

Adversarial Robustification via Text-to-Image Diffusion Models

Daewon Choi^{1*}, Jongheon Jeong^{2*}, Huiwon Jang¹, and Jinwoo Shin¹

¹ Korea Advanced Institute of Science and Technology (KAIST)

² Korea University

{daeone0920,huiwoen0516,jinwoos}@kaist.ac.kr jonghj@korea.ac.kr

Abstract. Adversarial robustness has been conventionally believed as a challenging property to encode for neural networks, requiring plenty of training data. In the recent paradigm of adopting off-the-shelf models, however, access to their training data is often infeasible or not practical, while most of such models are not originally trained concerning adversarial robustness. In this paper, we develop a scalable and model-agnostic solution to achieve adversarial robustness without using any data. Our intuition is to view recent text-to-image diffusion models as “adaptable” denoisers that can be optimized to specify target tasks. Based on this, we propose: (a) to initiate a denoise-and-classify pipeline that offers provable guarantees against adversarial attacks, and (b) to leverage a few synthetic reference images generated from the text-to-image model that enables novel adaptation schemes. Our experiments show that our data-free scheme applied to the pre-trained CLIP could improve the (provable) adversarial robustness of its diverse zero-shot classification derivatives (while maintaining their accuracy), significantly surpassing prior approaches that utilize the full training data. Not only for CLIP, we also demonstrate that our framework is easily applicable for robustifying other visual classifiers efficiently. Code is available at <https://github.com/ChoiDae1/robustify-T2I>.

Keywords: Adversarial robustness · Certified robustness · Text-to-image diffusion models · Denoised smoothing · Zero-shot robustification

1 Introduction

Arguably, recent breakthroughs in deep learning have been largely driven by massive data and model scaling [4, 28, 45, 57], which enabled many unprecedented capabilities in computer vision [27, 28, 45, 46, 52]. The *worst-case* behaviors of models at scale, however, have been relatively under-explored in the literature, despite their increasing practical relevance [6, 77, 80, 81]. *Adversarial robustness* [2, 5, 36, 58] is one of popular objectives in this context: specifically, it aims to build a model that makes consistent predictions for *every* input perturbation within a small, often imperceptible, bound. Although it has been demonstrated that many

* Equal contribution.




Fig. 1: Comparison of clean and robust ($\|\varepsilon\|_2 \leq 1.0$) accuracy on zero-shot classification: our framework (a) not only maintains the original accuracy of CLIP [45]; but also (b) significantly improves its robust accuracy, *e.g.* compared to Mao et al. [37].

types of (“natural”) robustness can benefit from proper data scaling [16, 45, 66], *e.g.*, combined with recent *vision-language models* [27, 34, 45, 56], the trend does not seem to hold for adversarial robustness so far [6, 37], particularly highlighting its challenging nature.

Despite being a desirable property, pursuing adversarial robustness in practice has been viewed as a costly design decision. This is possibly due to that most of existing techniques for adversarial robustness require a specialized, less-scalable training scheme [13, 36, 63] from enough data [47, 54], even followed by significant performance trade-offs [60, 75]. As a result, many of the existing pre-trained, off-the-shelf models widely used in the community remain susceptible to adversarial attacks [6, 31, 81]. Provided that the specific training data used for such off-the-shelf models is frequently not publicly accessible, it becomes increasingly challenging for end users to secure adversarial robustness for these models, *e.g.*, to incorporate them into security-concerned applications [9, 73].

A recent work [37] has attempted to address the challenge through an adversarial contrastive fine-tuning scheme on CLIP [45] using external data such as ImageNet, in order to transfer the obtained robustness to other zero-shot downstream tasks. Yet, the method is limited in a sense that: (a) it is applicable only for vision-language models, (b) requiring a substantial amount of training data, *e.g.*, as large as ImageNet in scale, to ensure its effectiveness. Furthermore, we observe that the approach is susceptible to an *overfitting* to the fine-tuning data: *e.g.*, the accuracy on downstream tasks often degrades after the fine-tuning (see Fig. 1a), and fails in transferring the robustness on datasets that significantly varies from ImageNet (see Tab. 2).

Contribution. In this paper, by leveraging recent text-to-image diffusion models, we propose a scalable framework that does not require any external datasets in *robustifying* image classifiers. Our framework is based on *denoised smoothing* [53], a recent technique that constructs a provably-robust classifier from neural networks (*i.e.*, *certified defense* [13, 35, 63, 65]) through a “denoise-and-classify” pipeline, with a denoiser model on top of classifier. Previous works upon denoised smoothing have only considered denoiser models that are optimized for

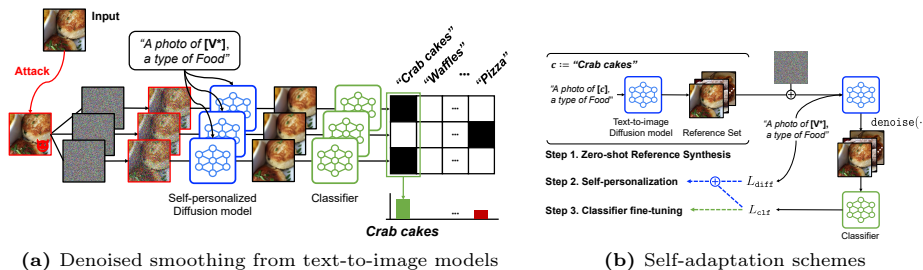


Fig. 2: An overview of the proposed framework: (a) during inference, we perform denoised smoothing with a self-personalized text-to-image diffusion model, having provable guarantees on adversarial robustness (Sec. 3.1); (b) by utilizing synthetic references from the text-to-image model, one can adapt both diffusion model and classifier for robustness (Sec. 3.2).

the target classification task [7, 26, 68], fully utilizing their training data. Here, we aim to extend its capability up to a next level, *i.e.*, robustifying without using data. To this end, we utilize *text-to-image super-resolution diffusion models* (like Imagen [52] or DeepFloyd-IF) with careful details we found, *e.g.*, on setting proper diffusion timestep. Upon the pipeline, we propose to utilize a few synthetic reference samples generated from the text-to-image model based on the textual labels of target task, and develop two adaptation schemes that boost the robustness: *viz.*, fine-tuning (a) of diffusion model via classifier-guided personalization, and (b) of classifier from denoised samples (see Fig. 2).

We conduct a comprehensive series of experiments to verify the effectiveness of our proposed framework, focusing on modern scenarios of classification using off-the-shelf models and a wide spectrum of datasets [3, 22, 29, 51]. First, we show that our method could significantly enhance CLIP with a new robustness-accuracy frontier, even outperforming the previous method [37] that utilizes the full ImageNet training data to robustify CLIP (see Fig. 1b). In particular, the effectiveness of our method is “provably” confirmed with the *certified* robust accuracy, which guarantees the lower-bound of empirical robust accuracy. The robustness gains from our framework are shown to be much more consistent across datasets compared to the prior work [37] that has been struggling to offer robustness when the target dataset significantly varies from those for fine-tuning, *e.g.*, EuroSAT upon ImageNet. Next, we verify that our proposed framework can also be effective in robustifying other generic vision classifiers, *e.g.*, an ImageNet pre-trained ResNet-50, even surpassing standard approaches to obtain adversarial robustness such as adversarial training [36].

To summarize, we make the following contributions:

1. To the best of our knowledge, our framework is the first approach toward robustifying any given (off-the-shelf) vision classifiers without using any data. We utilize recent text-to-image diffusion models as denoisers that can be adapted to target tasks via text-conditioning, and show that incorporating them into the inference of pre-trained classifiers can be a scalable approach.

