
Understanding the Intrinsic Robustness of Image Distributions using Conditional Generative Models

Xiao Zhang*
University of Virginia

Jinghui Chen*
University of California
Los Angeles

Quanquan Gu
University of California
Los Angeles

David Evans
University of Virginia

Abstract

Starting with Gilmer et al. (2018), several works have demonstrated the inevitability of adversarial examples based on different assumptions about the underlying input probability space. It remains unclear, however, whether these results apply to natural image distributions. In this work, we assume the underlying data distribution is captured by some conditional generative model, and prove intrinsic robustness bounds for a general class of classifiers, which solves an open problem in Fawzi et al. (2018). Building upon the state-of-the-art conditional generative models, we study the intrinsic robustness of two common image benchmarks under ℓ_2 perturbations, and show the existence of a large gap between the robustness limits implied by our theory and the adversarial robustness achieved by current state-of-the-art robust models.

1 Introduction

Deep neural networks (DNNs) have achieved remarkable performance on many visual (Sutskever et al., 2012; He et al., 2016) and speech (Hinton et al., 2012) recognition tasks, but recent studies have shown that state-of-the-art DNNs are surprisingly vulnerable to *adversarial perturbations*, small imperceptible input transformations that are designed to switch the prediction of the classifier (Szegedy et al., 2014; Goodfellow et al., 2015). This has led to a vigorous arms

race between heuristic defenses (Papernot et al., 2016; Madry et al., 2018; Chakraborty et al., 2018; Wang et al., 2019) that propose ways to defend against existing attacks and newly-devised attacks (Carlini and Wagner, 2017; Athalye et al., 2018; Tramer et al., 2020) that are able to penetrate such defenses. Reliable defenses appear to be elusive, despite progress on provable defenses, including formal verification (Katz et al., 2017; Tjeng et al., 2019) and relaxation-based certification methods (Sinha et al., 2018; Raghu et al., 2018; Wong and Kolter, 2018; Gowal et al., 2019; Wang et al., 2018). Even the strongest of these defenses leave large opportunities for adversaries to find adversarial examples, while suffering from high computation costs and scalability issues.

Witnessing the difficulties of constructing robust classifiers, a line of recent works (Gilmer et al., 2018; Fawzi et al., 2018; Mahloui et al., 2019a; Shafahi et al., 2019) aims to understand the limitations of robust learning by providing theoretical bounds on adversarial robustness for arbitrary classifiers. By imposing different assumptions on the underlying data distributions and allowable perturbations, all of these theoretical works show that no adversarially robust classifiers exist for an assumed metric probability space, as long as the perturbation strength is sublinear in the typical norm of the inputs. Although such impossibility results seem disheartening to the goal of building robust classifiers, it remains unknown to what extent real image distributions satisfy the assumptions needed to obtain these results.

In this paper, we aim to bridge the gap between the theoretical robustness analyses on well-behaved data distributions and the maximum achievable adversarial robustness, which we call *intrinsic robustness* (formally defined by Definition 3.2), for typical image distributions. More specifically, we assume the underlying data lie on a separable low-dimensional manifold, which can be captured using a conditional generative model, then systematically study the intrinsic robustness based on the conditional generating process from

*Equal Contribution

both theoretical and experimental perspectives. Our main contributions are:

- We prove a fundamental bound on intrinsic robustness (Section 4), provided that the underlying data distribution can be captured by a conditional generative model, solving an open problem in Fawzi et al. (2018).
- Building upon a trained conditional generative model that mimics the underlying data generating process, we empirically evaluate the intrinsic robustness on image distributions based on MNIST and ImageNet (Section 5.2). Our estimates of intrinsic robustness demonstrate that there is still a large gap between the limits implied by our theory and the state-of-the-art robustness achieved by robust training methods (Section 5.3).
- We theoretically characterize the fundamental relationship between the *in-distribution adversarial risk* (which restricts adversarial examples to lie on the image manifold, and is formally defined by Definition 3.3) and the intrinsic robustness (Remark 4.6), and propose an optimization method to search for in-distribution adversarial examples with respect to a given classifier. Our estimated in-distribution robustness for state-of-the-art adversarially trained classifiers, together with the derived intrinsic robustness bound, provide a better understanding on the intrinsic robustness for natural image distributions (Section 5.4).

Notation. We use lower boldfaced letters such as \mathbf{x} to denote vectors, and $[n]$ to denote the index set $\{1, 2, \dots, n\}$. For any $\mathbf{x} \in \mathcal{X}$ and $\epsilon \geq 0$, denote by $\mathcal{B}(\mathbf{x}, \epsilon, \Delta) = \{\mathbf{x}' \in \mathcal{X} : \Delta(\mathbf{x}, \mathbf{x}') \leq \epsilon\}$ the ϵ -ball around \mathbf{x} with radius ϵ in some distance metric Δ . When the metric is free of context, we simply write $\mathcal{B}(\mathbf{x}, \epsilon) = \mathcal{B}(\mathbf{x}, \epsilon, \Delta)$. We use $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ to denote the d -dimensional standard Gaussian distribution, and let ν_d be its probability measure. For the one dimensional case, we use $\Phi(x)$ to denote the cumulative distribution function (CDF) of $\mathcal{N}(0, 1)$, and use $\Phi^{-1}(x)$ to denote its inverse function. For any function $g : \mathcal{Z} \rightarrow \mathcal{X}$ and probability measure ν defined over \mathcal{Z} , $g_*(\nu)$ denotes the push-forward measure of ν . The ℓ_2 -norm of a vector $\mathbf{x} \in \mathbb{R}^n$ is defined as $\|\mathbf{x}\|_2 = (\sum_{i \in [n]} x_i^2)^{1/2}$.

