXiv:2001.03994*, 2020. 9, 20
- [62] Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. *Advances in Neural Information Processing Systems*, 33:2958–2969, 2020. 8, 20, 21, 24
- [63] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6023–6032, 2019. 20
- [64] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016. 20
- [65] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *International conference on machine learning*, pages 7354–7363. PMLR, 2019. 7
- [66] Huan Zhang, Tsui-Wei Weng, Pin-Yu Chen, Cho-Jui Hsieh, and Luca Daniel. Efficient neural network robustness certification with general activation functions. In *Advances in Neural Information Processing Systems*, pages 4944–4953, 2018. 2

## A Notations

Table 6: Main notations used in this paper

| Notation                                    | Description                                             |
|---------------------------------------------|---------------------------------------------------------|
| $d$                                         | dimensionality of the input vector                      |
| $K$                                         | number of output classes                                |
| $f : \mathbb{R}^d \rightarrow \mathbb{R}^K$ | neural network classifier                               |
| $x \in \mathbb{R}^d$                        | data sample                                             |
| $y$                                         | groundtruth class label                                 |
| $\delta \in \mathbb{R}^d$                   | input perturbation                                      |
| $\ \delta\ _p$                              | $\mathcal{L}_p$ norm of perturbation, $p \geq 1$        |
| $\Delta_{\min}$                             | minimum adversarial perturbation                        |
| $G$                                         | (conditional) generative model                          |
| $z \sim \mathcal{N}(0, I)$                  | latent vector sampled from Gaussian distribution        |
| $g$                                         | robustness score function defined in (3)                |
| $\Omega(f)/\widehat{\Omega}(f)$             | true/estimated global robustness defined in Section 3.1 |

## B More Motivations on using Generative Models for Robustness Evaluation

We emphasize the necessity of generative models using the points below.

1. *Global robustness assessment requires a GM.* The major focus and novelty of our study are to evaluate the global robustness with respect to the underlying true data distribution, and we propose to use a GM as a proxy. We argue that such a proxy is necessary to evaluate global robustness unless the true data distribution is known.
2. *GAN can provably match data distribution.* Recent works such as [53] and [34] have proved the convergence rate of approaching the true data distribution for a family of GANs under certain conditions. This will benefit global robustness evaluation (see Figure 3 for ablations on GAN variants).
3. *Privacy-sensitive remote model auditing.* As shown in Sec 4.5, synthetic data from generative models can facilitate the robustness evaluation of privacy-sensitive models.

### B.1 Related Works for Global Robustness Evaluation for Deep Neural Networks.

There are some works studying “global robustness”, while their contexts and scopes are different than ours. In [51], the global robustness is defined as the expectation of the maximal certified radius of  $\mathcal{L}_0$ -norm over a test dataset. Ours is not limited to a test set, and we take the novel perspective of the entire data distribution and use a generative model to define and evaluate global robustness. The other line of works considers deriving and computing the global Lipschitz constant of the classifier as a global certificate of robustness guarantee, as it quantifies the maximal change of the classifier with respect to the entire input space [33]. The computation can be converted as a semidefinite program (SDP) [16]. However, the computation of SDP is expensive and hard to scale to larger neural networks. Our method does not require computing the global Lipschitz constant, and our computation is as simple as data forward pass for model inference.

## C Proof of Theorem 1

In this section, we will give detailed proof for the certified global robustness estimate in Theorem 1. The proof contains three parts: (i) derive the local robustness certificate; (ii) derive the closed-form global Lipschitz constant; and (iii) prove the proposed global robustness estimate is a lower bound on the true global robustness.

We provide a proof sketch below:

1. We use the local robustness certificate developed in [59], which shows an expression of a certified (attack-proof)  $\mathcal{L}_p$ -norm bounded perturbation for any  $p \geq 1$ . The certificate is a function of the gap between the best and second-best class predictions, as well as a local Lipschitz constant associated with the gap function.