2. We further propose to generate a few reference samples re-utilizing the text-to-image diffusion model, and leverage them to adapt for individual tasks. We show that fine-tuning both the text-to-image diffusion model as well as classifier can be jointly beneficial to boost adversarial robustness.
3. We evaluate our framework in robustifying CLIP for a variety of zero-shot classification tasks: it could not only offer state-of-the-art robustness on the benchmark, but also show consistent gains across a wider range of datasets. We also verify the proposed framework offers robustness to generic vision classifiers other than CLIP, obtaining better robustness even compared to, *e.g.*, the popular adversarial training from scratch.

2 Preliminaries

Adversarial Robustness [5, 58] refers to a property of a classifier, say f , to make consistent prediction under *every* possible perturbations δ within a boundary, *e.g.*, an ℓ_2 -ball: it requires $f(x + \delta) = y$ for *any* $\|\delta\|_2 \leq \varepsilon$. In this respect, adversarial robustness of f for an input x can be measured by the *minimum-distance* of adversarial perturbation [5]:

$$R(x, y; f) := \min_{f(x') \neq y} \|x - x'\|_2. \quad (1)$$

One of the key challenges to achieving adversarial robustness is the hardness of accurately measuring Eq. (1), which has eventually falsified the robustness claim of many previous works in the literature [2, 59, 61].

Randomized Smoothing [13, 32] is currently one of the state-of-the-art methods in obtaining provable guarantees on adversarial robustness from neural network-based classifiers. Given a classifier f and an input x , randomized smoothing makes an inference by taking a majority vote of $f(x + \delta)$ for random Gaussian noise $\delta \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$. Specifically, it defines a *smoothed classifier* \hat{f} as follows:

$$\hat{f}(x) := \operatorname{argmax}_{c \in \mathcal{Y}} \mathbb{P}_{\delta \sim \mathcal{N}(0, \sigma^2 \mathbf{I})} [f(x + \delta) = c]. \quad (2)$$

Then, Cohen et al. [13] have shown that the adversarial robustness of \hat{f} at x is guaranteed by a lower-bound \underline{R} :

$$R(x, y; \hat{f}) \geq \sigma \cdot \Phi^{-1}(p_{\hat{f}}(x, y)) =: \underline{R}, \quad \text{where } p_{\hat{f}}(x, y) := \mathbb{P}_{\delta} [f(x + \delta) = y], \quad (3)$$

provided that $\hat{f}(x) = y$, otherwise $R(x, y; \hat{f}) := 0$. Here, Φ denotes the standard Gaussian CDF. Remark from Eq. (3) that the *certified robustness* \underline{R} of \hat{f} depends on $p_{\hat{f}}$, which is essentially the accuracy of f at noisy inputs $x + \delta$.

Denosed Smoothing [53] is a recent framework for randomized smoothing that has enabled a more scalable design. Specifically, it constructs the base classifier f for smoothing as a “denoise-and-classify” pipeline, which concatenates a

Gaussian denoiser, say $\mathbf{denoise}(\cdot)$, with any standard classifier f_{clf} as follows:

$$f(\hat{x}) := f_{\text{clf}}(\mathbf{denoise}(\hat{x})). \quad (4)$$

Here, a “good” denoiser $\mathbf{denoise}(\cdot)$ should accurately reconstruct the semantic of x from $\hat{x} := x + \delta$ with high probability of $\delta \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$. In practice, this often requires the denoiser to be sufficiently optimized for the input distribution of target tasks, otherwise the performance of \hat{f} is significantly limited by the denoiser [53]. For tasks where an accurate denoiser is available, however, denoised smoothing can offer a strong design for randomized smoothing: *e.g.*, Carlini et al. [7] achieved state-of-the-art certified robustness on ImageNet [51] by adopting a high-fidelity ImageNet diffusion model [14] as the denoiser. The idea of leveraging diffusion models for “denoise-and-classify” has been also considered as an empirical defense [41, 72] (*i.e.*, not certifiable): yet, such approaches also necessitate a separate diffusion model trained for the target dataset.

Diffusion Model [23, 40] aims to learn a generative distribution $p_{\text{data}}(x)$ by gradually denoising (or *reverse* process) from noisy inputs from so-called *diffusion* (or *forward*) process. Formally, it first defines x_t from a pure Gaussian noise $\varepsilon \sim \mathcal{N}(0, \mathbf{I})$ and a *timestep* $t \in [0, T]$, given x :

$$x_t := \sqrt{\alpha_t} \cdot x + \sqrt{1 - \alpha_t} \cdot \varepsilon, \quad (5)$$

where the factor $\alpha_t \in [0, 1]$ is a constant determined by t , which schedules the amount of noise typically as a monotonically decreasing function of t , *i.e.* x_t becomes noisier toward the unit Gaussian with increasing t . The key design of recent diffusion models is to parametrize a *noise estimator* $\varepsilon_\theta(x_t, t)$, trained over $p_{\text{data}}(x)$ and $t \in [0, T]$, which aims to predict the noise ε added to x_t given t .

Text-to-Image Diffusion Model [48, 52] is a particular instance of diffusion models, which have recently demonstrated remarkable capabilities in generating high-fidelity images from natural language descriptions. Specifically, it considers modified architectures for noise estimator to condition on text, *i.e.*, as the form of $\varepsilon_\theta(x_t, t, \tau_\theta(c))$, where τ_θ is a *text encoder* that maps a textual prompt c into an embedding vector. Depending on specific designs, the latest architectures for text-to-image diffusion models roughly fall into two categories to handle high-resolution inputs: (a) *latent diffusion models* [48], which first map x into a latent space of lower-resolution, and (b) *cascaded diffusion models* [24, 52], which train a lower-resolution diffusion model in the pixel space followed by multiple *super-resolution* diffusion models of increasing resolutions.

3 Text-to-Image Diffusion Models for Robustification

In this section, we introduce a scalable and model-agnostic framework to obtain adversarial robustness from image classifiers without accessing training data. Given an image classifier $f : \mathcal{X} \rightarrow \mathcal{Y}$ trained on a data distribution $p_{\text{data}}(x, y)$, we

aim to construct a new classifier \hat{f} that is adversarially-robust, without explicit knowledge of p_{data} , *e.g.*, for fine-tuning. The minimal assumption we make here is the textual knowledge of classes, *i.e.*, $\mathcal{Y} := \{c_i\}_{i=1}^K$ and each c_i is given in text – which is common in the recent literature of zero-shot classification [27, 34, 45]. From this information, \hat{f} is required to (a) maximize $\mathbb{E}_{(x,y)}[R(x,y;\hat{f})]$ Eq. (1), while (b) minimizing the accuracy trade-off from robustifying f .

The key ingredient of our proposed framework to this end is the recent *text-to-image diffusion models*, particularly those with *pixel-level, cascaded diffusion models*, *e.g.*, Imagen [52] or DeepFloyd-IF. Overall, our framework utilizes the model mainly in two different ways: (a) to be applied as the denoiser of the denoised smoothing pipeline Eq. (4) (see Sec. 3.1), and (b) to enable a fine-tuning of the classifier via personalization (see Sec. 3.2). The overall framework is illustrated in Fig. 2.

3.1 Denoised Smoothing from Text-to-Image Diffusion Models

We first propose to utilize recent text-to-image diffusion models by means of their performance as a “zero-shot” denoiser, so that it can be incorporated into the *denoised smoothing* pipeline Eq. (4). Consider an input x and its Gaussian perturbation, say $\hat{x} := x + \delta$, where $\delta \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ for a noise strength σ . Then, as observed by Carlini et al. [7], one can correspond \hat{x} into a *timestep* of diffusion process Eq. (5), *i.e.*, to $x_{\hat{t}}$ for some \hat{t} . Specifically, it follows by:

$$\sigma^2 = \frac{1 - \alpha_{\hat{t}}}{\alpha_{\hat{t}}}. \quad (6)$$

With this relationship, one can search over the noise schedule α_t of a diffusion model for its corresponding timestep \hat{t} given σ , which makes \hat{x} compatible with the inherent denoiser of diffusion models, *e.g.*, by scaling it with $\sqrt{\alpha_{\hat{t}}}$.