2 Related Work

Several recent works (Gilmer et al., 2018; Mahloujifar et al., 2019a; Shafahi et al., 2019; Dohmatob, 2019; Bhagoji et al., 2019) derived theoretical bounds on maximum achievable adversarial robustness using

isoperimetric inequality under different assumptions of the input space. For instance, based on the assumption that the input data are uniformly distributed over two concentric n -spheres (Gilmer et al., 2018) or the underlying metric probability space satisfies a concentrated property (Mahloujifar et al., 2019a), any classifier with constant test error was proven to be vulnerable to adversarial perturbations sublinear to the input dimension. Shafahi et al. (2019) showed that adversarial examples are inevitable, provided the maximum density of the underlying input distribution is small relative to uniform density. However, none of the above theoretical works provide any experiments to justify the imposed assumptions hold for real datasets, thus it is unclear whether the derived theoretical bounds are meaningful for typical image distributions. Our work belongs to this line of research, but encompasses the practical goal of understanding the robustness limits for real image distributions.

The most related literature to ours is Fawzi et al. (2018), which proved a classifier-independent upper bound on intrinsic robustness, provided the underlying distribution is well captured by a smoothed generative model with Gaussian latent space and small Lipschitz parameter. However, their proposed theory cannot be applied to image distributions that lie on a low-dimensional, non-smooth manifold, as their framework requires examples from different classes to be close enough in the latent space. In contrast, our proposed theoretical bounds on intrinsic robustness are more general in that they can be applied to non-smoothed data manifolds, such as image distributions generated by conditional models. In addition, we propose an empirical method to estimate the intrinsic robustness on the generated image distributions under worst-case ℓ_2 perturbations.

Mahloujifar et al. (2019b) proposed to understand the inherent limitations of robust learning using heuristic methods to measure the concentration of measure based on a given set of i.i.d. samples. However, it is unclear to what extent the estimated sample-based concentration approximates the actual intrinsic robustness with respect to the underlying data distribution. In comparison, we assume the underlying data distribution can be captured by a conditional generative model and directly study the robustness limit on the generated data distribution.

3 Preliminaries

We focus on the task of image classification. Let $(\mathcal{X}, \mu, \Delta)$ be a metric probability space, where $\mathcal{X} \subseteq \mathbb{R}^n$ denotes the input space, μ is a probability distribution over \mathcal{X} and Δ is some distance metric defined

on \mathcal{X} . Suppose there exists a ground-truth function, $f^* : \mathcal{X} \rightarrow [K]$, that gives a label to any image $\mathbf{x} \in \mathcal{X}$, where $[K]$ denotes the set of all possible class labels. The objective of classification is to learn a function $f : \mathcal{X} \rightarrow [K]$ that approximates f^* well. In the context of adversarial examples, f is typically evaluated based on *risk*, which captures the classification accuracy of f on normal examples, and *adversarial risk*, which captures the classifier’s robustness against adversarial perturbations:

Definition 3.1. Let $(\mathcal{X}, \mu, \Delta)$ be a metric probability space and f^* be the ground-truth classifier. For any classifier f , the *risk* of f is defined as:

$$\text{Risk}_\mu(f) = \Pr_{\mathbf{x} \sim \mu} [f(\mathbf{x}) \neq f^*(\mathbf{x})].$$

The *adversarial risk* of f against perturbations with strength ϵ in metric Δ is defined as:

$$\text{AdvRisk}_\mu^\epsilon(f) = \Pr_{\mathbf{x} \sim \mu} [\exists \mathbf{x}' \in \mathcal{B}(\mathbf{x}, \epsilon) \text{ s.t. } f(\mathbf{x}') \neq f^*(\mathbf{x}')].$$

Other definitions of adversarial risk also exist in literature, such as the definition used in Madry et al. (2018) and the one proposed in Fawzi et al. (2018). However, these definitions are equivalent to each other under the assumption that small perturbations do not change the ground-truth labels. Another closely-related definition for adversarial robustness is the expected distance to the nearest error (see Diochnos et al. (2018) for the relation between these definitions). Our results can be applied to this definition as well.

Under different assumptions of the metric probability space, previous works proved model-independent bounds on adversarial robustness. *Intrinsic robustness*, defined originally by Mahloujifar et al. (2019b), captures the maximum adversarial robustness that can be achieved for a given robust learning problem:

Definition 3.2. Using the same settings as in Definition 3.1 and let \mathcal{F} be some class of classifiers. The *intrinsic robustness* with respect to \mathcal{F} is defined as:

$$\text{Rob}_\mu^\epsilon(\mathcal{F}) = 1 - \inf_{f \in \mathcal{F}} \{\text{AdvRisk}_\mu^\epsilon(f)\}.$$

In this work, we consider the class of imperfect classifiers that have risk at least some $\alpha > 0$.

Motivated by the great success of producing natural-looking images using conditional generative adversarial nets (GANs) (Mirza and Osindero, 2014; Odena et al., 2017; Brock et al., 2019), we assume the underlying data distribution μ can be modeled by some *conditional generative model*. A generative model can be seen as a function $g : \mathcal{Z} \rightarrow \mathcal{X}$ that maps some latent distribution, usually assumed to be multivariate Gaussian, to some generated distribution over \mathcal{X} .

Conditional generative models incorporate the additional class information into the data generating process. A conditional generative model can be considered as a set of generative models $\{g_i\}_{i \in [K]}$, where images from the i -th class can be generated by transforming latent Gaussian vectors through g_i . More rigorously, we say a probability distribution μ can be generated by a conditional generative model $\{(g_i, p_i)\}_{i \in [K]}$, if $\mu = \sum_{i=1}^K p_i \cdot (g_i)_*(\nu_d)$, where K is the total number of different class labels, and $p_i \in [0, 1]$ represents the probability of sampling an image from class i .