2. We use Stein's Lemma [57] which states that the mean of a measurable function integrated over a zero-mean isotropic Gaussian distribution has a closed-form global Lipschitz constant in the  $\mathcal{L}_2$ -norm. This result helps avoid the computation of the local Lipschitz constant in Step 1 for global robustness evaluation using generative models.
3. We use the results from Steps 1 and 2 to prove that the proposed global robustness estimate  $\widehat{\Omega}(f)$  is a lower bound on the true global robustness  $\Omega(f)$  with respect to  $G$ .

### C.1 Local robustness certificate

In this part, we use the local robustness certificate in [59] to show an expression for local robustness certificate consisting of a gap function in model output and a local Lipschitz constant. The first lemma formally defines Lipschitz continuity and the second lemma introduces the local robustness certificate in [59].

**Lemma 1** (Lipschitz continuity in Gradient Form ([44])). *Let  $S \subset \mathbf{R}^d$  be a convex bound closed set and let  $f : S \rightarrow \mathbf{R}$  be a continuously differentiable function on an open set containing  $S$ . Then  $f$  is a Lipschitz continuous function if the following inequality holds for any  $x, y \in S$ :*

$$|f(x) - f(y)| \leq L_q \|x - y\|_p \quad (5)$$

where  $L_q = \max_{x \in S} \|\nabla f(x)\|_q$  : is the corresponding Lipschitz constant, and  $\nabla f(x) = (\frac{\partial f}{\partial x_1}, \dots, \frac{\partial f}{\partial x_d})^\top$  is the gradient of the function  $f(x)$ , and  $1/q + 1/p = 1$ ,  $p \geq 1$ ,  $q \leq \infty$ .

We say  $f$  is  $L_q$ -continuous in  $\mathcal{L}_p$  norm if (5) is satisfied.

**Lemma 2** (Formal guarantee on lower bound for untargeted attack of Theorem 3.2 in [59]). *Let  $x_0 \in \mathbf{R}^d$  and  $f : \mathbf{R}^d \rightarrow \mathbf{R}^K$  be a multi-class classifier, and  $f_i$  be the  $i$ -th output of  $f$ . For untargeted attack, to ensure that the adversarial examples can not be found for each class, for all  $\delta \in \mathbf{R}^d$ , the lower bound of minimum distortion can be expressed by:*

$$\|\delta\|_p \leq \min_{i \neq m} \frac{f_m(x_0) - f_i(x_0)}{L_q^i} \quad (6)$$

where  $m = \arg \max_{i \in \{1, \dots, K\}} f_i(x_0)$ ,  $1/q + 1/p = 1$ ,  $p \geq 1$ ,  $q \leq \infty$ , and  $L_q^i$  is the Lipschitz constant for the function  $f_m(x) - f_i(x)$  in  $L_q$  norm.

### C.2 Proof of closed-form global Lipschitz constant in the $L_2$ -norm over Gaussian distribution

In this part, we present two lemmas towards developing the global Lipschitz constant of a function smoothed by a Gaussian distribution.

**Lemma 3** (Stein's lemma [57]). *Given a soft classifier  $F : \mathbf{R}^d \rightarrow \mathbf{P}$ , where  $\mathbf{P}$  is the space of probability distributions over classes. The associated smooth classifier with parameter  $\sigma \geq 0$  is defined as:*

$$\bar{F} := (F * \mathcal{N}(0, \sigma^2 I))(x) = \mathbb{E}_{\delta \sim \mathcal{N}(0, \sigma^2 I)} [F(x + \delta)] \quad (7)$$

Then,  $\bar{F}$  is differentiable, and moreover,

$$\nabla \bar{F} = \frac{1}{\sigma^2} \mathbb{E}_{\delta \sim \mathcal{N}(0, \sigma^2 I)} [\delta \cdot F(x + \delta)] \quad (8)$$

In a lecture note<sup>1</sup>, Li used Stein's Lemma [57] to prove the following lemma:

**Lemma 4** (Proof of global Lipschitz constant). *Let  $\sigma \geq 0$ , let  $h : \mathbf{R}^d \rightarrow [0, 1]$  be measurable, and let  $H = h * \mathcal{N}(0, \sigma^2 I)$ . Then  $H$  is  $\sqrt{\frac{2}{\pi \sigma^2}}$ -continuous in  $L_2$  norm*