Now, we introduce the detailed design components in adopting text-to-image models, of the general form of $\varepsilon_{\theta}(x_t, t, \tau_{\theta}(c))$, to define **denoise**(\cdot) in Eq. (4).

Need for Pixel-based Diffusion Models. Existing off-the-shelf text-to-image diffusion models are often based on *latent diffusion model* architecture [48], *e.g.*, Stable Diffusion. Remark that, however, the pipeline of denoised smoothing Eq. (4) requires a denoiser to directly denoise a given noisy input \hat{x} , which is in *pixel-space*, making this kind of diffusion models incompatible for the pipeline. In this respect, our framework focuses on adopting *cascaded diffusion models* into the pipeline, such as Imagen [52] and DeepFloyd-IF, another popular design choice for recent text-to-image models and those indeed consist of pixel-level diffusion models (of different resolutions).

Super-resolution Diffusion Model as a Denoiser. More specifically, recall that cascaded diffusion models generally consist of (a) a low-resolution (*e.g.*, 64×64) text-conditional diffusion model, followed by (b) multiple stages of super-resolution diffusion models (*e.g.*, from 64×64 to 256×256) to enable higher-resolution generations in a scalable manner. Among these different diffusion models, our framework draws attention to the particular attribute of the

super-resolution models by means of effective denoisers. The choice is motivated by an intuition that super-resolution modules in cascaded diffusion models are more likely to be biased to “reconstruct” the original contents of x given \hat{x} in performing denoising, rather than generating new visual cues from scratch, which can be particularly beneficial in the pipeline of denoised smoothing.

Formally, super-resolution diffusion models in cascaded designs typically parameterize a noise estimator as follows:

$$\varepsilon_{\theta}(x_t, t, \tau_{\theta}(c)|\bar{x}_{t'}, t'), \quad (7)$$

where (x_t, t) is a noisy input and its timestep at the output resolution, and τ_{θ} is a text encoder for conditioning. The additional condition compared to the standard models, *i.e.*, $(\bar{x}_{t'}, t')$, is from the previous (lower-resolution) module, processed by (a) first interpolating the output of the previous module up to the output resolution, followed by (b) mixing with Gaussian noise using a certain timestep t' .

Timestep Correction. In our context of adapting the model for denoised smoothing, we propose to set *both inputs* of x_t and $\bar{x}_{t'}$ in Eq. (7) by $\sqrt{\alpha_{\hat{t}}}\cdot\hat{x}$, *i.e.*, by $x_t = \bar{x}_{t'} = \sqrt{\alpha_{\hat{t}}}\cdot\hat{x}$, where \hat{t} is the timestep searched with respect to Eq. (6). In this way, the super-resolution module ε_{θ} is “self-conditioned” by the information available from \hat{x} . A surprisingly important detail to make this design work is on the timestep t' : we find that setting a *higher value* of timestep for t' than \hat{t} , despite being $x_t = \bar{x}_{t'}$, is crucial for the denoising performance of the model. Specifically, we consider a *correction factor* $k > 1$ as a hyperparameter to scale t' , and propose to set:

$$t' := k \cdot \hat{t}. \quad (8)$$

This interesting behavior, specific to super-resolution diffusion models, can be explained by considering that $\bar{x}_{t'}$ in Eq. (7) is originally assumed to be “upsampled” before applying a noise. Therefore, $\bar{x}_{t'}$ is naturally expected to consist of narrower range of spatial frequencies, whereas the input given, $\sqrt{\alpha_{\hat{t}}}\cdot\hat{x}$, is directly from higher resolution: using higher values for t' is an effective way to reduce such excessive frequency information present in \hat{x} , given that it corresponds to an increased blurring in the denoising process.

Overall Pipeline. Putting together, our proposed denoised smoothing based pipeline is obtained using a text-conditional, super-resolution diffusion model ε_{θ} . Specifically, given a noisy input $\hat{x} := x + \delta$, where $\delta \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$, we define a denoiser function for Eq. (4) as follows:

$$\mathbf{denoise}_{\theta}(\hat{x}) := \hat{x} - \sigma \cdot \varepsilon_{\theta}(\sqrt{\alpha_{\hat{t}}}\hat{x}, \hat{t}, \tau_{\theta}(\mathbf{C}(\text{“ ”})))|_{\sqrt{\alpha_{\hat{t}}}\hat{x}, k\hat{t}}. \quad (9)$$

Here, $\mathbf{C}(c)$ is a pre-defined textual “template” that implants a given (textual) label c , specific per task. For example, we use $\mathbf{C}(c) := \text{“}A \textit{ photo of a } \{c\}, \textit{ a type of food.} \text{”}$ for the Food dataset [3] in our experiments, following Radford et al. [45]. For the case of Eq. (9), which considers a zero-shot case that the label is not given, we simply put the *empty string* “ ” for c . Once we have a concrete $\mathbf{denoise}(\cdot)$ at

hand, as Eq. (9), any classifier f that combines $\text{denoise}(\cdot)$ can now be robustified via randomized smoothing Eq. (2). The *smoothed classifier* it returns, \hat{f} , is provably robust within the certified radius it guarantees by Eq. (3) for each x . In practice, the overall smoothing procedure is statistically estimated with n *i.i.d.* Gaussian noise from $\mathcal{N}(0, \sigma^2 \mathbf{I})$: we provide the details for the estimation in Appendix A.5.

3.2 Self-adaptation Schemes

Upon our framework introduced in Sec. 3.1, we propose to re-utilize the text-to-image diffusion model to further improve its robustness. Consequently, we propose a two-step adaptation scheme of models, again using only the knowledge of textual label set $\mathcal{Y} = \{c_i\}_{i=1}^K$, *i.e.*, without using concrete data in \mathcal{X} .

Reference Set Synthesis. We start by leveraging the text-to-image model to synthesize a few reference images from the textual labels. Concretely, for a given textual label $c \sim \mathcal{Y}$, we obtain its corresponding prompt $\mathbf{C}(c)$ and use it to generate a synthetic image x^g by conditioning it into the text-to-image diffusion model. Repeating this process, we obtain high quality reference set $D^g = \{(x_i^g, c_i)\}_{i=1}^K$ only from the information of \mathcal{Y} .

Classifier-Guided Self-personalization. For a given reference set D^g , we next perform a fine-tuning of the text-to-image diffusion model. We adopt *Dream-Booth* [50] to this end, one of state-of-the-art method for personalizing text-to-image models. Specifically, it fine-tunes the given noise estimator network, ε_θ , with a special prompt combining a *unique identifier*, which is typically a list of meaningless characters (*e.g.*, “*sks*”), to implant the information of D^g . After the personalization, one can now use $\mathbf{C}(\text{“sks”})$ in Eq. (9) as a replacement of $\mathbf{C}(\text{“ ”})$ during its inference. Again, considering that ε_θ is a super-resolution diffusion model, we consider the following DreamBooth objective:

$$L_{\text{diff}}(\theta) := \mathbb{E}_{x^g, \varepsilon, t} [\|\varepsilon - \varepsilon_\theta(x_t^g, t, \tau_\theta(\mathbf{C}(\text{“sks”}))|x_t^g, kt)\|_2^2], \quad (10)$$

where $t \sim \mathcal{U}([0, T])$ is a random timestep, $\varepsilon \sim \mathcal{N}(0, \mathbf{I})$ is Gaussian noise, and $\tau_\theta(\mathbf{C}(\text{“sks”}))$ is the textual embedding from $\mathbf{C}(\text{“sks”})$ through the (frozen) text encoder τ_θ .

