DENSEPURE: UNDERSTANDING DIFFUSION MODELS TOWARDS ADVERSARIAL ROBUSTNESS

Anonymous authors
Paper under double-blind review

ABSTRACT

Diffusion models have been recently employed to improve certified robustness through the process of denoising. However, the theoretical understanding of why diffusion models are able to improve the certified robustness is still lacking, preventing from further improvement. In this study, we close this gap by analyzing the fundamental properties of diffusion models and establishing the conditions under which they can enhance certified robustness. This deeper understanding allows us to propose a new method DensePure, designed to improve the certified robustness of a pretrained model (i.e. classifier). Given an (adversarial) input, DensePure consists of multiple runs of denoising via the reverse process of the diffusion model (with different random seeds) to get multiple reversed samples, which are then passed through the classifier, followed by majority voting of inferred labels to make the final prediction. This design of using multiple runs of denoising is informed by our theoretical analysis of the conditional distribution of the reversed sample. Specifically, when the *data* density of a clean sample is high, its conditional density under the reverse process in a diffusion model is also high; thus sampling from the latter conditional distribution can purify the adversarial example and return the corresponding clean sample with a high probability. By using the highest density point in the conditional distribution as the reversed sample, we identify the robust region of a given instance under the diffusion model's reverse process. We show that this robust region is a union of multiple convex sets, and is potentially much larger than the robust regions identified in previous works. In practice, DensePure can approximate the label of the high density region in the conditional distribution so that it can enhance certified robustness. We conduct extensive experiments to demonstrate the effectiveness of DensePure by evaluating its certified robustness given a standard model via randomized smoothing. We show that DensePure is consistently better than existing methods on ImageNet, with 7% improvement on average.

1 INTRODUCTION

Diffusion models have been shown to be a powerful image generation tool (Ho et al., 2020; Song et al., 2021b) owing to their iterative diffusion and denoising processes. These models have achieved state-of-the-art performance on sample quality (Dhariwal & Nichol, 2021; Vahdat et al., 2021) as well as effective mode coverage (Song et al., 2021a). A diffusion model usually consists of two processes: (i) a forward diffusion process that converts data to noise by gradually adding noise to the input, and (ii) a reverse generative process that starts from noise and generates data by denoising one step at a time (Song et al., 2021b).

Given the natural denoising property of diffusion models, *empirical* studies have leveraged them to perform adversarial purification (Nie et al., 2022; Wu et al., 2022; Carlini et al., 2022). For instance, Nie et al. (2022) introduce a diffusion model based purification model *DiffPure*. They *empirically* show that by carefully choosing the amount of Gaussian noises added during the diffusion process, adversarial perturbations can be removed while preserving the true label semantics. Despite the significant empirical results, there is no provable guarantee of the achieved robustness. Carlini et al. (2022) instantiate the randomized smoothing approach with the diffusion model to offer a *provable guarantee* of model robustness against L_2-norm bounded adversarial example. However, they do

not provide a theoretical understanding of why and how the diffusion models contribute to such nontrivial certified robustness.

Our Approach. We theoretically analyze the fundamental properties of diffusion models to understand why and how it enhances certified robustness. This deeper understanding allows us to propose a new method DensePure to improve the certified robustness of any given classifier by more effectively using the diffusion model. An illustration of the DensePure framework is provided in Figure 1, where it consists of a pretrained diffusion model and a pretrained classifier. DensePure incorporates two steps: (i) using the reverse process of the diffusion model to obtain a sample of the posterior data distribution conditioned on the adversarial input; and (ii) repeating the reverse process multiple times with different random seeds to approximate the label of high density region in the conditional distribution via a majority vote. In particular, given an adversarial input, we repeatedly feed it into the reverse process of the diffusion model to get multiple reversed examples and feed them into the classifier to get their labels. We then apply the *majority vote* on the set of labels to get the final predicted label.

DensePure is inspired by our theoretical analysis, where we show that the diffusion model reverse process provides a conditional distribution of the reversed sample given an adversarial input, and sampling from this conditional distribution enhances the certified robustness. Specifically, we prove that when the data density of clean samples is high, it is a sufficient condition for the conditional density of the reversed samples to be also high. Therefore, in DensePure, samples from the conditional distribution can recover the ground-truth labels with a high probability.

For the convenience of understanding and rigorous analysis, we use the highest density point in the conditional distribution as the deterministic reversed sample for the classifier prediction. We show that the robust region for a given sample under the diffusion model's reverse process is the union of multiple convex sets, each surrounding a region around the ground-truth label. Compared with the robust region of previous work (Cohen et al., 2019), which only focuses on the neighborhood of *one* region with the ground-truth label, such union of multiple convex sets has the potential to provide a much larger robust region. Moreover, the characterization implies that the size of robust regions is affected by the relative density and the distance between data regions with the ground-truth label and those with other labels.

We conduct extensive experiments on ImageNet and CIFAR-10 datasets under different settings to evaluate the certifiable robustness of DensePure. In particular, we follow the setting from Carlini et al. (2022) and rely on randomized smoothing to certify robustness to adversarial perturbations bounded in the $\mathcal{L}_2$-norm. We show that DensePure achieves the new state-of-the-art *certified* robustness on the clean model without tuning any model parameters (off-the-shelf). On ImageNet, it achieves a consistently higher certified accuracy than the existing methods among every σ at every radius ϵ , 7% improvement on average.

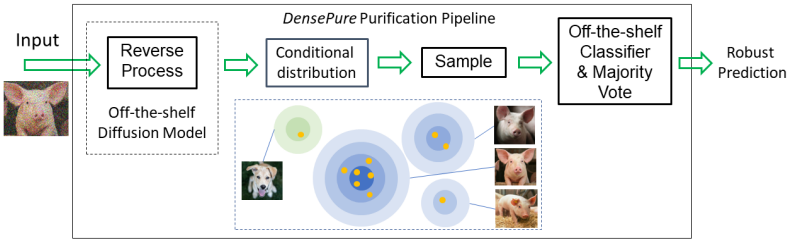

Figure 1: Pipeline of DensePure.

Technical Contributions. In this paper, we take the first step towards understanding the sufficient conditions of *adversarial purification* with diffusion models. We make contributions on both theoretical and empirical fronts: (1) We prove that under constrained data density property, an adversarial example can be recovered back to the original clean sample with high probability via the reverse process of a diffusion model. (2) In theory, we characterized the robust region for each point by further taking the highest density point in the conditional distribution generated by the reverse process as the reversed sample. (3) In practice, we proposed DensePure, which is a state-of-art adversarial purification pipeline directly leveraging the reverse process of a pre-trained diffusion model and label *majority vote*. (4) We demonstrated comparable performance of DensePure on CIFAR-10 and state-of-the-art performance on ImageNet.