<sup>1</sup><https://jerryzli.github.io/robust-ml-fall19/lec14.pdf>

### C.3 Proof of the proposed global robustness estimate $\widehat{\Omega}(f)$ is a lower bound on the true global robustness $\Omega(f)$ with respect to $G$

Recall that we assume a generative model  $G(\cdot)$  generates a sample  $G(z)$  with  $z \sim \mathcal{N}(0, I)$ . Following the form of Lemma 2 (but ignoring the local Lipschitz constant), let

$$g'(G(z)) = \max\{f_c(G(z)) - \max_{k \in \{1, \dots, K\}, k \neq c} f_k(G(z)), 0\} \quad (9)$$

denote the gap in the model likelihood of the correct class  $c$  and the most likely class other than  $c$  of a given classifier  $f$ , where the gap is defined to be 0 if the model makes an incorrect top-1 class prediction on  $G(z)$ . Then, using Lemma 4 with  $g'$ , we define

$$\mathbb{E}_{z \sim \mathcal{N}(0, I)}[g'(G(z))] = (g' \circ G) * \mathcal{N}(0, I) \quad (10)$$

[ZAITANG: Modified the equations] and thus  $\mathbb{E}_{z \sim \mathcal{N}(0, I)}[g'(G(z))]$  has a Lipschitz constant  $\sqrt{\frac{2}{\pi}}$  in  $\mathcal{L}_2$  norm. This implies that for any input perturbation  $\delta$ ,

$$|\mathbb{E}_{z \sim \mathcal{N}(0, I)}[g'(G(z) + \delta)] - \mathbb{E}_{z \sim \mathcal{N}(0, I)}[g'(G(z))]| \quad (11)$$

$$\leq \sqrt{\frac{2}{\pi}} \cdot \|\delta\|_2 \quad (12)$$

and therefore

$$\mathbb{E}_{z \sim \mathcal{N}(0, I)}[g'(G(z) + \delta)] \quad (13)$$

$$\geq \mathbb{E}_{z \sim \mathcal{N}(0, I)}[g'(G(z))] - \sqrt{\frac{2}{\pi}} \cdot \|\delta\|_2 \quad (14)$$

Note that if the right-hand side of (13) is greater than zero, this will imply the classifier attains a nontrivial positive mean gap with respect to the generative model. This condition holds for any  $\delta$  satisfying  $\|\delta\|_2 < \sqrt{\frac{\pi}{2}} \cdot \mathbb{E}_{z \sim \mathcal{N}(0, I)}[g'(G(z))]$ . Note that by definition any minimum perturbation on

$G(z)$  will be no smaller than  $\sqrt{\frac{\pi}{2}} \cdot \mathbb{E}_{z \sim \mathcal{N}(0, I)}[g'(G(z))]$  as it will make  $g'(G(z)) = 0$  almost surely.

Therefore, by defining  $g = \sqrt{\frac{\pi}{2}} \cdot g'$ , we conclude that the global robustness estimate  $\widehat{\Omega}(f)$  in (2) using the proposed local robustness score  $g$  defined in (3) is a certified lower bound on the true global robustness  $\Omega(f)$  with respect to  $G$ .

## D Proof of Theorem 2

To prove Theorem 2, we first define some notations as follows, with a slight abuse of the notation  $f$  as a generic function in this part. For a vector of independent random variables  $X = (X_1, \dots, X_n)$ , define  $X' = (X'_1, \dots, X'_n)$  to be i.i.d. to  $X$ ,  $x = (x_1, \dots, x_n) \in \mathbf{X}$ , and the sub-exponential norms  $\|\cdot\|_{\psi_2}$  for any random variable  $Z$  as

$$\|Z\|_{\psi_2} = \sup_{p \geq 1} \frac{\|Z\|_p}{\sqrt{p}} \quad (15)$$

Let  $f : X^n \mapsto \mathbf{R}$ . We further define the  $k$ -th centered conditional version of  $f$  as :

$$f_k(X) = f(X) - \mathbb{E}[f(X) | X_1, \dots, X_{k-1}, X_{k+1}, \dots, X_n] \quad (16)$$