To further boost capability of text-to-image diffusion model by means of a denoiser model, particularly in the context of *denoised smoothing*, we propose to regularize the model personalization objective, L_{diff} , with the *denoised classification loss*, namely as a *classifier-guided* regularization L_{clf} of personalization. The regularization essentially simulates the “denoise-and-classify” pipeline of denoised smoothing. Specifically, for a given reference image x^g in D^g , x^g is first processed into a noisy image x_t^g via Eq. (5) using random timestep $t \sim \mathcal{U}([0, T])$, followed by the (personalized version of) zero-shot denoising Eq. (9): obtaining a denoised image $\tilde{x}^g = \text{denoise}_\theta(\frac{1}{\sqrt{\alpha_t}} \cdot x_t^g)$. Therefore, given a classifier f_ψ , we propose to additionally minimize the following loss given the pair (\tilde{x}^g, c) :

$$L_{\text{clf}}(\theta, \psi) := \mathbb{E}_{(x^g, c) \sim D^g, t} [\mathbb{C}\mathbb{E}(f_\psi(\tilde{x}^g), c)], \quad (11)$$

where $\mathbb{CE}(\cdot, \cdot)$ is the cross-entropy loss. Overall, we minimize the following objective by combining the two losses:

$$\theta^* = \arg \min_{\theta} \{L_{\text{diff}}(\theta) + \lambda \cdot L_{\text{clf}}(\theta, \psi)\}, \quad (12)$$

where $\lambda > 0$ is a hyperparameter. The text encoder τ_{θ} of the text-to-diffusion model and the classifier f_{ψ} are fixed during the personalization.

Classifier Fine-tuning. Lastly, we also apply the denoised classification loss, L_{clf} , to further optimize the classifier side: even after the personalization of the diffusion model using D^g for more accurate denoising, the classifier may still be suboptimal due to the distribution mismatch between the clean and denoised images during denoised smoothing, as also suggested by Carlini et al. [7]. By also directly minimizing the loss for such “denoise-and-classify” images, one can further reduce the gap. Specifically, we minimize:

$$\psi^* = \arg \min_{\psi} L_{\text{clf}}(\theta^*, \psi). \quad (13)$$

Similarly, here we freeze the denoiser model θ^* during the optimization (of ψ).

4 Experiments

We verify the effectiveness of our proposed framework focusing on its ability of improving adversarial robustness without using external data. As far as we are aware, our setup has not been previously explored in the literature. For comparisons, we choose the following two recent methods as our closest baselines: (a) Mao et al. [37], which fine-tune CLIP [45] on ImageNet for empirical adversarial robustness on other (zero-shot) classification tasks; and (b) Carlini et al. [7], which consider an unconditional diffusion model optimized for target task in denoised smoothing, *e.g.*, on ImageNet. We provide the experimental details, *e.g.*, datasets, architectures, evaluation, fine-tuning, *etc.*, in Appendix A.

4.1 Robustification of CLIP

Firstly, we evaluate our framework upon CLIP-B/32 [45] for mainly comparison with Mao et al. [37]. Here, we not only consider standard zero-shot classification benchmarks, but also more domain-specialized datasets that significantly vary from ImageNet. These datasets are regarded as more challenging cases for Mao et al. [37] due to their reliance on ImageNet. We further evaluate the robustification performance of our framework on ImageNet compared to Mao et al. [37] and Carlini et al. [7] directly, both of which utilize the full ImageNet training data, contrary to ours that keeps the assumption of not using data.

Results on Standard Zero-shot Benchmarks. We evaluate the robustification performance of our proposed framework to CLIP-B/32 [45] covering an extensive zero-shot classification benchmark [3, 10, 11, 17, 29, 42, 44, 69]. Specifically, we compare how much our framework effective on improving the robust

Table 1: Robust and clean accuracy (%) on 8 zero-shot classification datasets using CLIP against ℓ_2 -adversary with $\varepsilon \in \{0.5, 1.0\}$. We additionally report certified accuracy at ε for “Ours” in parentheses (see “Certified”). Bold indicates the best.

(a) Robust accuracy (%)

	Method	STL	SUN	Cars	Food	Pets	Flower	DTD	Caltech	Average
$\varepsilon = 0.5$	CLIP	10.8	1.2	0.0	1.8	2.7	0.8	2.7	12.0	4.0
	CLIP-Smooth	42.6	23.7	14.3	8.9	36.4	16.6	10.2	44.8	24.7
	Mao et al. [37]	59.4	29.9	12.5	32.9	51.2	33.5	18.8	56.2	36.8
	Ours (Certified)	80.4 (66.0)	41.8 (32.1)	33.2 (28.4)	59.0 (45.7)	68.6 (60.8)	45.2 (34.9)	29.7 (23.0)	71.3 (65.1)	53.7 (44.5)
$\varepsilon = 1.0$	CLIP	2.4	0.0	0.0	0.2	0.2	0.2	1.0	7.8	1.5
	CLIP-Smooth	16.2	5.8	1.6	0.4	6.7	4.5	5.4	18.7	7.4
	Mao et al. [37]	21.2	11.0	2.8	10.3	23.2	14.4	12.0	33.9	16.1
	Ours (Certified)	66.0 (41.2)	38.3 (22.5)	27.0 (18.9)	47.3 (28.9)	59.6 (46.5)	33.1 (18.9)	24.9 (18.2)	64.1 (55.8)	45.0 (31.4)

(b) Clean accuracy (%)

	Method	STL	SUN	Cars	Food	Pets	Flower	DTD	Caltech	Average
-	CLIP	97.8	56.8	52.7	83.0	85.7	66.3	37.8	81.9	70.3
$\varepsilon = 0.5$	CLIP-Smooth	75.0	46.8	42.1	52.3	66.7	43.5	17.2	68.3	51.5
	Mao et al. [37]	94.8	60.0	48.7	69.7	80.8	57.7	34.0	79.7	65.7
	Ours (Certified)	94.8 (90.4)	58.6 (55.4)	54.1 (49.7)	80.2 (74.5)	83.6 (81.9)	61.4 (58.7)	42.7 (38.9)	81.7 (79.1)	69.6 (66.1)
$\varepsilon = 1.0$	CLIP-Smooth	32.4	27.7	40.8	31.3	43.8	36.8	7.0	54.0	34.2
	Mao et al. [37]	93.4	58.2	42.9	61.2	77.0	53.6	30.8	78.5	62.0
	Ours (Certified)	93.8 (80.2)	59.4 (53.6)	52.9 (45.5)	78.8 (64.8)	83.1 (77.7)	58.9 (48.3)	39.1 (32.4)	81.7 (75.7)	68.5 (59.8)

and clean accuracy of CLIP on these tasks without any data, considering two ℓ_2 -adversary of budget $\varepsilon \in \{0.5, 1.0\}$. We mainly compare with (a) Mao et al. [37], an adversarial fine-tuning scheme using ImageNet, as well as with the performance of (b) the vanilla CLIP zero-shot classification and (c) CLIP-Smooth, which directly applies randomized smoothing (Eq. (2)) to the CLIP model without using a denoiser.