2 PRELIMINARIES AND BACKGROUNDS

Continuous-Time Diffusion Model. The diffusion model has two components: the *diffusion process* followed by the *reverse process*. Given an input random variable $\mathbf{x}_0 \sim p$, the diffusion process adds isotropic Gaussian noises to the data so that the diffused random variable at time t is $\mathbf{x}_t = \sqrt{\alpha_t}(\mathbf{x}_0 + \boldsymbol{\epsilon}_t)$, s.t., $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \sigma_t^2 \boldsymbol{I})$, and $\sigma_t^2 = (1-\alpha_t)/\alpha_t$, and we denote $\mathbf{x}_t \sim p_t$. The forward diffusion process can also be defined by the stochastic differential equation

$$d\boldsymbol{x} = h(\boldsymbol{x}, t)dt + g(t)d\boldsymbol{w}, \tag{SDE}$$

where $\boldsymbol{x}_0 \sim p$, $h : \mathbb{R}^d \times \mathbb{R} \mapsto \mathbb{R}^d$ is the drift coefficient, $g : \mathbb{R} \mapsto \mathbb{R}$ is the diffusion coefficient, and $\boldsymbol{w}(t) \in \mathbb{R}^n$ is the standard Wiener process.

Under mild conditions B.1, the reverse process exists and removes the added noise by solving the reverse-time SDE (Anderson, 1982)

$$d\hat{\boldsymbol{x}} = [h(\hat{\boldsymbol{x}}, t) - g(t)^2 \nabla_{\hat{\boldsymbol{x}}} \log p_t(\hat{\boldsymbol{x}})]dt + g(t)d\overline{\boldsymbol{w}}, \tag{reverse-SDE}$$

where dt is an infinitesimal reverse time step, and $\overline{\boldsymbol{w}}(t)$ is a reverse-time standard Wiener process.

In our context, we use the conventions of VP-SDE (Song et al., 2021b) where $h(\boldsymbol{x}; t) := -\frac{1}{2}\gamma(t)x$ and $g(t) := \sqrt{\gamma(t)}$ with $\gamma(t)$ positive and continuous over $[0, 1]$, such that $x(t) = \sqrt{\alpha_t}x(0) + \sqrt{1-\alpha_t}\boldsymbol{\epsilon}$ where $\alpha_t = e^{-\int_0^t \gamma(s)ds}$ and $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$. We use $\{\mathbf{x}_t\}_{t\in[0,1]}$ and $\{\hat{\mathbf{x}}_t\}_{t\in[0,1]}$ to denote the diffusion process and the reverse process generated by SDE and reverse-SDE respectively, which follow the same distribution.

Discrete-Time Diffusion Model (or DDPM (Ho et al., 2020)). DDPM constructs a discrete Markov chain $\{\mathbf{x}_0, \mathbf{x}_1, \cdots, \mathbf{x}_i, \cdots, \mathbf{x}_N\}$ as the forward process for the training data $\mathbf{x}_0 \sim p$, such that $\mathbb{P}(\mathbf{x}_i|\mathbf{x}_{i-1}) = \mathcal{N}(\mathbf{x}_i; \sqrt{1-\beta_i}\mathbf{x}_{i-1}, \beta_i I)$, where $0 < \beta_1 < \beta_2 < \cdots < \beta_N < 1$ are predefined noise scales such that $\mathbf{x}_N$ approximates the Gaussian white noise. Denote $\overline{\alpha}_i = \prod_{i=1}^N (1-\beta_i)$, we have $\mathbb{P}(\mathbf{x}_i|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_i; \sqrt{\overline{\alpha}_i}\mathbf{x}_0, (1-\overline{\alpha}_i)\boldsymbol{I})$, i.e., $\mathbf{x}_t(\mathbf{x}_0, \epsilon) = \sqrt{\overline{\alpha}_i}\mathbf{x}_0 + (1-\overline{\alpha}_i)\boldsymbol{\epsilon}, \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$.

The reverse process of DDPM learns a reverse direction variational Markov chain $p_{\boldsymbol{\theta}}(\mathbf{x}_{i-1}|\mathbf{x}_i) = \mathcal{N}(\mathbf{x}_{i-1}; \boldsymbol{\mu}_{\boldsymbol{\theta}}(\mathbf{x}_i, i), \Sigma_{\boldsymbol{\theta}}(\mathbf{x}_i, i))$. Ho et al. (2020) defines $\boldsymbol{\epsilon}_{\boldsymbol{\theta}}$ as a function approximator to predict $\boldsymbol{\epsilon}$ from $\boldsymbol{x}_i$ such that $\boldsymbol{\mu}_{\boldsymbol{\theta}}(\mathbf{x}_i, i) = \frac{1}{\sqrt{1-\beta_i}}\left(\mathbf{x}_i - \frac{\beta_i}{\sqrt{1-\overline{\alpha}_i}}\boldsymbol{\epsilon}_{\boldsymbol{\theta}}(\mathbf{x}_i, i)\right)$. Then the reverse time samples are generated by $\hat{\mathbf{x}}_{i-1} = \frac{1}{\sqrt{1-\beta_i}}\left(\hat{\mathbf{x}}_i - \frac{\beta_i}{\sqrt{1-\overline{\alpha}_i}}\boldsymbol{\epsilon}_{\boldsymbol{\theta}^*}(\hat{\mathbf{x}}_i, i)\right) + \sqrt{\beta_i}\boldsymbol{\epsilon}, \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, I)$, and the optimal parameters $\boldsymbol{\theta}^*$ are obtained by solving $\boldsymbol{\theta}^* := \arg\min_{\boldsymbol{\theta}} \mathbb{E}_{\mathbf{x}_0, \boldsymbol{\epsilon}}\left[||\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_{\boldsymbol{\theta}}(\sqrt{\overline{\alpha}_i}\mathbf{x}_0 + (1-\overline{\alpha}_i), i)||_2^2\right]$.

Randomized Smoothing. Randomized smoothing is used to certify the robustness of a given classifier against L_2-norm based perturbation. It transfers the classifier f to a smooth version $g(\boldsymbol{x}) = \arg\max_c \mathbb{P}_{\boldsymbol{\epsilon}\sim\mathcal{N}(\mathbf{0}, \sigma^2\boldsymbol{I})}(f(\boldsymbol{x} + \boldsymbol{\epsilon}) = c)$, where g is the smooth classifier and σ is a hyper-parameter of the smooth classifier g, which controls the trade-off between robustness and accuracy. Cohen et al. (2019) shows that $g(x)$ induces the certifiable robustness for $\boldsymbol{x}$ under the L_2-norm with radius R, where $R = \frac{\sigma}{2}\left(\Phi^{-1}(p_A) - \Phi^{-1}(p_B)\right)$; p_A and p_B are probability of the most probable class and "runner-up" class respectively; Φ is the inverse of the standard Gaussian CDF. The p_A and p_B can be estimated with arbitrarily high confidence via Monte Carlo method (Cohen et al., 2019).