Based on the conditional generative model, we introduce the definition of *in-distribution adversarial risk*:

Definition 3.3. Consider the same settings as in Definition 3.1. Suppose μ can be captured by a conditional generative model $\{(g_i, p_i)\}_{i \in [K]}$. For any given classifier f , the *in-distribution adversarial risk* of f against ϵ -perturbations is defined as:

$$\begin{aligned} \text{In-AdvRisk}_\mu^\epsilon(f) = & \Pr_{(\mathbf{x}, i) \sim \mu} [\exists \mathbf{z}' \in \mathcal{Z} \text{ s.t. } g_i(\mathbf{z}') \in \mathcal{B}(\mathbf{x}, \epsilon) \\ & \text{and } f(g_i(\mathbf{z}')) \neq f^*(g_i(\mathbf{z}'))]. \end{aligned}$$

Given the fact that the in-distribution adversarial risk restricts the adversarial examples to be on the image manifold, it holds that, for any classifier f , $\text{In-AdvRisk}_\mu^\epsilon(f) \leq \text{AdvRisk}_\mu^\epsilon(f)$. As will be shown in the next section, such a notion of in-distribution adversarial risk is closely related to the intrinsic robustness for the considered class of imperfect classifiers.

4 Main Theoretical Results

In this section, we present our main theoretical results on intrinsic robustness, provided the underlying distribution can be modeled by some conditional generative model (our results and proof techniques could also be easily applied to unconditional generative models). Based on the underlying generative process, the following local Lipschitz condition connects perturbations in the image space to the latent space.

Condition 4.1. Let $g : \mathbb{R}^d \rightarrow \mathcal{X}$ be a generative model that maps the latent Gaussian distribution ν_d to some generated distribution. Consider Euclidean distance as the distance metric for \mathbb{R}^d , and Δ as the metric for \mathcal{X} . Given $r > 0$, g is said to be $L(r)$ -locally Lipschitz with probability at least $1 - \delta$, if it satisfies

$$\begin{aligned} \Pr_{\mathbf{z} \sim \nu_d} [\forall \mathbf{z}' \in \mathcal{B}(\mathbf{z}, r), \\ \Delta(g(\mathbf{z}'), g(\mathbf{z})) \leq L(r) \|\mathbf{z}' - \mathbf{z}\|_2] \geq 1 - \delta. \end{aligned}$$

As the main tool for bounding the intrinsic robustness, we present the Gaussian Isoperimetric inequality

for the sake of completeness. This inequality, proved by Borell (1975) and Sudakov and Tsirelson (1978), bounds the minimum expansion of any subset with respect to the standard Gaussian measure.

Lemma 4.2 (Gaussian Isoperimetric Inequality). Consider metric probability space $(\mathbb{R}^d, \nu_d, \|\cdot\|_2)$, where ν_d is the probability measure for d -dimensional standard Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$, and $\|\cdot\|_2$ denotes the Euclidean distance. For any subset $\mathcal{E} \subseteq \mathbb{R}^d$ and $r \geq 0$, let $\mathcal{E}_r = \{\mathbf{z} \in \mathbb{R}^d : \exists \mathbf{z}' \in \mathcal{E}, \text{ s.t. } \|\mathbf{z} - \mathbf{z}'\|_2 \leq r\}$ be the r -expansion of \mathcal{E} , then it holds that

$$\nu_d(\mathcal{E}_r) \geq \Phi(\Phi^{-1}(\nu_d(\mathcal{E})) + r), \quad (4.1)$$

where $\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x \exp(-u^2/2) \cdot du$ is the CDF of $\mathcal{N}(0, 1)$, and $\Phi^{-1}(x)$ denotes its inverse function.

In particular, when \mathcal{E} belongs to the set of half-spaces, the equality is achieved in (4.1).

Making use of the Gaussian Isoperimetric Inequality and the local Lipschitz condition of the conditional generator, the following theorem proves a lower bound on the (in-distribution) adversarial risk for any given classifier, provided the underlying distribution can be captured by a conditional generative model.

Theorem 4.3. Let $(\mathcal{X}, \mu, \Delta)$ be a metric probability space and $f^* : \mathcal{X} \rightarrow [K]$ be the underlying ground-truth. Suppose μ can be generated by a conditional generative model $\{(g_i, p_i)\}_{i \in [K]}$. Given $\epsilon > 0$, suppose there exist constants $r > 0$ and $\delta \in (0, 1]$ such that for any $i \in [K]$, g_i satisfies $L_i(r)$ -local Lipschitz property with probability at least $1 - \delta$ and $r \cdot L_i(r) \geq \epsilon$. Then for any classifier f , it holds that

$$\begin{aligned} \text{AdvRisk}_\mu^\epsilon(f) &\geq \text{In-AdvRisk}_\mu^\epsilon(f) \\ &\geq \sum_{i=1}^K p_i \cdot \Phi\left(\Phi^{-1}(\text{Risk}_{\mu_i}(f)) + \frac{\epsilon}{L_i(r)}\right) - \delta, \end{aligned}$$

where $\mu_i = (g_i)_*(\nu_d)$ is the pushforward measure of ν_d though g_i , for any $i \in [K]$.

We provide a proof in Appendix A.1. Theorem 4.3 suggests the (in-distribution) adversarial risk is related to the risk on each data manifold and the ratio between the perturbation strength and the Lipschitz constant.

The following theorem, proved in Appendix A.2, gives a theoretical upper bound on the intrinsic robustness with respect to the class of imperfect classifiers.