**Lemma 5** (Concentration inequality from Theorem 3.1 in [40]). *Let  $f : X^n \mapsto \mathbf{R}$  and  $X = (X_1, \dots, X_n)$  be a vector of independent random variables with values in a space  $\mathbb{X}$ . Then for any  $t > 0$  we have*

$$\Pr(f(X) - \mathbb{E}[f(X')] > t) \leq \exp \left( \frac{-t^2}{32e \left\| \sum_k \|f_k(X)\|_{\psi_2}^2 \right\|_{\infty}} \right) \quad (17)$$

Recall that we aim to derive a probabilistic guarantee on the sample mean of the local robustness score in (3) from a  $K$ -way classifier with its outputs bounded by  $[0, 1]^K$ . Following the definition of  $g$  (for simplicity, ignoring the constant  $\sqrt{\pi/2}$ ), the sample mean  $f$  can be expressed as:

$$f(X) = \frac{1}{n} \sum_{i=1}^n g(X_i) \quad (18)$$

where  $X_i \sim \mathcal{N}(0, I)$ .

Following the definition of (16),

$$f_k(X) = f(X) - \mathbb{E}[f(X)|X_1, \dots, X_{k-1}, X_{k+1}, \dots, X_n] \quad (19)$$

$$= \frac{1}{n} [g(X_k) - g(X'_k)] \leq \frac{1}{n} \quad (20)$$

This implies  $f_k(X)$  is bounded by  $\frac{1}{n}$ , i.e.,  $\|f_k(X)\|_\infty \leq \frac{1}{n}$ , and also  $\|f_k(X)\|_{\psi_2} \leq \frac{1}{n}$ .

Squaring over  $\|f_k(X)\|_{\psi_2}$  gives

$$\|f_k(X)\|_{\psi_2}^2 \leq \frac{1}{n^2} \quad (21)$$

As a result,

$$\left\| \sum_k \|f_k(X)\|_{\psi_2}^2 \right\|_\infty \leq n \cdot \frac{1}{n^2} = \frac{1}{n} \quad (22)$$

Divide both side of (22) and multiply with  $\frac{-t^2}{32e}$  gives:

$$\frac{-t^2}{32e \left\| \sum_k \|f_k(X)\|_{\psi_2}^2 \right\|_\infty} \leq \frac{-t^2 n}{32e} \quad (23)$$

Take exponential function over both side of (23) gives

$$\exp \left( \frac{-t^2}{32e \left\| \sum_k \|f_k(X)\|_{\psi_2}^2 \right\|_\infty} \right) \leq \exp \left( \frac{-t^2 n}{32e} \right) \quad (24)$$

Recall Lemma 5, since this bound holds on both sides of the central mean, we rewrite it as:

$$\text{Prob}(|f(X) - \mathbb{E}[f(X')]| > t) \leq 2 \exp \left( \frac{-t^2}{32e \left\| \sum_k \|f_k(X)\|_{\psi_2}^2 \right\|_\infty} \right) \quad (25)$$

Hence to ensure that given a statistical tolerance  $\epsilon > 0$  with  $\delta$  as the maximum outage probability, i.e.,  $\text{Prob}(|f(X) - E[f(X')]| > \epsilon) \leq \delta$ , we have

$$2 \cdot \exp \left( \frac{-\epsilon^2}{32e \left\| \sum_k \|f_k(X)\|_{\psi_2}^2 \right\|_\infty} \right) \leq 2 \exp \left( \frac{-\epsilon^2 n}{32e} \right) \quad (26)$$

$$\leq \delta \quad (27)$$

Finally, (26) implies that the sample complexity to reach the  $(\epsilon, \delta)$  condition is  $n \geq \frac{32e \cdot \log(2/\delta)}{\epsilon^2}$ .

Figure 5 shows the flow chart of Algorithm 1.