In Tab. 1, we compare the robust and clean accuracy of our framework with the baselines, respectively. We observe that the vanilla CLIP model is originally vulnerable to adversarial attacks: their average robust accuracy at $\varepsilon = 0.5$ is decreased from the clean accuracy by 66.3% ($70.3\% \rightarrow 4.0\%$), near the chance level. Although CLIP-Smooth obtains better average robustness than vanilla CLIP ($4.0\% \rightarrow 24.7\%$ at $\varepsilon = 0.5$) through randomized smoothing, however, it is still susceptible to perturbations with a higher bound, *i.e.*, $\varepsilon = 1.0$. Mao et al. [37] fairly outperforms these baselines: nevertheless, we observe that it tends to exhibit insufficient robustness gains for “domain-specific” datasets, *e.g.*, Cars and DTD, as further confirmed in Tab. 2.

Our framework shows a significant improvement in robustness over other baselines across entire datasets and ε . For instance, we obtain 16.9% average robust accuracy gain at $\varepsilon = 0.5$ compared to Mao et al. [37] and this discrepancy becomes more larger ($16.9\% \rightarrow 28.9\%$) as ε is larger ($0.5 \rightarrow 1.0$). Moreover, the certified robust accuracy of our framework also outperforms the (empirical)

Table 2: Robust and clean accuracy (%) on three domain-specialized benchmarks using CLIP against ℓ_2 -adversary with $\varepsilon \in \{0.5, 1.0\}$. We report certified accuracy in parentheses. Bold indicates the best and runner-up is underlined.

Method	Robust accuracy (%)			Clean accuracy (%)		
	CropDisease	EuroSAT	ISIC	CropDisease	EuroSAT	ISIC
$\varepsilon = 0.5$						
CLIP	0.0	0.0	0.0	20.9	<u>42.6</u>	<u>27.3</u>
CLIP-Smooth	1.8	<u>13.6</u>	<u>9.0</u>	7.1	16.6	14.3
Mao et al. [37]	<u>2.2</u>	0.6	4.2	16.0	25.0	26.9
Ours (Certified)	11.5 (4.3)	29.0 (11.0)	17.6 (5.8)	<u>20.3</u> (16.8)	45.2 (39.0)	35.7 (33.7)
$\varepsilon = 1.0$						
CLIP	0.0	0.0	0.0	20.9	<u>42.6</u>	27.3
CLIP-Smooth	<u>1.2</u>	<u>11.8</u>	<u>4.6</u>	4.3	18.2	10.6
Mao et al. [37]	0.6	0.0	2.2	<u>16.0</u>	20.4	22.0
Ours (Certified)	4.9 (0.8)	28.2 (5.0)	8.4 (1.4)	15.8 (8.4)	44.8 (37.8)	<u>26.3</u> (22.0)

Table 3: Robust and clean accuracy (%) on ImageNet using CLIP against ℓ_2 -adversary with $\varepsilon \in \{0.5, 1.0\}$. We report certified accuracy in parentheses. Bold indicates the best and runner-up is underlined.

Method	Data-free?	Robust accuracy (%)		Clean accuracy (%)	
		$\varepsilon = 0.5$	$\varepsilon = 1.0$	$\varepsilon = 0.5$	$\varepsilon = 1.0$
CLIP	✓	1.4	0.2	58.2	58.2
CLIP-Smooth	✓	16.8 (9.8)	2.2 (1.2)	45.2 (25.0)	35.2 (3.8)
Ours (w/o adapt)	✓	<u>40.0</u> (29.6)	31.0 (17.6)	56.2 (50.8)	55.2 (42.0)
Ours	✓	42.6 (34.2)	<u>31.4</u> (20.6)	<u>57.6</u> (53.4)	<u>56.2</u> (46.0)
Mao et al. [37]	✗	26.0	12.3	51.2	47.2
Carlini et al. [7]	✗	38.6 (30.2)	32.4 (19.8)	54.4 (49.8)	53.6 (44.2)

robust accuracy of other baselines across all datasets and ε , *i.e.*, the “lower-bound” robust accuracy already outperforms the empirical robustness of other baselines.

These all results show strong adversarial robustness without external data is possible via our framework. Moreover, both the average empirical and certified clean accuracy of our framework not only surpass other baselines but also outperform the clean accuracy of vanilla CLIP in some datasets, such as Cars (52.7% \rightarrow 54.1%), DTD (37.8% \rightarrow 42.7%). These results indicate the flexibility of our framework, which is able to trade-off between robust and clean accuracy.

Results on Domain-Specific Benchmarks. Next, we focus our evaluation on several domain-specific datasets as more challenging but practical scenarios, namely on CropDiseases [38], EuroSAT [22], and ISIC [12]: *e.g.*, EuroSAT consists of photos specifically taken from satellites. In Tab. 2, we observe that Mao et al. [37] exhibits particularly bad, even near-zero, robustness on these datasets, possibly due to the model itself being fine-tuned only from a domain represented by ImageNet. Our framework still shows consistent robustness gains here: *e.g.*, it offers a significantly higher robustness in EuroSAT at $\varepsilon = 1.0$ of 28.2%. Similarly to Tab. 1, our framework also maintains the clean accuracy of CLIP, notably even surpassing it on EuroSAT and ISIC.

Results on ImageNet. Finally, we show that our robustification scheme (without using any data) can be competitive and even better compared to those di-

Table 4: Robust and clean accuracy (%) ResNet-50 on ImageNet across training schemes, using ℓ_2 -adversary with $\varepsilon \in \{0.5, 1.0\}$. We report certified accuracy in parentheses. Bold indicates the best and runner-up is underlined.

Method	Data-free?	Robust accuracy (%)		Clean accuracy (%)	
		$\varepsilon = 0.5$	$\varepsilon = 1.0$	$\varepsilon = 0.5$	$\varepsilon = 1.0$
Standard Training	✗	5.2	1.0	74.4	74.4
+ Ours (w/o adapt)	✓	<u>56.2</u> (47.0)	44.2 (27.4)	<u>73.0</u> (67.0)	68.8 (57.2)
+ Ours	✓	57.0 (50.4)	47.8 (34.0)	70.4 (68.2)	<u>71.8</u> (60.8)
Adversarial Training [36]	✗	51.0	<u>46.8</u>	55.0	55.0
Randomized Smoothing [13]	✗	55.2 (48.6)	43.8 (37.0)	65.4 (66.8)	55.4 (57.0)
Carlini et al. [7]	✗	<u>56.2</u> (49.2)	45.2 (33.2)	72.6 (67.4)	70.0 (57.8)

rectly accessing to training data. We consider ImageNet [51] for the evaluation, considering that Mao et al. [37] fine-tunes directly on ImageNet. In addition to Mao et al. [37], we also consider Carlini et al. [7] as another baseline, by considering a denoise-and-classify pipeline that combines CLIP with an unconditional ImageNet diffusion model [14]: this can provide a clearer comparison on the effectiveness of our zero-shot denoised smoothing (Sec. 3.1).

In Tab. 3, we report robust and clean accuracy, comparing our framework with these baselines. Even while Mao et al. [37] and Carlini et al. [7] directly access the ImageNet training data, our framework achieves better (or competitive) performance to them. Compared with Mao et al. [37], we obtain not only 16.6% and 19.1% gains in empirical robust accuracy at $\varepsilon = 0.5$ and 1.0 but also 8.2% and 8.3% even from certified robust accuracy. Although Carlini et al. [7] gets a slightly better empirical robust accuracy at $\varepsilon = 1.0$, our framework still achieves a higher certified robust accuracy. Here, we notice that our proposed self-adaptation schemes (Sec. 3.2) play a crucial role for the gains: *e.g.*, it contributes to a 4.6% gain in certified accuracy (29.6% \rightarrow 34.2%) at $\varepsilon = 0.5$. Regarding the clean accuracy, our framework notably shows only a 0.6% gap $\varepsilon = 0.5$ compared to CLIP. All these results demonstrate the superiority of our framework in ensuring sufficient robust and clean accuracy, even compared with models directly accessing training data.