Theorem 4.4. Under the same setting as in Theorem 4.3, let $L_{\max}(r) = \max_{i \in [K]} L_i(r)$. Consider the class of imperfect classifiers $\mathcal{F}_\alpha = \{f : \text{Risk}_\mu(f) \geq \alpha\}$ with $\alpha > 0$, then the intrinsic robustness with respect to

\mathcal{F}_α can be bounded as,

$$\begin{aligned} \text{Rob}_\mu^\epsilon(\mathcal{F}_\alpha) &\leq 1 + \delta - \min_{i \in [K]} \left\{ p_i \cdot \Phi\left(\Phi^{-1}\left(\frac{\alpha}{p_i}\right) + \frac{\epsilon}{L_{\max}(r)}\right) \right\}, \end{aligned}$$

provided that $\alpha/p_i \leq 1$ for any $i \in [K]$. In addition, if we consider the family of classifiers that have conditional risk at least α for each class, namely $\tilde{\mathcal{F}}_\alpha = \{f : \text{Risk}_{\mu_i}(f) \geq \alpha, \forall i \in [K]\}$, then the intrinsic robustness with respect to $\tilde{\mathcal{F}}_\alpha$ can be bounded by

$$\text{Rob}_\mu^\epsilon(\tilde{\mathcal{F}}_\alpha) \leq 1 + \delta - \sum_{i=1}^K p_i \cdot \Phi\left(\Phi^{-1}(\alpha) + \frac{\epsilon}{L_{\max}(r)}\right).$$

Remark 4.5. Theorem 4.4 shows that if the data distribution can be captured by a conditional generative model, the intrinsic robustness bound with respect to imperfect classifiers will largely depend on the ratio ϵ/L_{\max} . For instance, if we assume the ratio $\epsilon/L_{\max} = 1$, then Theorem 4.4 suggests that no classifier with initial risk at least 5% can achieve robust accuracy exceeding 75% for the assumed data generating process. In addition, if we assume the local Lipschitz parameter L_{\max} is some constant, then adversarial robustness is indeed not achievable for high-dimensional data distributions, provided the perturbation strength ϵ is sublinear to the input dimension, which is the typical setting considered.

Remark 4.6. The intrinsic robustness is closely related to the in-distribution adversarial risk. For the class of classifiers \mathcal{F}_α , one can prove that the intrinsic robustness is equivalent to the maximum achievable in-distribution adversarial robustness:

$$\text{Rob}_\mu^\epsilon(\mathcal{F}_\alpha) = 1 - \inf_{f \in \mathcal{F}_\alpha} \{\text{In-AdvRisk}_\mu^\epsilon(f)\}. \quad (4.2)$$

Trivially, $\text{AdvRisk}_\mu^\epsilon(f) \geq \text{In-AdvRisk}_\mu^\epsilon(f)$ holds for any f . For a given $f \in \mathcal{F}_\alpha$, one can construct an $h_f \in \mathcal{F}_\alpha$ such that $h_f(\mathbf{x}) = f(\mathbf{x})$ if $\mathbf{x} \in \mathcal{E}_f \cap \mathcal{M}$ and $h_f(\mathbf{x}) = f^*(\mathbf{x})$ otherwise, where $\mathcal{E}_f = \{\mathbf{x} \in \mathcal{X} : f(\mathbf{x}) \neq f^*(\mathbf{x})\}$ denotes the error region of f and \mathcal{M} is the considered image manifold. The construction immediately suggests $\text{In-AdvRisk}_\mu^\epsilon(f) = \text{AdvRisk}_\mu^\epsilon(h_f)$, which implies,

$$\begin{aligned} \inf_{f \in \mathcal{F}_\alpha} \{\text{In-AdvRisk}_\mu^\epsilon(f)\} &= \inf_{f \in \mathcal{F}_\alpha} \{\text{AdvRisk}_\mu^\epsilon(h_f)\} \\ &\geq \inf_{f \in \mathcal{F}_\alpha} \{\text{AdvRisk}_\mu^\epsilon(f)\}. \end{aligned}$$

Combining both directions proves the soundness of (4.2). This equivalence suggests the in-distribution adversarial robustness of any classifier in \mathcal{F}_α can be viewed as a lower bound on the actual intrinsic robustness, which motivates us to study the intrinsic robustness by estimating the in-distribution adversarial robustness of trained robust models in our experiments.



(a) ACGAN Generated MNIST



(b) BigGAN Generated ImageNet

Figure 1: Illustration of the generated images using different conditional models. For BigGAN generated images, we select 10 specific classes from the 1000 ImageNet classes (corresponding to the 10 image classes in CIFAR-10).

5 Experiments

This section provides our empirical evaluations of the intrinsic robustness on real image distributions to evaluate the tightness of our bound. We test our bound on two image distributions generated using MNIST (LeCun et al., 1998) and ImageNet (Deng et al., 2009) datasets. Code for all our experiments is available at <https://github.com/xiaozhanguva/Intrinsic-Rob>.

5.1 Conditional GAN Models

Instead of directly evaluating the robustness on real datasets, we make use of conditional GAN models to generate datasets from the learned data distributions and evaluate the robustness of several state-of-the-art robust models trained on the generated dataset for a fair comparison with the theoretical robustness limits. Note that this approach is only feasible with conditional generative models as unconditional models cannot provide the corresponding labels for the generated data samples. For MNIST, we adopt ACGAN (Odena et al., 2017) which features an additional auxiliary classifier for better conditional image generation. The ACGAN model generates 28×28 images from a 100-dimension latent space concatenated with an additional 10-dimension one-hot encoding of the conditional class labels. For ImageNet, we adopt the BigGAN model (Brock et al., 2019) which is the state-of-the-art GAN model in conditional image generation. It generates 128×128 images from a 120-dimension latent space. We down-sampled the generated images to 32×32 for efficiency propose. We consider a standard Gaussian¹

¹The original BigGAN model uses truncated Gaussian. We adapted it to standard Gaussian distribution.

as the latent distribution for both conditional generative models. Figure 1 shows examples of the generated MNIST and ImageNet images. For both figures, each column of images corresponds to a particular label class of the considered dataset.