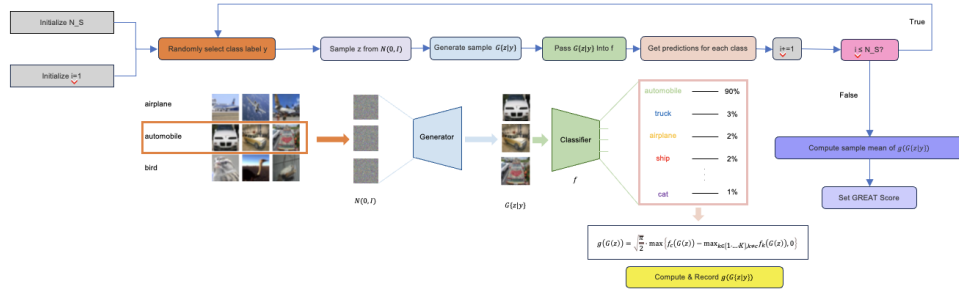


Figure 5: The Flow Chart of GREAT Score.

Table 7: Spearman’s rank correlation coefficient on CIFAR-10 using GREAT Score, RobustBench (with test set), and Auto-Attack (with generated samples) with different calibration methods.

|                                        | GREAT Score v.s.<br>RobustBench Correlation | GREAT Score v.s.<br>AutoAttack Correlation | RobustBench v.s<br>AutoAttack Correlation |
|----------------------------------------|---------------------------------------------|--------------------------------------------|-------------------------------------------|
| softmax with temperature               | -0.5024                                     | -0.5334                                    | 0.7296                                    |
| sigmoid with temperature               | 0.7083                                      | 0.3641                                     | 0.7296                                    |
| sigmoid with temperature after softmax | -0.2525                                     | -0.2722                                    | 0.7296                                    |
| softmax with temperature after sigmoid | 0.8971                                      | 0.6941                                     | 0.7296                                    |

## E Comparison between CW Attack and GREAT Score

We provide a detailed comparison of the time complexity between GREAT Score and CW Attack.

The time complexity of the GREAT Score algorithm is determined by the number of iterations (generated samples) in the loop, which is denoted as  $N_S$ . Within each iteration, the algorithm performs operations such as random selection, sampling from a Gaussian distribution, generating samples, and predicting class labels using the classifier. We assume these operations have constant time complexity  $I$  and absorb them in the big  $O$  notation. Additionally, the algorithm computes the sample mean of the recorded statistics, which involves summing and dividing the values. As there are  $N_S$  values to sum and divide, this step has a time complexity of  $O(N_S)$ . Therefore, the overall time complexity of the algorithm can be approximated as  $O(N_S \cdot I)$ .

Using our notation, consider a  $K$ -way classifier  $f$ . Let  $x$  be a data sample and  $y$  be its top-1 classification label. Denote  $\delta$  as the adversarial perturbation. The untargeted CW Attack ( $L_2$  norm) solves the following optimization objective:

$$\delta^* = \arg \min_{\delta} (\|\delta\|_2^2 + \alpha \cdot \max\{f_y(x + \delta) - \max_{k \in \{1, \dots, K\}, k \neq y} f_k(x + \delta), 0\}) \quad (28)$$

where  $f_k(\cdot)$  is the prediction of the  $k$ -th class, and  $\alpha > 0$  is a hyperparameter.

For CW attack, the optimization process iteratively finds the adversarial perturbation. The number of iterations required depends on factors such as the desired level of attack success and the convergence criteria. Each iteration involves computing gradients, updating variables, and evaluating the objective function. It also involves a hyperparameter  $\alpha$  search stage to adjust the weighted loss function.

Specifically, let  $B$  be the complexity of backpropagation,  $T_g$  be the number of iterative optimizations, and  $T_b$  be the number of binary search steps for  $\alpha$ . The dominant computation complexity of CW attack for  $N_S$  samples is in the order of  $O(N_S \cdot T_g \cdot T_b \cdot B)$ . Normally,  $T_g$  is set to 1000, and  $T_b$  is set to 9. Therefore, CW attack algorithm is much more time-consuming than GREAT Score.

## F Best Calibration Coefficient on different activation methods

Table 7 shows the best ranking coefficient we achieved on each calibration option for CIFAR-10. Among all these four calibration choices, we found that Sigmoid then Temperature Softmax achieves the best result.