4.2 Robustification of Generic Vision Classifiers

We further validate the scalability of our framework in robustifying classifiers other than CLIP, particularly considering ResNet-50 pre-trained on ImageNet. Here, we regard Carlini et al. [7] as a baseline again. For a thorough comparison, we consider two additional baselines: Adversarial Training [36] and Randomized Smoothing [13], both trained from scratch using the full ImageNet training data.

Tab. 4 reports the results in robust and clean accuracy on ImageNet. Compared with the standard training, we obtain significant 51.8% (5.2% \rightarrow 57.0% at $\varepsilon = 0.5$) and 46.8% (1.0% \rightarrow 47.8% at $\varepsilon = 1.0$) gains in robust accuracy. These also outperform other baselines as well, *e.g.* it surpasses Carlini et al. [7] in even certified robust accuracy. Again, the self-adaptation schemes contributes significantly to the gains, *e.g.* by 6.6% at $\varepsilon = 1.0$.




Fig. 3: Qualitative comparisons of denoised images on varying correct factor k . We compared the denoised outputs from Eq. (9) under Gaussian noise of $\sigma = 0.25$.

4.3 Ablation Study

We perform an ablation study to investigate the individual effectiveness of our framework. Here, we main compare the certified test accuracy of smoothed classifiers on ImageNet, as well as the *average certified radius* (ACR; higher is better) [74], for a collective view of overall certified robustness.

Timestep Correction. We further analyze the influence of correction factor k in Eq. (8) both qualitatively and quantitatively. Specifically, we compare three different correction factors $k \in \{0.5, 1.0, 1.8\}$, including our default choice of $k = 1.8$ in our experiments.³ Here, $k = 1.0$ means that t' is not corrected when used for Eq. (9). In Fig. 3, we observe that using a higher value for correction factor k leads to clearer denoised outputs. Tab. 5 also confirms that the certified robustness denoised smoothing obtained is significantly impacted by the quality of the denoised samples. These results show that adjusting timestep t' is a critical design choice to make existing super-resolution diffusion models as a powerful denoiser model for denoised smoothing.

Adaptation Schemes. In Tab. 6, we perform a component-wise analysis on our adaptation schemes (in Sec. 3.2): *viz.*, fine-tuning (a) text-to-image diffusion model and (b) classifier. First, we confirm fine-tuning the text-to-image model obtains consistent gains in certified accuracy across datasets and perturbation bound ε , compared to the baseline without adaptation. The performance gains are further strengthened by also adapting the classifier.⁴ We also remark that both adaptation schemes (Eq. (10) and Eq. (11)) randomize the timestep $t \sim \mathcal{U}([0, T])$ during their fine-tuning, *i.e.*, the observed gains in $\sigma = \{0.25, 0.5\}$ are from a single fine-tuning of the same pipeline but only using different timesteps at denoising.

³ We chose k by comparing the CLIP accuracy of denoised images from a small subset of ImageNet, as reported in Appendix C.1.

⁴ We provide the complete results of certified accuracy across noise scale σ and perturbation bound ε , in Tab. 7 of Appendix C.

Table 5: Ablation study on k in Eq. (8). Bold indicates the best.

σ	k	ACR	Certified accuracy at ε (%)					
			0.0	0.25	0.5	0.75	1.0	1.25
0.25	0.5	0.152	36.2	23.0	14.3	6.6		
	1.0	0.229	46.0	35.6	25.2	13.4		
	1.8	0.277	50.8	41.8	29.6	19.2		
0.50	0.5	0.080	13.2	10.2	6.0	4.3	3.4	2.0
	1.0	0.200	28.8	22.0	16.4	11.2	8.4	6.0
	1.8	0.367	42.0	35.4	29.2	23.8	17.6	13.0

Table 6: Ablation study on the proposed self-adaptation schemes. Bold and underline indicate the best and runner-up, respectively.

ε	Adapt.		Certified accuracy at ε (%)			
	T2I	CLIP	ImageNet	STL	SUN	Food
0.5	\times	\times	29.6	55.2	28.3	43.6
	\checkmark	\times	<u>31.8</u>	66.0	30.7	43.8
	\checkmark	\checkmark	34.2	66.0	32.1	45.7
1.0	\times	\times	17.6	27.0	17.3	21.8
	\checkmark	\times	<u>19.4</u>	40.8	<u>19.3</u>	27.3
	\checkmark	\checkmark	20.6	41.2	22.5	28.9

5 Related Work

Adversarial Robustness. Since the observation of adversarial examples [20, 58], there have been continuous efforts in achieving adversarial robustness from neural networks, either in forms of *empirical defenses* [2, 5, 59] mainly based upon *adversarial training* [36, 64, 67, 75, 78], or *certified defenses* [13, 33, 65, 70, 76], depending on whether the robustness claim is provable or not. One of common beliefs in the literature has been that adversarial robustness is a property that has to be learned using concrete data [8, 21, 54]. In this work, we move away from this assumption, proposing a novel direction of robustifying with no data.

Zero-shot Visual Recognition. Traditionally, zero-shot classification [30, 43], which aims to identify novel categories not present during training, has been a challenging task in computer vision. In recent years, large-scale vision-language models [27, 45] have demonstrated remarkable capabilities in this regard, *e.g.*, compared to prior arts [1, 18, 25, 39, 49, 55]. Upon these advances in obtaining high *accuracy* under zero-shot, our work further questions whether it is possible to obtain *adversarial robustness* in a zero-shot manner as well – showing that large-scale text-to-image models can be a way to achieve this.

Transfer Learning for Robustness. Another line of research focuses on the *transferability* of adversarial robustness, *i.e.*, by utilizing an external source of data for robustness [8, 21]. For example, several works [15, 19, 62, 71] considered a meta-learning based robust training aiming for *few-shot* adaptation of adversarial robustness. Yet, they commonly rely on costly meta-training procedure from scratch, limiting their applicability to larger models. More related to our work, Mao et al. [37] have recently proposed an adversarial contrastive fine-tuning for vision-language models, in order to transfer adversarial robustness to other zero-shot classification tasks. Again, in contrast to ours, the approach still requires a substantial amount of training data, *e.g.*, as large as ImageNet in scale.

6 Conclusion

In this paper, we introduce a new formulation of robustifying vision classifiers without external data, as a more realistic concern in the era of adopting off-the-shelf models. We propose a simple-yet-effective approach to this problem, which incorporates recent text-to-image diffusion models into the inference of a classifier in novel ways. Our approach is applicable for any off-the-shelf classifiers, making it as favorable and practical to obtain (provable) adversarial robustness when the use of external data is either limited or impractical for users. We hope our approach paves the way toward more reliable and secure AI systems, along the way towards robustifying existing components that have been considered powerful but fragile against adversarial attacks. We also believe our proposal suggests interesting future research, such as extending the framework to robustify commercial, black-box APIs [53]. For example, this may require further techniques, such as zeroth-order optimization [79], as the adaptation scheme.

Acknowledgements

This work was partially supported by Center for Applied Research in Artificial Intelligence (CARAI) grant funded by Defense Acquisition Program Administration (DAPA) and Agency for Defense Development (ADD) (UD230017TD), and by Institute for Information & communications Technology Promotion (IITP) grant funded by the Korea government (MSIT) (No.RS-2019-II190075, Artificial Intelligence Graduate School Program (KAIST)). Jongheon Jeong acknowledges support from Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. RS-2019-II190079, Artificial Intelligence Graduate School Program (Korea University)), and from Korea University grant K2405671. We are grateful to the Center for AI Safety (CAIS) for generously providing compute resources that supported a significant portion of the experiments conducted in this work. We thank Jihoon Tack for providing helpful feedbacks and suggestions in preparing an earlier version of our manuscript.