5.2 Local Lipschitz Constant Estimation

From Theorem 4.4, we observe that given a class of classifiers with risk at least α , the derived intrinsic robustness upper bound is mainly decided by the perturbation strength ϵ and the local Lipschitz constant $L(r)$. While ϵ is usually predesignated in common robustness evaluation settings, the local Lipschitz constant $L(r)$ is unknown for most real world tasks. Computing an exact Lipschitz constant of a deep neural network is a difficult open problem. Thus, instead of obtaining the exact value, we approximate $L(r)$ using a sample-based approach with respect to the generative models.

Recalling Definition 4.1, we consider Δ as the ℓ_2 distance and $g(\mathbf{z})$ and $g(\mathbf{z}')$ are easy to compute via the generator network. Computing $L(r)$, however, is much more complicated as it requires obtaining a maximum value within a radius- r ball. To deal with this, our approach approximates $L(r)$ by sampling N points in the neighborhood around \mathbf{z} and takes the maximum value as the estimation of the true maximum value within the ball. Since the definition of local Lipschitz is probabilistic, we take multiple samples of the latent vectors \mathbf{z} to estimate the local Lipschitz constant $L(r)$. The estimation procedure is summarized in Algorithm 1, which gives an underestimate of the underlying truth. Developing better Lipschitz estimation methods is an active area in machine learning research, but is not the main focus of this work.

Table 1: Local Lipschitz constants of ACGAN model on MNIST classes with $r = 0.5$ and $\delta = 0.001$.

Class	digit 0	digit 1	digit 2	digit 3	digit 4	digit 5	digit 6	digit 7	digit 8	digit 9
$L(r)$	7.9	8.6	8.3	7.8	10.3	11.0	9.5	7.8	9.3	10.9

 Table 2: Local Lipschitz constants of BigGAN on 10 selected ImageNet classes with $r = 0.5$ and $\delta = 0.001$.

Class	airliner	jeep	goldfinch	tabby cat	hartebeest	Maltese dog	bullfrog	sorrel	pirate ship	pickup
$L(r)$	13.1	14.5	11.7	12.4	10.4	11.3	9.4	13.0	13.1	14.9

Algorithm 1 Local Lipschitz Estimation

Input: number of samples S , number of local neighbors per sample N , r , δ

for $i = 1, \dots, S$ **do**

 Generate a latent space sample \mathbf{z}_i

 Generate N samples $\{\hat{\mathbf{z}}_i^j\}_{j=1}^N$ within $\mathcal{B}_r(\mathbf{z}_i)$

$L_i = \max_j \frac{\|g(\hat{\mathbf{z}}_i^j) - g(\mathbf{z}_i)\|_2}{\|\hat{\mathbf{z}}_i^j - \mathbf{z}_i\|_2}$

end for

Output: $(1 - \delta)$ -percentile of $\{L_i\}_{i=1}^S$

Tables 1 and 2 summarize the local Lipschitz constants estimated for the trained ACGAN and BigGAN generators conditioned on each class. For both conditional generators, we set $S = 1000$, $N = 2000$, $r = 0.5$ and $\delta = 0.001$ in Algorithm 1 for Lipschitz estimation. For BigGAN, the specifically selected 10 classes from ImageNet are reported in Table 2. In addition, the reported estimates are averaged over 10 repeated trials, where the standard deviation varies from 0.2 to 0.6 for ACGAN and varies from 0.3 to 1.1 for BigGAN.

Compared with unconditional generative models, conditional ones generate each class using a separate generator. Thus, the local Lipschitz constant of each class-conditioned generator is expected to be smaller than that of unconditional ones, as the within-class variation is usually much smaller than the between-class variation for a given classification dataset. For instance, we trained an unconditional GAN generator (Goodfellow et al., 2014) on MNIST dataset, which yields an overall local Lipschitz constant of 27.01 from Algorithm 1 under the same parameter settings. If we plug in this estimated Lipschitz constant into the theoretical results in Fawzi et al. (2018), the implied intrinsic robustness bound is in fact vacuous (above 1) with perturbations strength $\epsilon \leq 3.0$ in ℓ_2 distance.

5.3 Comparisons with Robust Classifiers

We compare our derived intrinsic robustness upper bound with the empirical adversarial robustness achieved by the current state-of-the-art defense meth-

ods under ℓ_2 perturbations. Specifically, we consider three robust training methods: *LP-Certify*: optimization-based certified robust defense (Wong et al., 2018); *Adv-Train*: PGD attack based adversarial training (Madry et al., 2018); and *TRADES*: adversarial training by accuracy and robustness trade-off (Zhang et al., 2019). We adopt these robust training methods to train robust classifiers over a set of generated training images and evaluate their robustness on the corresponding generated test set.

For MNIST, we use our trained ACGAN model to generate 10 classes of hand-written digits with 60,000 training images and 10,000 testing images. For ImageNet, we use the BigGAN model to generate 10 selected classes of images, which contains 50,000 images for training set and 10,000 images for test set. We refer to the 10-class BigGAN generated dataset as ‘ImageNet10’. We set $\epsilon = 3.0$ for training robust models using Adv-Train and TRADES for both generated datasets, whereas we only train the LP-based certified robust classifier with $\epsilon = 2.0$ on generated MNIST data, as it is not able to scale with ImageNet10 as well as generated MNIST with larger ϵ (see Appendix B.1 for all the selected hyper-parameters and network architectures).