## G Detailed descriptions of the Models

We provide the detail description for classifiers on RobustBench in what follows. The classifiers for CIFAR-10 are mentioned first and the last paragraph provides descriptions for ImageNet classifiers.

- *Rebuffi et al. [46]*: Rebuffi et al. [46] proposed a fixing data augmentation method such as using CutMix [63] and GANs to prevent over-fitting. There are 4 models recorded in [46]: Rebuffi\_extra uses extra data from Tiny ImageNet in training, while Rebuffi\_70\_ddpm uses synthetic data from DDPM. Rebuffi\_70\_ddpm/Rebuffi\_28\_ddpm/Rebuffi\_R18 varies in the network architecture. They use WideResNet-70-16 [64]/WideResNet-28-10 [64]/PreActResNet-18 [24].
- *Gowal et al. [21]*: Gowal et al. [21] studied various training settings such as training losses, model sizes, and model weight averaging. Gowal\_extra differs from Gowal in using extra data from Tiny ImageNet for training.
- *Augustin et al. [3]*: Augustin et al. [3] proposed RATIO, which trains with an out-Of-distribution dataset. Augustin\_WRN\_extra uses the out-of-distribution data samples for training while Augustin\_WRN does not.
- *SehWag et al. [54]*: SehWag et al. [54] found that a proxy distribution containing extra data can help to improve the robust accuracy. SehWag/SehWag\_R18 uses WideResNet-34-10 [64]/ResNet-18 [23], respectively.
- *Rade et al. [45]*: Rade [45] incorporates wrongly labeled data samples for training.
- *Wu et al. [62]*: Wu2020Adversarial [62] regularizes weight loss landscape.
- *LIBRARY*: Engstrom2019Robustness<sup>2</sup> is a package used to train and evaluate the robustness of neural network.
- *Rice et al. [47]*: Rice2020Overfitting [47] uses early stopping in reduce over-fitting during training.
- *Rony et al. [50]*: Rony2019Decoupling [50] generates gradient-based attacks for robust training.
- *Ding et al. [50]*: Ding2020MMA [14] enables adaptive selection of perturbation level during training.

For the 5 ImageNet models, Trans [52] incorporates transfer learning with adversarial training. Its model variants Trans1/Trans2/Trans3 use WideResNet-50-2 [64]/ResNet-50 [23]/ResNet-18 [23]. LIBRARY means using the package mentioned in Group of other models to train on ImageNet. Fast [61] means fast adversarial training. There is no  $\mathcal{L}_2$ -norm benchmark for ImageNet on RobustBench, so we use the  $\mathcal{L}_\infty$ -norm benchmark.

## H Approximation Error and Sample Complexity

Figure 6 presents the sample complexity as analyzed in Theorem 2 with varying approximation error ( $\epsilon$ ) and three confidence parameters ( $\delta$ ) for quantifying the difference between the sample mean and the true mean for global robustness estimation. As expected, smaller  $\delta$  or smaller  $\epsilon$  will lead to higher sample complexity.




Figure 6: The relationship between the approximation error ( $\epsilon$ ) and sample complexity in Theorem 2, with three different confidence levels:  $\delta = \{5, 15, 25\}\%$ .

<sup>2</sup><https://github.com/MadryLab/robustness>

Table 8: Group-wise time efficiency evaluation on CIFAR-10 using GREAT Score and Auto-Attack (with 500 generated samples).