3 THEORETICAL ANALYSIS

In this section, we theoretically analyze why and how the diffusion model can enhance the robustness of a given classifier. We will analyze directly on SDE and reverse-SDE as they generate the same stochastic processes $\{\mathbf{x}_t\}_{t\in[0,T]}$ and the literature works establish an approximation on reverse-SDE (Song et al., 2021b; Ho et al., 2020).

We first show that given a diffusion model, solving reverse-SDE will generate a conditional distribution based on the scaled adversarial sample, which will have high density on data region with high *data* density and near to the adversarial sample in Theorem 3.1. See detailed conditions in B.1.

Theorem 3.1. *Under conditions B.1, solving equation reverse-SDE starting from time t and sample* $\boldsymbol{x}_{a,t} = \sqrt{\alpha_t}\boldsymbol{x}_a$ *will generate a reversed random variable* $\hat{\mathbf{x}}_0$ *with density* $\mathbb{P}(\hat{\mathbf{x}}_0 = \boldsymbol{x}|\hat{\mathbf{x}}_t = \boldsymbol{x}_{a,t}) \propto$

$p(\boldsymbol{x}) \cdot \frac{1}{\sqrt{(2\pi\sigma_t^2)^n}} \exp\left(\frac{-||\boldsymbol{x}-\boldsymbol{x}_a||_2^2}{2\sigma_t^2}\right)$, *where p is the data distribution, $\sigma_t^2 = \frac{1-\alpha_t}{\alpha_t}$ is the variance of Gaussian noise added at time t in the diffusion process.*

Proof. (sketch) Under conditions B.1, we know $\{\mathbf{x}_t\}_{t\in[0,1]}$ and $\{\hat{\mathbf{x}}_t\}_{t\in[0,1]}$ follow the same distribution, and then the rest proof follows Bayes' Rule. □

Please see the full proofs of this and the following theorems in Appendix B.2.

Remark 1. *Note that $\mathbb{P}(\hat{\mathbf{x}}_0 = \boldsymbol{x}|\hat{\mathbf{x}}_t = \boldsymbol{x}_{a,t}) > 0$ if and only if $p(\boldsymbol{x}) > 0$, thus the generated reverse sample will be on the data region where we train classifiers.*

In Theorem 3.1, the conditional density $\mathbb{P}(\hat{\mathbf{x}}_0 = \boldsymbol{x}|\hat{\mathbf{x}}_t = \boldsymbol{x}_{a,t})$ is high if both $p(\boldsymbol{x})$ and the Gaussian term have high values, i.e., $\boldsymbol{x}$ has high *data* density and is close to the adversarial sample $\boldsymbol{x}_a$. The latter condition is reasonable since adversarial perturbations are typically bounded due to budget constraints. So the above argument implies that a reversed sample will more likely to have the ground-truth label if data region with the ground-truth label has high *data* density. For the sake of theoretical analysis and understanding, we take the point with highest conditional density $\mathbb{P}(\hat{\mathbf{x}}_0 = \boldsymbol{x}|\hat{\mathbf{x}}_t = \boldsymbol{x}_{a,t})$ as the reversed sample, defined as $\mathcal{P}(\boldsymbol{x}_a; t) := \arg\max_{\boldsymbol{x}} \mathbb{P}(\hat{\mathbf{x}}_0 = \boldsymbol{x}|\hat{\mathbf{x}}_t = \boldsymbol{x}_{a,t})$. $\mathcal{P}(\boldsymbol{x}_a; t)$ is a representative of the high density data region in the conditional distribution and $\mathcal{P}(\cdot; t)$ is a deterministic purification model. In the following, we characterize the robust region for data region with ground-truth label under $\mathbb{P}(\cdot; t)$. The robust region and robust radius for a general deterministic purification model given a classifier are defined below.

Definition 3.2 (Robust Region and Robust Radius)**.** *Given a classifier f and a point $\boldsymbol{x}_0$, let $\mathcal{G}(\boldsymbol{x}_0) := \{\boldsymbol{x} : f(\boldsymbol{x}) = f(\boldsymbol{x}_0)\}$ be the data region where samples have the same label as $\boldsymbol{x}_0$. Then given a deterministic purification model $\mathcal{P}(\cdot\,; \psi)$ with parameter ψ, we define the robust region of $\mathcal{G}(\boldsymbol{x}_0)$ under $\mathcal{P}$ and f as $\mathcal{D}_{\mathcal{P}}^f(\mathcal{G}(\boldsymbol{x}_0); \psi) := \{\boldsymbol{x} : f(\mathcal{P}(\boldsymbol{x}; \psi)) = f(\boldsymbol{x}_0)\}$, i.e., the set of $\boldsymbol{x}$ such that purified sample $\mathcal{P}(\boldsymbol{x}; \psi)$ has the same label as $\boldsymbol{x}_0$ under f. Further, we define the robust radius of $\boldsymbol{x}_0$ as $r_{\mathcal{P}}^f(\boldsymbol{x}_0; \psi) := \max\left\{r : \boldsymbol{x}_0 + ru \in \mathcal{D}_{\mathcal{P}}^f(\boldsymbol{x}_0; \psi)\ ,\ \forall ||u||_2 \leq 1\right\}$, i.e., the radius of maximum inclined ball of $\mathcal{D}_{\mathcal{P}}^f(\boldsymbol{x}_0; \psi)$ centered around $\boldsymbol{x}_0$. We will omit $\mathcal{P}$ and f when it is clear from the context and write $\mathcal{D}(\mathcal{G}(\boldsymbol{x}_0); \psi)$ and $r(\boldsymbol{x}_0; \psi)$ instead.*

Remark 2. *In Definition 3.2, the robust region (resp. radius) is defined for each class (resp. point). When using the point with highest $\mathbb{P}(\hat{\mathbf{x}}_0 = \boldsymbol{x}|\hat{\mathbf{x}}_t = \boldsymbol{x}_{a,t})$ as the reversed sample, $\psi := t$.*

Now given a sample $\boldsymbol{x}_0$ with ground-truth label, we are ready to characterize the robust region $\mathcal{D}(\mathcal{G}(\boldsymbol{x}_0); \psi)$ under purification model $\mathcal{P}(\cdot; t)$ and classifier f. Intuitively, if the adversarial sample $\boldsymbol{x}_a$ is near to $\boldsymbol{x}_0$ (in Euclidean distance), $\boldsymbol{x}_a$ keeps the same label semantics of $\boldsymbol{x}_0$ and so as the purified sample $\mathcal{P}(\boldsymbol{x}_a; t)$, which implies that $f(\mathcal{P}(\boldsymbol{x}_a; \psi)) = f(\boldsymbol{x}_0)$. However, the condition that $\boldsymbol{x}_a$ is near to $\boldsymbol{x}_0$ is sufficient but not necessary since we can still achieve $f(\mathcal{P}(\boldsymbol{x}_a; \psi)) = f(\boldsymbol{x}_0)$ if $\boldsymbol{x}_a$ is near to any sample $\tilde{\boldsymbol{x}}_0$ with $f(\mathcal{P}(\tilde{\boldsymbol{x}}_a; \psi)) = f(\boldsymbol{x}_0)$. In the following, we will show that the robust region $\mathcal{D}(\mathcal{G}(\boldsymbol{x}_0); \psi)$ is the union of the convex robust sub-regions surrounding every $\tilde{\boldsymbol{x}}_0$ with the same label as $\boldsymbol{x}_0$. The following theorem characterizes the convex robust sub-region and robust region respectively.