A commonly-used method to evaluate the robustness of a given model is by performing carefully-designed adversarial attacks. Here we adopt the PGD attack (Madry et al., 2018), and report the robust accuracy (classification accuracy on inputs generated using the PGD attack) as the empirically measured model robustness. We test both the natural classification accuracy and the robustness of the aforementioned adversarially trained classifiers under ℓ_2 perturbations with perturbation strength ϵ selected from $\{1.0, 2.0, 3.0\}$. See Appendix B.1 for PGD parameter settings.

Table 3 compares the empirically measured robustness of the trained robust classifiers and the derived theoretical upper bound on intrinsic robustness. More specifically, $\epsilon = 0$ corresponds to the standard classification. For empirically measured robust accuracy when $\epsilon > 0$, we report both the mean and the standard deviation over 10 repeated trials. For computing

Table 3: Comparisons between the empirically measured robustness of adversarially trained classifiers and the implied theoretical intrinsic robustness bound on the conditional generated datasets.

	ACGAN generated MNIST				BigGAN generated ImageNet10			
	$\epsilon = 0.0$	$\epsilon = 1.0$	$\epsilon = 2.0$	$\epsilon = 3.0$	$\epsilon = 0.0$	$\epsilon = 1.0$	$\epsilon = 2.0$	$\epsilon = 3.0$
LP-Certify	88.3%	$74.0 \pm 0.4\%$	$51.1 \pm 0.6\%$	$23.5 \pm 0.3\%$	-	-	-	-
Adv-Train	97.2%	$93.1 \pm 0.2\%$	$83.5 \pm 0.3\%$	$58.9 \pm 0.4\%$	82.1%	$67.8 \pm 0.3\%$	$47.1 \pm 0.4\%$	$23.4 \pm 0.4\%$
TRADES	98.3%	$94.8 \pm 0.2\%$	$81.8 \pm 0.4\%$	$57.7 \pm 0.4\%$	83.4%	$68.5 \pm 0.3\%$	$49.1 \pm 0.5\%$	$27.8 \pm 0.5\%$
Our Bound	-	98.2%	97.8%	97.2%	-	83.5%	81.8%	80.0%

our theoretical robust bounds, we set the risk threshold $\alpha = 0.015$ for generated MNIST and $\alpha = 0.15$ for ImageNet10, to reflect the best natural accuracy achieved by the considered robust classifiers.

Under most settings, there exists a large gap between the robust limit implied by our theory and the best adversarial robustness achieved by state-of-the-art robust classifiers. For instance, Adv-Train and TRADES only achieve less than 50% robust accuracy on the generated ImageNet10 data with $\epsilon = 2.0$, whereas the estimated robustness bound is as high as 81.8%. The gap becomes even larger when we increase the perturbation strength ϵ . In contrast to the previous theoretical results on artificial distributions, for these image classification problems we cannot simply conclude from the intrinsic robustness bound that adversarial examples are inevitable. This huge gap between the empirical robustness of the best current image classifiers and the estimated theoretical bound suggests that either there is a way to train better robust models or that there exist other explanations for the inherent limitations of robust learning against adversarial examples.

5.4 In-distribution Adversarial Robustness

In Section 5.3, we empirically show the unconstrained robustness of existing robust classifiers is far below the intrinsic robustness upper bound implied by our theory for real distributions. However, it is not clear whether the reason is that current robust training methods are far from perfect, or that our derived upper bound is not tight enough due to the Lipschitz relaxation step used for proving such bound. In this section, we empirically study the in-distribution adversarial risk for a better characterization of the actual intrinsic robustness. As shown in Remark 4.6, the in-distribution adversarial robustness of any classifier with risk at least α can be regarded as a lower bound for the intrinsic robustness $\text{Rob}_\mu^\epsilon(\mathcal{F}_\alpha)$. This provides us a more accurate characterization of the intrinsic robustness bound and enables better understanding of intrinsic robustness.

While there are many types of attack algorithms in the literature that can be used to evaluate the uncon-

strained robustness of a given classifier in the image space, little has been done in terms of how to evaluate the in-distribution robustness. In order to empirically evaluate the in-distribution robustness, we straightforwardly formulate the following optimization problem to find adversarial examples on the image manifold:

$$\min_{\mathbf{z}} \mathcal{L}(f(G(\mathbf{z}, y)), y) \text{ s.t. } \|G(\mathbf{z}, y) - \mathbf{x}\|_2 \leq \epsilon, \quad (5.1)$$

where $\mathbf{z} \in \mathbb{R}^d$, \mathbf{x} is the data sample in the image space to be attacked, f is the given classifier, and \mathcal{L} denotes the adversarial loss function. The goal of (5.1) is to optimize the latent vector to lower the adversarial loss (make the robust classifier mis-classify some generated images) while keeping the distance between the generated image and the test image within ϵ perturbation limit. The key difficulty in solving (5.1) lies in the fact that we cannot perform any type of projection operations as we are optimizing over \mathbf{z} but the constraints are imposed on the generated image space $G(\mathbf{z}, y)$. This prohibits the use of common attack algorithms such as PGD. In order to solve (5.1), we transform (5.1) into the following Lagrangian formulation:

$$\min_{\mathbf{z}} \|G(\mathbf{z}, y) - \mathbf{x}\|_2 + \lambda \cdot \mathcal{L}(f(G(\mathbf{z}, y)), y). \quad (5.2)$$

This formulation ignores the perturbation constraint of ϵ and tries to find the in-distribution adversarial examples with the smallest possible perturbation. In order to evaluate the intrinsic robustness under a given ϵ perturbation budget, we need to further check all in-distribution adversarial examples found and only count those with perturbations within the ϵ constraint. Note that even though (5.2) provides us a feasible way to compute the in-distribution robustness of a classifier, equation (5.2) itself could be hard to solve in general. First, it is not obvious how to initialize \mathbf{z} . Random initialization of \mathbf{z} could lead to bad local optima which prevent the optimizer from efficiently solving (5.2) or even finding a \mathbf{z} that could make $G(\mathbf{z}, y)$ close enough to \mathbf{x} . Second, the hyper-parameter λ could be quite sensitive to different test examples. Failing to choose a proper λ could also lead to failures in finding in-distribution adversarial examples within ϵ