| Model Name                  | GREAT Score(Per Sample)(s) | AutoAttack(Per Sample)(s) |
|-----------------------------|----------------------------|---------------------------|
| Rebuffi_extra [46]          | 0.038                      | 60.872                    |
| Gowal_extra [21]            | 0.034                      | 59.586                    |
| Rebuffi_70_ddpm [46]        | 0.034                      | 61.3362                   |
| Rebuffi_28_ddpm [46]        | 0.011                      | 10.3828                   |
| Augustin_WRN_extra [3]      | 0.013                      | 10.096                    |
| Sehwag [54]                 | 0.011                      | 10.3662                   |
| Augustin_WRN [3]            | 0.011                      | 10.1056                   |
| Rade [45]                   | 0.007                      | 4.4114                    |
| Rebuffi_R18[46]             | 0.008                      | 4.4644                    |
| Gowal [21]                  | 0.034                      | 60.746                    |
| Sehwag_R18 [54]             | 0.007                      | 3.8652                    |
| Wu2020Adversarial [62]      | 0.012                      | 10.9826                   |
| Augustin2020Adversarial [3] | 0.014                      | 6.9148                    |
| Engstrom2019Robustness [15] | 0.012                      | 6.6462                    |
| Rice2020Overfitting [47]    | 0.007                      | 3.5776                    |
| Rony2019Decoupling [50]     | 0.010                      | 8.5834                    |
| Ding2020MMA [14]            | 0.008                      | 3.6194                    |

## I Complete Run-time Results

The complete run-time results of Figure 4 are given in Table 8:

## J Sample Complexity and GREAT Score




Figure 7: The relation of GREAT Score and sample complexity using CIFAR-10 and Rebuffi\_extra model over (500-10000) range. The data points refer to the mean value for GREAT Score, and the error bars refers to the standard derivation for GREAT Score.

Figure 7 reports the mean and variance of GREAT Score with a varying number of generated data samples using CIFAR-10 and the Rebuffi\_extra model, ranging from 500 to 10000 with 500 increment. Figure 8 reports the mean and variance of GREAT Score ranging from 50 to 1000 with 50 increment. The results show that the statistics of GREAT Score are quite stable even with a small number of data samples.

## K GREAT Score Evaluation on the Original Test Samples of CIFAR-10

Besides evaluating the GREAT Score on the generated samples from GAN, we also run the evaluation process on 500 test samples of CIFAR-10. Table 9 shows the evaluated GREAT Score.




Figure 8: The relation of GREAT Score and sample complexity using CIFAR-10 and Rebuffi\_extra model over (50-1000) range. The data points refer to the mean value for GREAT Score, and the error bars refers to the standard derivation for GREAT Score.



Figure 9: Generated Images for old subgroup.

## L Generated Images from Facial GAN Models

We show the generated images from four groups in what follows.

## M Impact Statements

As this work focuses on quantifying and scoring the global robustness of neural network classifiers, we do not currently foresee any negative impact based on our work. We envision our work to be used in model auditing settings such as model cards.

## N Scalability Analysis

To evaluate the scalability of the GREAT Score framework, we conducted experiments using three ResNet variants (ResNet50, ResNet101, and ResNet152) with varying dataset sizes ranging from 500 to 2000 images. The computation times were measured in milliseconds without implementing any attack mechanisms.

Table 10 presents the detailed computational performance across different configurations.

Our experimental results demonstrate a linear increase in computation time with respect to both dataset size and model complexity. More sophisticated architectures like ResNet152 required proportionally more processing time compared to simpler ones like ResNet50. This linear scalability indicates that



Figure 10: Generated Images for young subgroup.



Figure 11: Generated Images for with-eyeglasses subgroup.



Figure 12: Generated Images for without-eyeglasses subgroup.

the GREAT Score framework efficiently handles larger datasets and more complex models, making it suitable for large-scale applications.

Table 9: GREAT Score on CIFAR-10. The results are averaged over 500 original test samples.