Theorem 3.3. *Under conditions B.1 and classifier f, let $\boldsymbol{x}_0$ be the sample with ground-truth label and $\boldsymbol{x}_a$ be the adversarial sample, then (i) the purified sample $\mathcal{P}(\boldsymbol{x}_a; t)$ will have the ground-truth label if $\boldsymbol{x}_a$ falls into the following convex set,*

$$\mathcal{D}_{sub}(\boldsymbol{x}_0; t) := \bigcap_{\{\boldsymbol{x}_0' : f(\boldsymbol{x}_0') \neq f(\boldsymbol{x}_0)\}} \left\{\boldsymbol{x}_a : (\boldsymbol{x}_a - \boldsymbol{x}_0)^\top (\boldsymbol{x}_0' - \boldsymbol{x}_0) < \sigma_t^2 \log\left(\frac{p(\boldsymbol{x}_0)}{p(\boldsymbol{x}_0')}\right) + \frac{||\boldsymbol{x}_0' - \boldsymbol{x}_0||_2^2}{2}\right\},$$

and further, (ii) the purified sample $\mathcal{P}(\boldsymbol{x}_a; t)$ will have the ground-truth label if and only if $\boldsymbol{x}_a$ falls into the following set, $\mathcal{D}(\mathcal{G}(\boldsymbol{x}_0); t) := \bigcup_{\tilde{\boldsymbol{x}}_0 : f(\tilde{\boldsymbol{x}}_0) = f(\boldsymbol{x}_0)} \mathcal{D}_{sub}(\tilde{\boldsymbol{x}}_0; t)$. In other words, $\mathcal{D}(\mathcal{G}(\boldsymbol{x}_0); t)$ is the robust region for data region $\mathcal{G}(\boldsymbol{x}_0)$ under $\mathcal{P}(\cdot; t)$ and f.

Proof. (sketch) (i). Each convex half-space defined by the inequality corresponds to a $\boldsymbol{x}_0'$ such that $f(\boldsymbol{x}_0') \neq f(\boldsymbol{x}_0)$ where $\boldsymbol{x}_a$ within satisfies $\mathbb{P}(\hat{\mathbf{x}}_0 = \boldsymbol{x}_0|\hat{\mathbf{x}}_t = \boldsymbol{x}_{a,t}) > \mathbb{P}(\hat{\mathbf{x}}_0 = \boldsymbol{x}_0' \mid \hat{\mathbf{x}}_t = \boldsymbol{x}_{a,t})$. This implies that $\mathcal{P}(\boldsymbol{x}_a; t) \neq \boldsymbol{x}_0'$ and $f(\mathcal{P}(\boldsymbol{x}_a; \psi)) = f(\boldsymbol{x}_0)$. The convexity is due to that the intersection

of convex sets is convex. (ii). The "if" follows directly from (i). The "only if" holds because if $\boldsymbol{x}_a \notin \mathcal{D}(\mathcal{G}(\boldsymbol{x}_0); t)$, then exists $\tilde{\boldsymbol{x}}_1$ such that $f(\tilde{\boldsymbol{x}}_1) \neq f(\boldsymbol{x}_0)$ and $\mathbb{P}(\hat{\mathbf{x}}_0 = \tilde{\boldsymbol{x}}_1 | \hat{\mathbf{x}}_t = \boldsymbol{x}_{a,t}) > \mathbb{P}(\hat{\mathbf{x}}_0 = \tilde{\boldsymbol{x}}_0 | \hat{\mathbf{x}}_t = \boldsymbol{x}_{a,t}), \forall \tilde{\boldsymbol{x}}_0$ s.t. $f(\tilde{\boldsymbol{x}}_0) = f(\boldsymbol{x}_0)$, and thus $f(\mathcal{P}(\boldsymbol{x}_a; \psi)) \neq f(\boldsymbol{x}_0)$. □

Remark 3. *Theorem 3.3 implies that when data region $\mathcal{G}(\boldsymbol{x}_0)$ has higher data density and larger distances to data regions with other labels, it tends to have larger robust region and points in data region tends to have larger radius. Since adversarial attack typically has small magnitude, with large robust region, the adversarial sample can be recovered to the clean sample with a high probability.*

In the literature, people focus more on the robust radius (lower bound) $r(\mathcal{G}(\boldsymbol{x}_0); t)$ (Cohen et al., 2019; Carlini et al., 2022), which can be obtained by finding the maximum inclined ball inside $\mathcal{D}(\mathcal{G}(\boldsymbol{x}_0); t)$ centering $\boldsymbol{x}_0$. Note that although $\mathcal{D}_{\text{sub}}(\boldsymbol{x}_0; t)$ is convex, $\mathcal{D}(\mathcal{G}(\boldsymbol{x}_0); t)$ is generally not. Therefore, finding $r(\mathcal{G}(\boldsymbol{x}_0); t)$ is a non-convex optimization problem. In particular, it can be formulated into a disjunctive optimization problem with integer indicator variables, which is typically NP-hard to solve. One alternative could be finding the maximum inclined ball in $\mathcal{D}_{\text{sub}}(\boldsymbol{x}_0; t)$, which can be formulated into a convex optimization problem whose optimal value provides a lower bound for $r(\mathcal{G}(\boldsymbol{x}_0); t)$. However, $\mathcal{D}(\mathcal{G}(\boldsymbol{x}_0); t)$ has the potential to provide much larger robustness radius because it might connect different convex robust sub-regions into one, as shown in Figure 2.

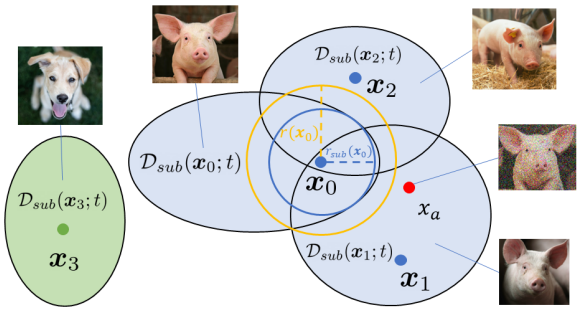

Figure 2: An illustration of the robust region $\mathcal{D}(\boldsymbol{x}_0; t) = \bigcup_{i=1}^{3} \mathcal{D}_{sub}(\boldsymbol{x}_i; t)$, where $\boldsymbol{x}_0, \boldsymbol{x}_1, \boldsymbol{x}_2$ are samples with ground-truth label and $\boldsymbol{x}_3$ is a sample with another label. $\boldsymbol{x}_a = \boldsymbol{x}_0 + \boldsymbol{\epsilon}_a$ is an adversarial sample such that $\mathcal{P}(\boldsymbol{x}_a; t) = \boldsymbol{x}_1 \neq \boldsymbol{x}_0$ and thus the classification is correct but $\boldsymbol{x}_a$ is not reversed back to $\boldsymbol{x}_0$. $r_{sub}(\boldsymbol{x}_0) < r(\boldsymbol{x}_0)$ shows our claim that the union leads to a larger robust radius.