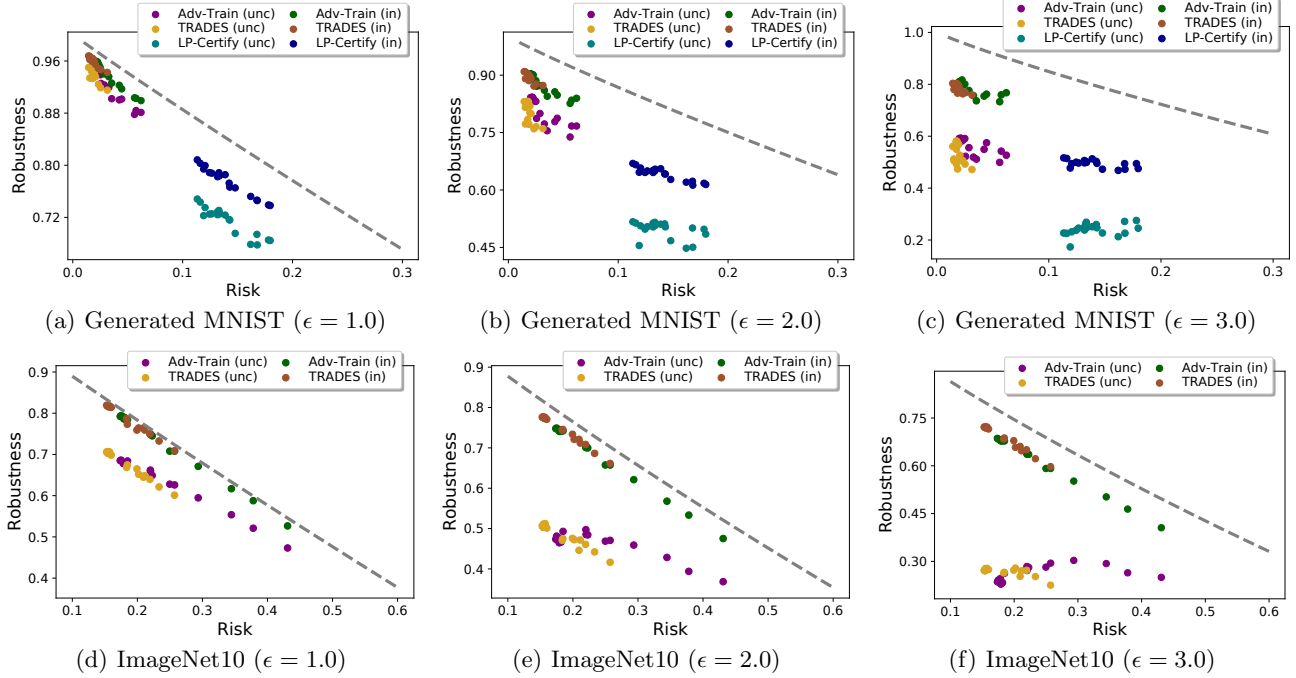


Figure 2: Comparisons between the theoretical intrinsic robustness bound and the empirically estimated unconstrained/in-distribution adversarial robustness, denoted as “unc” and “in” in the legend, of models produced during robust training on the generated data under ℓ_2 . In each subfigure, the dotted curve line represents the theoretical bound on intrinsic robustness with horizontal axis denoting the different choice of α .

constraint. In order to tackle the aforementioned challenges, we propose to solve another optimization problem for the initialization of \mathbf{z} and adopt binary search for the best choice of λ (see Appendix B.2 for more details of our implementation).

Figure 2 summarizes results from our empirical evaluations on intrinsic robustness of the generated MNIST and ImageNet10 data. We evaluate the empirical robustness of three types of robust training methods at different time points during the training procedure. To be more specific, we evaluate the robustness of the intermediate models produced every 5 training epochs. For each method, we plot both the unconstrained robustness measured by PGD attacks and the in-distribution robustness measured using the aforementioned strategies. In addition, based on the local Lipschitz constants estimated in Section 5.2, we plot the implied theoretical bound on intrinsic robustness as the dotted line curve for direct comparison.

Compared with the intrinsic robustness upper bound (dotted curve line), the unconstrained robustness of various robustly-trained models is much smaller, and the gap between them becomes more obvious as we increase ϵ . This aligns with our observations in Section 5.3. However under all the considered settings, the estimated in-distribution adversarial robustness is much higher than the unconstrained one and closer to

the theoretical upper bound, especially for the ImageNet10 data. Note that according to Remark 4.6, the actual intrinsic robustness $\text{Rob}_\mu^\epsilon(\mathcal{F}_\alpha)$ should lie between the in-distribution robustness of any given classifier with risk at least α and the derived intrinsic robustness upper bound. Observing the big gap between the estimated in-distribution and unconstrained robustness of various robustly trained models, one would expect the current state-of-the-art robust models are still far from approaching the actual intrinsic robustness limit for real image distributions.

6 Conclusions

We studied the intrinsic robustness of typical image distributions using conditional generative models. By deriving theoretical upper bounds on intrinsic robustness and providing empirical estimates on the generated image distributions, we observed a large gap between the theoretical intrinsic robust limit and the best robustness achieved by state-of-the-art robust classifiers. Our results imply that the inevitability of adversarial examples claimed in recent theoretical studies, such as Fawzi et al. (2018), do not apply to real image distributions, and suggest that there is a need for deeper understanding on the intrinsic robustness limitations for real data distributions.