| Model Name                  | RobustBench<br>Accuracy(%) | AutoAttack<br>Accuracy(%) | GREAT<br>Score | Test Samples<br>GREAT Score |
|-----------------------------|----------------------------|---------------------------|----------------|-----------------------------|
| Rebuffi_extra [46]          | 82.32                      | 87.20                     | 0.507          | 0.465                       |
| Gowal_extra [21]            | 80.53                      | 85.60                     | 0.534          | 0.481                       |
| Rebuffi_70_ddpm [46]        | 80.42                      | 90.60                     | 0.451          | 0.377                       |
| Rebuffi_28_ddpm [46]        | 78.80                      | 90.00                     | 0.424          | 0.344                       |
| Augustin_WRN_extra [3]      | 78.79                      | 86.20                     | 0.525          | 0.525                       |
| Sehwag [54]                 | 77.24                      | 89.20                     | 0.227          | 0.227                       |
| Augustin_WRN [3]            | 76.25                      | 86.40                     | 0.583          | 0.489                       |
| Rade [45]                   | 76.15                      | 86.60                     | 0.413          | 0.331                       |
| Rebuffi_R18[46]             | 75.86                      | 87.60                     | 0.369          | 0.297                       |
| Gowal [21]                  | 74.50                      | 86.40                     | 0.124          | 0.109                       |
| Sehwag_R18 [54]             | 74.41                      | 88.60                     | 0.236          | 0.176                       |
| Wu2020Adversarial [62]      | 73.66                      | 84.60                     | 0.128          | 0.106                       |
| Augustin2020Adversarial [3] | 72.91                      | 85.20                     | 0.569          | 0.493                       |
| Engstrom2019Robustness [15] | 69.24                      | 82.20                     | 0.160          | 0.127                       |
| Rice2020Overfitting [47]    | 67.68                      | 81.80                     | 0.152          | 0.120                       |
| Rony2019Decoupling [50]     | 66.44                      | 79.20                     | 0.275          | 0.221                       |
| Ding2020MMA [14]            | 66.09                      | 77.60                     | 0.112          | 0.08                        |

Table 10: Computation time (ms) for different ResNet models and dataset sizes

| Dataset Size | ResNet50 | ResNet101 | ResNet152 |
|--------------|----------|-----------|-----------|
| 500          | 3274     | 6251      | 9149      |
| 1000         | 6529     | 12528     | 18339     |
| 1500         | 9785     | 18838     | 27481     |
| 2000         | 12960    | 24917     | 36588     |

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: Yes, the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope.

In the abstract and introduction, we introduce the GREAT Score, a novel and computation-efficient attack-independent metric for global robustness evaluation against adversarial perturbations. This claim is supported throughout the paper by detailing the methodology and implementation of the GREAT Score, which utilizes off-the-shelf generative models such as GANs. The paper discusses the theoretical guarantees associated with GREAT Score's estimation of true global robustness, reinforcing the validity of this claim.

The lightweight and scalable nature of the GREAT Score is emphasized, noting that it only requires model predictions on generated data samples. This is thoroughly validated through extensive experimental results on CIFAR-10 and ImageNet datasets, where we demonstrate high consistency between GREAT Score and the attack-based model rankings on RobustBench. These results substantiate the claim that GREAT Score can serve as an efficient alternative for robustness benchmarks.

Furthermore, the paper explores the novel application of GREAT Score for evaluating the robustness of online facial recognition APIs. This application is detailed in the results section, providing additional evidence of the metric's versatility and practical utility.

Therefore, the main claims in the abstract and introduction are well-supported by the comprehensive experimental results and detailed analysis presented in the paper, ensuring an accurate reflection of our contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: Yes, please see Limitations in Section 5.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.

- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Please see Section 3 for all the Theory Assumptions, we provide the proof for them in Appendix C and D .

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Yes, please check Section 4.1 and Appendix for experiment details.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.

- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

## 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [\[Yes\]](#)

Justification: Please see the attached code in the supplementary material. We also prepare a document alongside the code for instructions to reproduce the experiments.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: We have provided a experiment setup in Section 4.1. Besides, we give a detail explanation in each subsection of the experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.

- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We provide a probabilistic guarantee on our GREAT Score evaluation in Theorem 2, which can be translated to error bars.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Please check Section 4.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have exactly follow the code of Ethics of NeurIPS.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.

- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

#### 10. **Broader Impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Please see Section M.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

#### 11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: As our work do not directly release models and data, hence no high risk for misuse of it.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

#### 12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We have cite all the assets we used in the paper. We use all licensed datasets under MIT License.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

### 13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We have provided our code for evaluation robustness framework. Although our method involved using generative model to generate new data. We actually share the generation process, not the data itself. Therefore, no new data is introduced here.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No Human Subjects concerned.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: No Human Subjects concerned.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.