In practice, we cannot guarantee to establish an exact reverse process like reverse-SDE but instead try to establish an approximate reverse process to mimic the exact one. As long as the approximate reverse process is close enough to the exact reverse process, they will generate close enough conditional distributions based on the adversarial sample. Then the density and locations of the data regions in two conditional distributions will not differ much and so is the robust region for each data region. We take the score-based diffusion model in Song et al. (2021b) for an example and demonstrate Theorem 3.4 to bound the KL-divergnece between conditional distributions generated by reverse-SDE and score-based diffusion model. Ho et al. (2020) showed that using variational inference to fit DDPM is equivalent to optimizing an objective resembling score-based diffusion model with a specific weighting scheme, so the results can be extended to DDPM.

Theorem 3.4. *Under score-based diffusion model Song et al. (2021b) and conditions B.1, we have $D_{KL}(\mathbb{P}(\hat{\mathbf{x}}_0 = \boldsymbol{x} \mid \hat{\mathbf{x}}_t = \boldsymbol{x}_{a,t}) \| \mathbb{P}(\mathbf{x}_0^\theta = \boldsymbol{x} \mid \mathbf{x}_t^\theta = \boldsymbol{x}_{a,t})) = \mathcal{J}_{\text{SM}}(\theta, t; \lambda(\cdot))$, where $\{\hat{\boldsymbol{x}}_\tau\}_{\tau \in [0,t]}$ and $\{\boldsymbol{x}_\tau^\theta\}_{\tau \in [0,t]}$ are stochastic processes generated by reverse-SDE and score-based diffusion model respectively, $\mathcal{J}_{\text{SM}}(\theta, t; \lambda(\cdot)) := \frac{1}{2} \int_0^t \mathbb{E}_{p_\tau(\mathbf{x})} \left[\lambda(\tau) \left\| \nabla_{\mathbf{x}} \log p_\tau(\mathbf{x}) - \boldsymbol{s}_\theta(\mathbf{x}, \tau) \right\|_2^2 \right] \mathrm{d}\tau$, $\boldsymbol{s}_\theta(\mathbf{x}, \tau)$ is the score function to approximate $\nabla_{\mathbf{x}} \log p_\tau(\mathbf{x})$, and $\lambda : \mathbb{R} \to \mathbb{R}$ is any weighting scheme used in the training score-based diffusion models.*

Proof. (sketch) Let $\boldsymbol{\mu}_t$ and $\boldsymbol{\nu}_t$ be the path measure for reverse processes $\{\hat{\mathbf{x}}_\tau\}_{\tau \in [0,t]}$ and $\{\mathbf{x}_\tau^\theta\}_{\tau \in [0,t]}$ respectively based on the $\boldsymbol{x}_{a,t}$. Under conditions B.1, $\boldsymbol{\mu}_t$ and $\boldsymbol{\nu}_t$ are uniquely defined and the KL-divergence can be computed via the Girsanov theorem Oksendal (2013). □

Remark 4. *Theorem 3.4 shows that if the training loss is smaller, the conditional distributions generated by reverse-SDE and score-based diffusion model are closer, and are the same if the training*

loss is zero. Furthermore, by the Pinsker's inequality, the total variation (a distance metric) is upper bounded by $D_{TV}(\mathbb{P}(\hat{\mathbf{x}}_0 = \boldsymbol{x} \mid \hat{\mathbf{x}}_t = \boldsymbol{x}_{a,t}) \| \mathbb{P}(\mathbf{x}_0^\theta = \boldsymbol{x} \mid \mathbf{x}_t^\theta = \boldsymbol{x}_{a,t})) \leq \sqrt{\frac{1}{2}\mathcal{J}_{\mathrm{SM}}(\theta, t; \lambda(\cdot))}$.

4 DENSEPURE

Inspired by the theoretical analysis, we introduce DensePure and show how to calculate its certified robustness radius via the randomized smoothing algorithm.

Framework. Our framework, DensePure, consists of two components: (1) an off-the-shelf diffusion model with reverse process $\mathbf{rev}$ and (2) an off-the-shelf base classifier f.

The pipeline of DensePure is shown in Figure 1. Given an input $\boldsymbol{x}$, we feed it into the reverse process $\mathbf{rev}$ of the diffusion model to get the reversed sample $\mathbf{rev}(\boldsymbol{x})$ and then repeat the above process K times to get K reversed samples $\{\mathbf{rev}(\boldsymbol{x})_1, \cdots, \mathbf{rev}(\boldsymbol{x})_K\}$. We feed the above K reversed samples into the classifier to get the corresponding prediction $\{f(\mathbf{rev}(\boldsymbol{x})_1), \cdots, f(\mathbf{rev}(\boldsymbol{x})_K)\}$ and then apply the *majority vote*, termed **MV**, on these predictions to get the final predicted label $\hat{y} = \mathbf{MV}(\{f(\mathbf{rev}(\boldsymbol{x})_1), \cdots, f(\mathbf{rev}(\boldsymbol{x})_K)\}) = \arg\max_c \sum_{i=1}^{K} \mathbf{1}\{f(\mathbf{rev}(\boldsymbol{x})_i) = c\}$.

Certified Robustness of DensePure with Randomized Smoothing.

In this paragraph, we will illustrate the algorithm to calculate certified robustness of DensePure via RS, which offers robustness guarantees for a model under a L_2-norm ball.

In particular, we follow the similar setting of Carlini et al. (2022) which uses a DDPM-based diffusion model. The overall algorithm contains three steps:

(1) Our framework estimates n, the number of steps used for the reverse process of DDPM-based diffusion model. Since Randomized Smoothing (Cohen et al., 2019) adds Gaussian noise $\boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \boldsymbol{I})$, to data input $\boldsymbol{x}$ to get the randomized data input, $\boldsymbol{x}_{\text{rs}} = \boldsymbol{x} + \boldsymbol{\epsilon}$, we map between the noise required by the randomized example $\boldsymbol{x}_{\text{rs}}$ and the noise required by the diffused data $\boldsymbol{x}_n$ (i.e., $\boldsymbol{x}_n \sim \mathcal{N}(\boldsymbol{x}_n; \sqrt{\overline{\alpha}_n}\boldsymbol{x}_0, (1 - \overline{\alpha}_n)\boldsymbol{I})$) with n step diffusion processing so that $\overline{\alpha}_n = \frac{1}{1+\sigma^2}$. In this way, we can compute the corresponding timestep n, where $n = \arg\min_s \{|\overline{\alpha}_s - \frac{1}{1+\sigma^2}| \mid s \in [N]\}$.