Acknowledgements

This research was sponsored in part by the National Science Foundation SaTC-1717950 and SaTC-1804603, and additional support from Amazon, Baidu, and Intel. The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing any funding agencies.

References

- Athalye, A., Carlini, N., and Wagner, D. (2018). Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *International Conference on Machine Learning (ICML)*.
- Bhagoji, A. N., Cullina, D., and Mittal, P. (2019). Lower bounds on adversarial robustness from optimal transport. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Borell, C. (1975). The Brunn-Minkowski inequality in Gauss space. *Inventiones mathematicae*, 30(2):207–216.
- Brock, A., Donahue, J., and Simonyan, K. (2019). Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations (ICLR)*.
- Carlini, N. and Wagner, D. (2017). Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy*.
- Chakraborty, A., Alam, M., Dey, V., Chattopadhyay, A., and Mukhopadhyay, D. (2018). Adversarial attacks and defences: A survey. *arXiv preprint arXiv:1810.00069*.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). ImageNet: A Large-Scale Hierarchical Image Database. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Diochnos, D., Mahloujifar, S., and Mahmoody, M. (2018). Adversarial risk and robustness: General definitions and implications for the uniform distribution. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Dohmatob, E. (2019). Generalized no free lunch theorem for adversarial robustness. In *International Conference on Machine Learning (ICML)*.
- Fawzi, A., Fawzi, H., and Fawzi, O. (2018). Adversarial vulnerability for any classifier. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Gilmer, J., Metz, L., Faghri, F., Schoenholz, S. S., Raghu, M., Wattenberg, M., and Goodfellow, I. (2018). Adversarial spheres. *arXiv preprint arXiv:1801.02774*.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Goodfellow, I., Shlens, J., and Szegedy, C. (2015). Explaining and harnessing adversarial examples. In *International Conference on Learning Representations (ICLR)*.
- Gowal, S., Dvijotham, K., Stanforth, R., Bunel, R., Qin, C., Uesato, J., Mann, T., and Kohli, P. (2019). Scalable verified training for provably robust image classification. In *International Conference on Computer Vision (ICCV)*.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hinton, G., Deng, L., Yu, D., Dahl, G., Mohamed, A.-r., Jaitly, N., Senior, A., Vanhoucke, V., Nguyen, P., Kingsbury, B., et al. (2012). Deep neural networks for acoustic modeling in speech recognition. *IEEE Signal processing magazine*, 29.
- Katz, G., Barrett, C., Dill, D. L., Julian, K., and Kochenderfer, M. J. (2017). Reluplex: An efficient SMT solver for verifying deep neural networks. In *International Conference on Computer Aided Verification*.
- LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., et al. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. (2018). Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations (ICLR)*.
- Mahloujifar, S., Diochnos, D. I., and Mahmoody, M. (2019a). The curse of concentration in robust learning: Evasion and poisoning attacks from concentration of measure. In *AAAI Conference on Artificial Intelligence*.
- Mahloujifar, S., Zhang, X., Mahmoody, M., and Evans, D. (2019b). Empirically measuring concentration: Fundamental limits on intrinsic robustness. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Mirza, M. and Osindero, S. (2014). Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*.
- Odena, A., Olah, C., and Shlens, J. (2017). Conditional image synthesis with auxiliary classifier gans. In *International Conference on Machine Learning (ICML)*.

- Papernot, N., McDaniel, P., Wu, X., Jha, S., and Swami, A. (2016). Distillation as a defense to adversarial perturbations against deep neural networks. In *IEEE Symposium on Security and Privacy*.
- Raghunathan, A., Steinhardt, J., and Liang, P. (2018). Certified defenses against adversarial examples. In *International Conference on Learning Representations (ICLR)*.
- Shafahi, A., Huang, W. R., Studer, C., Feizi, S., and Goldstein, T. (2019). Are adversarial examples inevitable? In *International Conference on Learning Representations (ICLR)*.
- Sinha, A., Namkoong, H., and Duchi, J. (2018). Certifying some distributional robustness with principled adversarial training. In *International Conference on Learning Representations (ICLR)*.
- Sudakov, V. N. and Tsirelson, B. S. (1978). Extremal properties of half-spaces for spherically invariant measures. *Journal of Soviet Mathematics*, 9(1):9–18.
- Sutskever, I., Hinton, G. E., and Krizhevsky, A. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems (NeurIPS)*.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. (2014). Intriguing properties of neural networks. In *International Conference on Learning Representations (ICLR)*.
- Tjeng, V., Xiao, K. Y., and Tedrake, R. (2019). Evaluating robustness of neural networks with mixed integer programming. In *International Conference on Learning Representations (ICLR)*.
- Tramer, F., Carlini, N., Brendel, W., and Madry, A. (2020). On adaptive attacks to adversarial example defenses. *arXiv preprint arXiv:2002.08347*.
- Wang, S., Chen, Y., Abdou, A., and Jana, S. (2018). MixTrain: Scalable training of formally robust neural networks. *arXiv preprint arXiv:1811.02625*.
- Wang, Y., Ma, X., Bailey, J., Yi, J., Zhou, B., and Gu, Q. (2019). On the convergence and robustness of adversarial training. In *International Conference on Machine Learning (ICML)*.
- Wong, E. and Kolter, Z. (2018). Provable defenses against adversarial examples via the convex outer adversarial polytope. In *International Conference on Machine Learning (ICML)*.
- Wong, E., Schmidt, F., Metzen, J. H., and Kolter, J. Z. (2018). Scaling provable adversarial defenses. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Zhang, H., Yu, Y., Jiao, J., Xing, E., El Ghaoui, L., and Jordan, M. (2019). Theoretically principled trade-off between robustness and accuracy. In *International Conference on Machine Learning (ICML)*.