References

1. Akata, Z., Reed, S., Walter, D., Lee, H., Schiele, B.: Evaluation of output embeddings for fine-grained image classification. In: CVPR (2015)
2. Athalye, A., Carlini, N., Wagner, D.: Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In: Proceedings of the 35th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 80, pp. 274–283. PMLR (10–15 Jul 2018), <http://proceedings.mlr.press/v80/athalye18a.html>
3. Bossard, L., Guillaumin, M., Van Gool, L.: Food-101 – Mining discriminative components with random forests. In: ECCV (2014)
4. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in Neural Information Processing Systems* **33**, 1877–1901 (2020)
5. Carlini, N., Athalye, A., Papernot, N., Brendel, W., Rauber, J., Tsipras, D., Goodfellow, I., Madry, A., Kurakin, A.: On evaluating adversarial robustness. arXiv preprint arXiv:1902.06705 (2019)
6. Carlini, N., Nasr, M., Choquette-Choo, C.A., Jagielski, M., Gao, I., Awadalla, A., Koh, P.W., Ippolito, D., Lee, K., Tramer, F., et al.: Are aligned neural networks adversarially aligned? arXiv preprint arXiv:2306.15447 (2023)
7. Carlini, N., Tramer, F., Dvijotham, K.D., Rice, L., Sun, M., Kolter, J.Z.: (Certified!!) adversarial robustness for free! In: ICLR (2023)
8. Carmon, Y., Raghunathan, A., Schmidt, L., Duchi, J.C., Liang, P.S.: Unlabeled data improves adversarial robustness. *Advances in Neural Information Processing Systems* **32** (2019)
9. Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M., Elhadad, N.: Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In: Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 1721–1730 (2015)
10. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., Vedaldi, A.: Describing textures in the wild. In: CVPR (2014)

11. Coates, A., Ng, A., Lee, H.: An analysis of single layer networks in unsupervised feature learning. In: AISTATS (2011)
12. Codella, N.C., Gutman, D., Celebi, M.E., Helba, B., Marchetti, M.A., Dusza, S.W., Kalloo, A., Liopyris, K., Mishra, N., Kittler, H., et al.: Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic). In: 2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018). pp. 168–172. IEEE (2018)
13. Cohen, J., Rosenfeld, E., Kolter, Z.: Certified adversarial robustness via randomized smoothing. In: ICML. pp. 1310–1320. PMLR (2019)
14. Dhariwal, P., Nichol, A.Q.: Diffusion models beat GANs on image synthesis. In: NeurIPS (2021)
15. Dong, J., Wang, Y., Lai, J., Xie, X.: Improving adversarially robust few-shot image classification with generalizable representations. In: CVPR. pp. 9015–9024. IEEE (2022)
16. Fang, A., Ilharco, G., Wortsman, M., Wan, Y., Shankar, V., Dave, A., Schmidt, L.: Data determines distributional robustness in contrastive language image pre-training (CLIP). In: International Conference on Machine Learning. pp. 6216–6234. PMLR (2022)
17. Fei-Fei, L., Fergus, R., Perona, P.: Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. Computer Vision and Pattern Recognition Workshop (2004)
18. Frome, A., Corrado, G.S., Shlens, J., Bengio, S., Dean, J., Ranzato, M., Mikolov, T.: DeViSE: A deep visual-semantic embedding model. In: NeurIPS (2013)
19. Goldblum, M., Fowl, L., Goldstein, T.: Adversarially robust few-shot learning: A meta-learning approach. In: NeurIPS (2020)
20. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: International Conference on Learning Representations (2015)
21. Gowal, S., Rebuffi, S.A., Wiles, O., Stimberg, F., Calian, D.A., Mann, T.A.: Improving robustness using generated data. *Advances in Neural Information Processing Systems* **34**, 4218–4233 (2021)
22. Helber, P., Bischke, B., Dengel, A., Borth, D.: EuroSAT: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE J. Sel. Topics in Appl. Earth Observ. and Remote Sensing* **12**(7), 2217–2226 (2019)
23. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020)
24. Ho, J., Saharia, C., Chan, W., Fleet, D.J., Norouzi, M., Salimans, T.: Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research* **23** (2022), <http://jmlr.org/papers/v23/21-0635.html>
25. Huang, H., Wang, C., Yu, P.S., Wang, C.D.: Generative dual adversarial network for generalized zero-shot learning. In: CVPR (2019)
26. Jeong, J., Shin, J.: Multi-scale diffusion denoised smoothing. In: NeurIPS (2023)
27. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q.V., Sung, Y., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: ICML (2021)
28. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. *arXiv:2304.02643* (2023)
29. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: Proceedings of the IEEE International Conference on Computer Vision Workshops. pp. 554–561 (2013)

30. Lampert, C.H., Nickisch, H., Harmeling, S.: Learning to detect unseen object classes by between-class attribute transfer. In: CVPR (2009)
31. Lan, L.C., Zhang, H., Wu, T.R., Tsai, M.Y., Wu, I., Hsieh, C.J., et al.: Are AlphaZero-like agents robust to adversarial perturbations? *Advances in Neural Information Processing Systems* **35**, 11229–11240 (2022)
32. Lecuyer, M., Atlidakis, V., Geambasu, R., Hsu, D., Jana, S.: Certified robustness to adversarial examples with differential privacy. In: 2019 IEEE Symposium on Security and Privacy (SP). pp. 656–672. IEEE (2019)
33. Leino, K., Wang, Z., Fredrikson, M.: Globally-robust neural networks. In: International Conference on Machine Learning. pp. 6212–6222. PMLR (2021)
34. Li, J., Li, D., Xiong, C., Hoi, S.C.H.: BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: ICML (2022)
35. Li, L., Xie, T., Li, B.: SoK: Certified robustness for deep neural networks. In: IEEE Symposium on Security and Privacy. pp. 1289–1310. IEEE (2023)
36. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. In: ICLR (2018)
37. Mao, C., Geng, S., Yang, J., Wang, X., Vondrick, C.: Understanding zero-shot adversarial robustness for large-scale models. In: ICLR (2023)
38. Mohanty, S.P., Hughes, D.P., Salathé, M.: Using deep learning for image-based plant disease detection. *Frontiers in plant science* **7**, 1419 (2016)
39. Ni, J., Zhang, S., Xie, H.: Dual adversarial semantics-consistent network for generalized zero-shot learning. In: NeurIPS (2019)
40. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: ICML (2021)
41. Nie, W., et al.: Diffusion models for adversarial purification. In: ICML (2022)
42. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: Proceedings of the Indian Conference on Computer Vision, Graphics and Image Processing (2008)
43. Palatucci, M., Pomerleau, D., Hinton, G.E., Mitchell, T.M.: Zero-shot learning with semantic output codes. In: NeurIPS (2009)
44. Parkhi, O.M., Vedaldi, A., Zisserman, A., Jawahar, C.V.: Cats and dogs. In: IEEE Conference on Computer Vision and Pattern Recognition (2012)
45. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML. pp. 8748–8763. PMLR (2021)
46. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: ICML (2021)
47. Rebuffi, S.A., Gowal, S., Calian, D.A., Stimberg, F., Wiles, O., Mann, T.A.: Data augmentation can improve robustness. *Advances in Neural Information Processing Systems* **34**, 29935–29948 (2021)
48. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR. pp. 10684–10695 (2022)
49. Romera-Paredes, B., Torr, P.H.S.: An embarrassingly simple approach to zero-shot learning. In: ICML (2015)
50. Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dream-Booth: Fine tuning text-to-image diffusion models for subject-driven generation. In: CVPR. pp. 22500–22510 (2023)
51. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision* **115**(3), 211–252 (2015). <https://doi.org/10.1007/s11263-015-0816-y>

52. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, S.K.S., Lopes, R.G., Karagol-Ayan, B., Salimans, T., Ho, J., Fleet, D.J., Norouzi, M.: Photorealistic text-to-image diffusion models with deep language understanding. In: *NeurIPS* (2022), http://papers.nips.cc/paper_files/paper/2022/hash/ec795aeadae0b7d230fa35cbaf04c041-Abstract-Conference.html
53. Salman, H., Sun, M., Yang, G., Kapoor, A., Kolter, J.Z.: Denoised smoothing: A provable defense for pretrained classifiers. *Advances in Neural Information Processing Systems* **33**, 21945–21957 (2020)
54. Schmidt, L., Santurkar, S., Tsipras, D., Talwar, K., Madry, A.: Adversarially robust generalization requires more data. *Advances in Neural Information Processing Systems* **31** (2018)
55. Schönfeld, E., Ebrahimi, S., Sinha, S., Darrell, T., Akata, Z.: Generalized zero- and few-shot learning via aligned variational autoencoders. In: *CVPR* (2019)
56. Singh, A., Hu, R., Goswami, V., Couairon, G., Galuba, W., Rohrbach, M., Kiela, D.: FLAVA: A foundational language and vision alignment model. In: *CVPR* (2022)
57. Singh, M., Duval, Q., Alwala, K.V., Fan, H., Aggarwal, V., Adcock, A., Joulin, A., Dollár, P., Feichtenhofer, C., Girshick, R., Girdhar, R., Misra, I.: The effectiveness of MAE pre-pretraining for billion-scale pretraining. *arXiv preprint arXiv:2303.13496* (2023)
58. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R.: Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199* (2013)
59. Tramèr, F., Carlini, N., Brendel, W., Madry, A.: On adaptive attacks to adversarial example defenses. In: *Advances in Neural Information Processing Systems*. vol. 33 (2020)
60. Tsipras, D., Santurkar, S., Engstrom, L., Turner, A., Madry, A.: Robustness may be at odds with accuracy. In: *International Conference on Learning Representations* (2019), <https://openreview.net/forum?id=SyxAb30cY7>
61. Uesato, J., O’donoghue, B., Kohli, P., Oord, A.: Adversarial risk and the dangers of evaluating against weak attacks. In: *International Conference on Machine Learning*. pp. 5025–5034. PMLR (2018)
62. Wang, R., Xu, K., Liu, S., Chen, P.Y., Weng, T.W., Gan, C., Wang, M.: On fast adversarial robustness adaptation in model-agnostic meta-learning. In: *ICLR* (2021)
63. Wang, S., Zhang, H., Xu, K., Lin, X., Jana, S., Hsieh, C.J., Kolter, J.Z.: Beta-CROWN: Efficient bound propagation with per-neuron split constraints for neural network robustness verification. *Advances in Neural Information Processing Systems* **34**, 29909–29921 (2021)
64. Wang, Y., Zou, D., Yi, J., Bailey, J., Ma, X., Gu, Q.: Improving adversarial robustness requires revisiting misclassified examples. In: *International Conference on Learning Representations* (2020), <https://openreview.net/forum?id=rk10g6EFwS>
65. Wong, E., Kolter, J.Z.: Provable defenses against adversarial examples via the convex outer adversarial polytope. In: *ICML. JMLR Workshop and Conference Proceedings*, vol. 80, pp. 5283–5292. JMLR (2018)
66. Wortsman, M., Ilharco, G., Kim, J.W., Li, M., Kornblith, S., Roelofs, R., Lopes, R.G., Hajishirzi, H., Farhadi, A., Namkoong, H., et al.: Robust fine-tuning of zero-shot models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7959–7971 (2022)

67. Wu, D., Xia, S.T., Wang, Y.: Adversarial weight perturbation helps robust generalization. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) *Advances in Neural Information Processing Systems*. vol. 33, pp. 2958–2969. Curran Associates, Inc. (2020), <https://proceedings.neurips.cc/paper/2020/file/1ef91c212e30e14bf125e9374262401f-Paper.pdf>
68. Xiao, C., Chen, Z., Jin, K., Wang, J., Nie, W., Liu, M., Anandkumar, A., Li, B., Song, D.: DensePure: Understanding diffusion models for adversarial robustness. In: *The Eleventh International Conference on Learning Representations* (2022)
69. Xiao, J., Hays, J., Ehinger, K.A., Oliva, A., Torralba, A.: SUN database: Large-scale scene recognition from abbey to zoo. In: *CVPR* (2010)
70. Xiao, K.Y., Tjeng, V., Shafiullah, N.M.M., Madry, A.: Training for faster adversarial robustness verification via inducing ReLU stability. In: *International Conference on Learning Representations* (2019), <https://openreview.net/forum?id=BJfIVjAcKm>
71. Yin, C., Tang, J., Xu, Z., Wang, Y.: Adversarial meta-learning. arXiv preprint arXiv:1806.03316 (2018)
72. Yoon, J., Hwang, S.J., Lee, J.: Adversarial purification with score-based generative models. In: *International Conference on Machine Learning*. pp. 12062–12072. PMLR (2021)
73. Yurtsever, E., Lambert, J., Carballo, A., Takeda, K.: A survey of autonomous driving: Common practices and emerging technologies. *IEEE Access* **8**, 58443–58469 (2020)
74. Zhai, R., Dan, C., He, D., Zhang, H., Gong, B., Ravikumar, P., Hsieh, C.J., Wang, L.: MACER: Attack-free and scalable robust training via maximizing certified radius. In: *ICLR* (2020)
75. Zhang, H., Yu, Y., Jiao, J., Xing, E., El Ghaoui, L., Jordan, M.: Theoretically principled trade-off between robustness and accuracy. In: *International Conference on Machine Learning*. pp. 7472–7482. PMLR (2019)
76. Zhang, H., Chen, H., Xiao, C., Goyal, S., Stanforth, R., Li, B., Boning, D., Hsieh, C.J.: Towards stable and efficient training of verifiably robust neural networks. In: *International Conference on Learning Representations* (2020), <https://openreview.net/forum?id=Skxuk1rFwB>
77. Zhang, J., Yi, Q., Sang, J.: Towards adversarial attack on vision-language pre-training models. In: *Proceedings of the 30th ACM International Conference on Multimedia*. pp. 5005–5013 (2022)
78. Zhang, J., Xu, X., Han, B., Niu, G., Cui, L., Sugiyama, M., Kankanhalli, M.: Attacks which do not kill training make adversarial learning stronger. In: III, H.D., Singh, A. (eds.) *Proceedings of the 37th International Conference on Machine Learning*. *Proceedings of Machine Learning Research*, vol. 119, pp. 11278–11287. PMLR (13–18 Jul 2020)
79. Zhang, Y., et al.: How to robustify black-box ML models? A zeroth-order optimization perspective. In: *ICLR* (2022)
80. Zhou, Z., Hu, S., Li, M., Zhang, H., Zhang, Y., Jin, H.: AdvCLIP: Downstream-agnostic adversarial examples in multimodal contrastive learning. In: *Proceedings of the 31st ACM International Conference on Multimedia*. pp. 6311–6320 (2023)
81. Zou, A., Wang, Z., Kolter, J.Z., Fredrikson, M.: Universal and transferable adversarial attacks on aligned language models. arXiv preprint arXiv:2307.15043 (2023)