(2). Given the above calculated timestep n, we scale $\boldsymbol{x}_{rs}$ with $\sqrt{\overline{\alpha}_n}$ to obtain the scaled randomized smoothing sample $\sqrt{\overline{\alpha}_n}\boldsymbol{x}_{rs}$. Then we feed $\sqrt{\overline{\alpha}_n}\boldsymbol{x}_{rs}$ into the reverse process of the diffusion model by K-times to get the reversed sample set $\{\hat{\boldsymbol{x}}_0^1, \hat{\boldsymbol{x}}_0^2, \cdots, \hat{\boldsymbol{x}}_0^i, \cdots, \hat{\boldsymbol{x}}_0^K\}$.

(3). We feed the obtained reversed sample set into a standard *off-the-shelf* classifier f to get the corresponding predicted labels $\{f(\hat{\boldsymbol{x}}_0^1), f(\hat{\boldsymbol{x}}_0^2), \ldots, f(\hat{\boldsymbol{x}}_0^i), \ldots, f(\hat{\boldsymbol{x}}_0^K)\}$, and apply *majority vote*, denoted $\mathbf{MV}(\cdots)$, on these predicted labels to get the final label for $\boldsymbol{x}_{rs}$.

Fast Sampling. To calculate the reversed sample, the standard reverse process of DDPM-based models require repeatedly applying a "single-step" operation n times to get the reversed sample $\hat{\boldsymbol{x}}_0$ (i.e., $\hat{\boldsymbol{x}}_0 = \mathbf{Reverse}(\cdots \mathbf{Reverse}(\cdots \mathbf{Reverse}(\mathbf{Reverse}(\sqrt{\overline{\alpha}_n}\boldsymbol{x}_{rs}; n); n-1); \cdots; i); \cdots 1)$). Here $\hat{\boldsymbol{x}}_{i-1} = \mathbf{Reverse}(\hat{\boldsymbol{x}}_i; i)$ is equivalent to sample $\hat{\boldsymbol{x}}_{i-1}$ from $\mathcal{N}(\hat{\boldsymbol{x}}_{i-1}; \boldsymbol{\mu}_{\boldsymbol{\theta}}(\hat{\boldsymbol{x}}_i, i), \boldsymbol{\Sigma}_{\boldsymbol{\theta}}(\hat{\boldsymbol{x}}_i, i))$, where $\boldsymbol{\mu}_{\boldsymbol{\theta}}(\hat{\boldsymbol{x}}_i, i) = \frac{1}{\sqrt{1-\beta_i}}\left(\hat{\boldsymbol{x}}_i - \frac{\beta_i}{\sqrt{1-\overline{\alpha}_i}}\boldsymbol{\epsilon}_{\boldsymbol{\theta}}(\hat{\boldsymbol{x}}_i, i)\right)$ and $\boldsymbol{\Sigma}_{\boldsymbol{\theta}} := \exp(v \log \beta_i + (1-v) \log \widetilde{\beta}_i)$. Here v is a parameter learned by DDPM and $\widetilde{\beta}_i = \frac{1-\overline{\alpha}_{i-1}}{1-\overline{\alpha}_i}$.

To reduce the time complexity, we use the uniform sub-sampling strategy from Nichol & Dhariwal (2021). We uniformly sample a subsequence with size b from the original N-step the reverse process. Note that Carlini et al. (2022) set $b = 1$ for the "one-shot" sampling, in this way, $\hat{\boldsymbol{x}}_0 = \frac{1}{\sqrt{\overline{\alpha}_n}}(\boldsymbol{x}_n - \sqrt{1-\overline{\alpha}_n}\boldsymbol{\epsilon}_{\boldsymbol{\theta}}(\sqrt{\overline{\alpha}_n}\boldsymbol{x}_{rs}, n))$ is a deterministic value so that the reverse process does not obtain a posterior data distribution conditioned on the input. Instead, we can tune the number of the sub-sampled DDPM steps to be larger than one ($b > 1$) to sample from a posterior data distribution conditioned on the input. The details about the fast sampling are shown in appendix C.2.

5 EXPERIMENTS

In this section, we use DensePure to evaluate certified robustness on two standard datasets, CIFAR-10 (Krizhevsky et al., 2009) and ImageNet (Deng et al., 2009).

Method	Off-the-shelf	Certified Accuracy at ϵ(%) CIFAR-10 0.25	0.5	0.75	1.0	ImageNet 0.5	1.0	1.5	2.0	3.0
PixelDP (Lecuyer et al., 2019)	✗	$^{(71.0)}$22.0	$^{(44.0)}$2.0	-	-	$^{(33.0)}$16.0	-	-	-	-
RS (Cohen et al., 2019)	✗	$^{(75.0)}$61.0	$^{(75.0)}$43.0	$^{(65.0)}$32.0	$^{(65.0)}$23.0	$^{(67.0)}$49.0	$^{(57.0)}$37.0	$^{(57.0)}$29.0	$^{(44.0)}$19.0	$^{(44.0)}$12.0
SmoothAdv (Salman et al., 2019a)	✗	$^{(82.0)}$68.0	$^{(76.0)}$54.0	$^{(68.0)}$41.0	$^{(64.0)}$32.0	$^{(63.0)}$54.0	$^{(56.0)}$42.0	$^{(56.0)}$34.0	$^{(41.0)}$26.0	$^{(41.0)}$18.0
Consistency (Jeong & Shin, 2020)	✗	$^{(77.8)}$68.8	$^{(75.8)}$58.1	$^{(72.9)}$48.5	$^{(52.3)}$37.8	$^{(55.0)}$50.0	$^{(55.0)}$44.0	$^{(55.0)}$34.0	$^{(41.0)}$24.0	$^{(41.0)}$17.0
MACER (Zhai et al., 2020)	✗	$^{(81.0)}$71.0	$^{(81.0)}$59.0	$^{(66.0)}$46.0	$^{(66.0)}$38.0	$^{(68.0)}$57.0	$^{(64.0)}$43.0	$^{(64.0)}$31.0	$^{(48.0)}$25.0	$^{(48.0)}$14.0
Boosting (Horváth et al., 2021)	✗	$^{(83.4)}$70.6	$^{(76.8)}$60.4	$^{(71.6)}$**52.4**	$^{(73.0)}$**38.8**	$^{(65.6)}$57.0	$^{(57.0)}$44.6	$^{(57.0)}$38.4	$^{(44.6)}$28.6	$^{(38.6)}$21.2
SmoothMix (Jeong et al., 2021)	✓	$^{(77.1)}$67.9	$^{(77.1)}$57.9	$^{(74.2)}$47.7	$^{(61.8)}$37.2	$^{(55.0)}$50.0	$^{(55.0)}$43.0	$^{(55.0)}$38.0	$^{(40.0)}$26.0	$^{(40.0)}$17.0
Denoised (Salman et al., 2020)	✓	$^{(72.0)}$56.0	$^{(62.0)}$41.0	$^{(62.0)}$28.0	$^{(44.0)}$19.0	$^{(60.0)}$33.0	$^{(38.0)}$14.0	$^{(38.0)}$6.0	-	-
Lee (Lee, 2021)	✓	60.0	42.0	28.0	19.0	41.0	24.0	11.0	-	-
Carlini (Carlini et al., 2022)	✓	$^{(88.0)}$73.8	$^{(88.0)}$56.2	$^{(88.0)}$41.6	$^{(74.2)}$31.0	$^{(77.0)}$71.0	$^{(74.0)}$54.0	$^{(74.0)}$46.0	$^{(59.0)}$29.0	$^{(59.0)}$22.0
Ours	✓	$^{(87.6)}$**76.6**	$^{(87.6)}$**64.6**	$^{(87.6)}$50.4	$^{(73.6)}$37.4	$^{(80.0)}$**76.0**	$^{(75.0)}$**62.0**	$^{(75.0)}$**49.0**	$^{(61.0)}$**37.0**	$^{(61.0)}$**26.0**

Table 1: Certified accuracy compared with existing works. The certified accuracy at $\epsilon = 0$ for each model is in the parentheses. The certified accuracy for each cell is from the respective papers except Carlini et al. (2022). Our diffusion model and classifier are the same as Carlini et al. (2022), where the off-the-shelf classifier uses ViT-based architectures trained on a large dataset (ImageNet-22k).

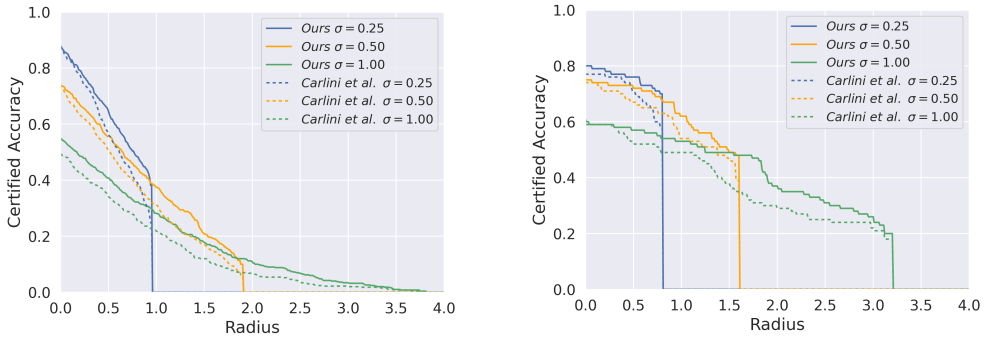

CIFAR-10 ImageNet

Figure 3: Comparing our method vs Carlini et al. (2022) on CIFAR-10 and ImageNet. The lines represent the certified accuracy with different L_2 perturbation bound with different Gaussian noise $\sigma \in \{0.25, 0.50, 1.00\}$.

Experimental settings We follow the experimental setting from Carlini et al. (2022). Specifically, for CIFAR-10, we use the 50-M unconditional improved diffusion model from Nichol & Dhariwal (2021) as the diffusion model. We select ViT-B/16 model Dosovitskiy et al. (2020) pretrained on ImageNet-21k and finetuned on CIFAR-10 as the classifier, which could achieve 97.9% accuracy on CIFAR-10. For ImageNet, we use the unconditional 256×256 guided diffusion model from Dhariwal & Nichol (2021) as the diffusion model and pretrained BEiT large model (Bao et al., 2021) trained on ImageNet-21k as the classifier, which could achieve 88.6% top-1 accuracy on validation set of ImageNet-1k. We select three different noise levels $\sigma \in \{0.25, 0.5, 1.0\}$ for certification. For the parameters of DensePure , we set $K = 40$ and b = 10 except the results in ablation study. The details about the baselines are in the appendix.

5.1 Main Results

We compare our results with other baselines. The results are shown in Table 1.

For CIFAR-10, comparing with the models which are *carefully* trained with randomized smoothing techniques in an end-to-end manner (i.e., w/o off-the-shelf classifier), we observe that our method with the standard off-the-shelf classifier outperforms them at smaller $\epsilon = \{0.25, 0.5\}$ on both CIFAR-10 and ImageNet datasets while achieves comparable performance at larger $\epsilon = \{0.75, 1.0\}$. Comparing with the non-diffusion model based methods with off-the-shelf classifier (i.e., Denoised (Salman et al., 2020) and Lee (Lee, 2021)), both our method and Carlini et al. (2022) are significantly better than them. These results verify the non-trivial adversarial robustness improvements introduced from the diffusion model. For ImageNet, our method is consistently better than all priors with a large margin.

Since both Carlini et al. (2022) and DensePure use the diffusion model, to better understand the importance of our design, that approximates the label of the high density region in the conditional distribution, we compare DensePure with Carlini et al. (2022) in a more fine-grained manner.

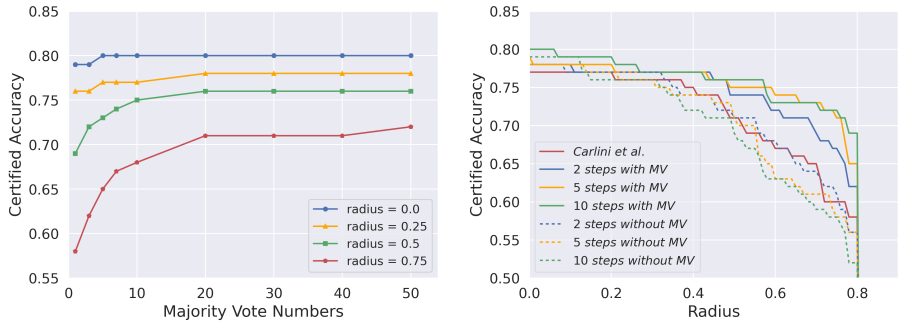

Figure 4: Ablation study on ImageNet. The left image shows the certified accuracy among different vote numbers with different radius $\epsilon \in \{0.0, 0.25, 0.5, 0.75\}$. The right image shows the certified accuracy with different fast sampling steps b.

We show detailed certified robustness of the model among different σ at different radius for CIFAR-10 in Figure 3-left and for ImageNet in Figure 3-right. We also present our results of certified accuracy at different ϵ in Appendix D.3. From these results, we find that our method is still consistently better at most ϵ (except $\epsilon = 0$) among different σ. The performance margin between ours and Carlini et al. (2022) will become even larger with a large ϵ. These results further indicate that although the diffusion model improves model robustness, leveraging the posterior data distribution conditioned on the input instance (like DensePure) via reverse process instead of using single sample ((Carlini et al., 2022)) is the key for better robustness. Additionally, we use the off-the-shelf classifiers, which are the VIT-based architectures trained a larger dataset. In the later ablation study section, we select the CNN-based architecture wide-ResNet trained on standard dataset from scratch. Our method still achieves non-trivial robustness. Further, our experiments in Appendix D.7 shows that removing the diffusion model from DensePure deteriorates the performance and thus our design is non-trivial.

5.2 Ablation Study

Voting samples (K) We first show how K affects the certified accuracy. For efficiency, we select $b = 10$. We conduct experiments for both datasets. We show the certified accuracy among different r at $\sigma = 0.25$ in Figure 4. The results for $\sigma = 0.5, 1.0$ and CIFAR-10 are shown in the Appendix D.4. Comparing with the baseline (Carlini et al., 2022), we find that a larger majority vote number leads to a better certified accuracy. It verifies that DensePure indeed benefits the adversarial robustness and making a good approximation of the label with high density region requires a large number of voting samples. We find that our certified accuracy will almost converge at $r = 40$. Thus, we set $r = 40$ for our experiments. The results with other σ show the similar tendency. To further improve the time efficency, we can use K-Consensus (Horváth et al., 2021). It accelerates the majority vote process by $45\% \sim 60\%$ with a negligible performance drop. The experimental details and results are in Appendix D.8.

Fast sampling steps (b) To investigate the role of b, we conduct additional experiments with $b \in \{2, 5\}$ at $\sigma = 0.25$. The results on ImageNet are shown in Figure 4 and results for $\sigma = 0.5, 1.0$ and CIFAR-10 are shown in the Appendix D.5. By observing results *with* majority vote, we find that a larger b can lead to a better certified accuracy since a larger b generates images with higher quality. By observing results *without* majority vote, the results show opposite conclusions where a larger b leads to a lower certified accuracy, which contradicts to our intuition. We guess the potential reason is that though more sampling steps can normally lead to better image recovery quality, it also brings more randomness, increasing the probability that the reversed image locates into a data region with the wrong label. These results further verify that majority vote is necessary for a better performance.

Different architectures One advantage of DensePure is to use the off-the-shelf classifier so that it can plug in any classifier. We choose Convolutional neural network (CNN)-based architectures: Wide-ResNet28-10 (Zagoruyko & Komodakis, 2016) for CIFAR-10 with 95.1% accuracy and Wide-ResNet50-2 for ImageNet with 81.5% top-1 accuracy, at $\sigma = 0.25$. The results are shown in Table 2 and Figure E in Appendix D.6. Results for more model architectures and σ of ImageNet are also shown in Appendix D.6. We show that our method can enhance the certified robustness of any given classifier trained on the original data distribution. Noticeably, although the performance of CNN-

Datasets	Methods	Model	Certified Accuracy at $\epsilon(\%)$ 0.0	0.25	0.5	0.75	Model	0.0	0.25	0.5	0.75
CIFAR-10	Carlini (Carlini et al., 2022)	ViT-B/16	**93.0**	76.0	57.0	47.0	WRN28-10	86.0	66.0	55.0	37.0
	Ours	ViT-B/16	92.0	**82.0**	**69.0**	**56.0**	WRN28-10	**90.0**	**77.0**	**63.0**	**50.0**
ImageNet	Carlini (Carlini et al., 2022)	BEiT	77.0	76.0	71.0	60.0	WRN50-2	73.0	67.0	57.0	48.0
	Ours	BEiT	**80.0**	**78.0**	**76.0**	**71.0**	WRN50-2	**81.0**	**72.0**	**66.0**	**61.0**

Table 2: Certified accuracy of our method among different classifier. BeiT and ViT are pre-trained on a larger dataset ImageNet-22k and fine-tuned at ImageNet-1k and CIFAR-10 respectively. WideResNet is trained on ImageNet-1k for ImageNet and trained on CIFAR-10 from scratch for CIFAR-10.

based classifier is lower than Transformer-based classifier, DensePure with CNN-based model as the classifier can outperform Carlini et al. (2022) with ViT-based model as the classifier (except $\epsilon = 0$ for CIFAR-10).

6 Related Work

Using an off-the-shelf generative model to purify adversarial perturbations has become an important direction in adversarial defense. Previous works have developed various purification methods based on different generative models, such as GANs (Samangouei et al., 2018), autoregressive generative models (Song et al., 2018), and energy-based models (Du & Mordatch, 2019; Grathwohl et al., 2020; Hill et al., 2021). More recently, as diffusion models (or score-based models) achieve better generation quality than other generative models (Ho et al., 2020; Dhariwal & Nichol, 2021), many works consider using diffusion models for adversarial purification (Nie et al., 2022; Wu et al., 2022; Sun et al., 2022) Although they have found good empirical results in defending against existing adversarial attacks (Nie et al., 2022), there is no provable guarantee about the robustness about such methods. On the other hand, certified defenses provide guarantees of robustness (Mirman et al., 2018; Cohen et al., 2019; Lecuyer et al., 2019; Salman et al., 2020; Horváth et al., 2021; Zhang et al., 2018; Raghunathan et al., 2018a;b; Salman et al., 2019b; Wang et al., 2021). They provide a lower bounder of model accuracy under constrained perturbations. Among them, approaches Lecuyer et al. (2019); Cohen et al. (2019); Salman et al. (2019a); Jeong & Shin (2020); Zhai et al. (2020); Horváth et al. (2021); Jeong et al. (2021); Salman et al. (2020); Lee (2021); Carlini et al. (2022) based on randomized smoothing (Cohen et al., 2019) show the great scalability and achieve promising performance on large network and dataset. The most similar work to us is Carlini et al. (2022), which uses diffusion models combined with standard classifiers for certified defense. They view diffusion model as blackbox without having a theoretical under- standing of why and how the diffusion models contribute to such nontrivial certified robustness.

7 Conclusion

In this work, we theoretically prove that the diffusion model could purify adversarial examples back to the corresponding clean sample with high probability, as long as the data density of the corresponding clean samples is high enough. Our theoretical analysis characterizes the conditional distribution of the reversed samples given the adversarial input, generated by the diffusion model reverse process. Using the highest density point in the conditional distribution as the deterministic reversed sample, we identify the robust region of a given instance under the diffusion model reverse process, which is potentially much larger than previous methods. Our analysis inspires us to propose an effective pipeline DensePure, for adversarial robustness. We conduct comprehensive experiments to show the effectiveness of DensePure by evaluating the certified robustness via the randomized smoothing algorithm. Note that DensePure is an off-the-shelf pipeline that does not require training a smooth classifier. Our results show that DensePure achieves the new SOTA certified robustness for perturbation with $\mathcal{L}_2$-norm. We hope that our work sheds light on an in-depth understanding of the diffusion model for adversarial robustness.

Limitations. The time complexity of DensePure is high since it requires repeating the reverse process multiple times. In this paper, we use fast sampling to reduce the time complexity and show that the setting ($b = 2$ and $K = 10$) can achieve nontrivial certified accuracy. We leave the more advanced fast sampling strategy as the future direction.