

AdvDiff: Generating Unrestricted Adversarial Examples using Diffusion Models

Xuelong Dai¹, Kaisheng Liang¹, and Bin Xiao¹

The Hong Kong Polytechnic University
xuelong.dai@connect.polyu.hk, cskliang@comp.polyu.edu.hk,
b.xiao@polyu.edu.hk

Abstract. Unrestricted adversarial attacks present a serious threat to deep learning models and adversarial defense techniques. They pose severe security problems for deep learning applications because they can effectively bypass defense mechanisms. However, previous attack methods often directly inject Projected Gradient Descent (PGD) gradients into the sampling of generative models, which are not theoretically provable and thus generate unrealistic examples by incorporating adversarial objectives, especially for GAN-based methods on large-scale datasets like ImageNet. In this paper, we propose a new method, called AdvDiff, to generate unrestricted adversarial examples with diffusion models. We design two novel adversarial guidance techniques to conduct adversarial sampling in the reverse generation process of diffusion models. These two techniques are effective and stable in generating high-quality, realistic adversarial examples by integrating gradients of the target classifier interpretably. Experimental results on MNIST and ImageNet datasets demonstrate that AdvDiff is effective in generating unrestricted adversarial examples, which outperforms state-of-the-art unrestricted adversarial attack methods in terms of attack performance and generation quality.

Keywords: Unrestricted Adversarial Attacks · Diffusion Models · Interpretable Adversarial Diffusion Sampling

1 Introduction

While the deep learning (DL) community continues to explore the wide range of applications of DL models, researchers [37] have demonstrated that these models are highly susceptible to deception by adversarial examples. Adversarial examples are generated by adding perturbations to clean data. The perturbed examples can deceive DL classifiers with high confidence while remaining imperceptible to humans. Many strong attack methods [3, 8, 13, 21, 22, 25] are proposed and investigated to improve the robustness of DL models.

In contrast to existing perturbation-based adversarial attacks, Song et al. [35] found that using a well-trained generative adversarial network with an auxiliary classifier (AC-GAN) [29] can directly generate new adversarial examples without perturbing the clean data. These newly generated examples are considered **unrestricted** as they are obtained by optimizing input noise vectors

without any norm restrictions. Compared to traditional adversarial examples, unrestricted adversarial examples [9, 30] are more aggressive against current adversarial defenses. A malicious adversary can also generate an unlimited number of unrestricted adversarial examples using a trained GAN.

Diffusion models [17] are likelihood-based generative models proposed recently, which emerged as a strong competitor to GANs. Diffusion models have outperformed GANs for image synthesis tasks [12, 19, 31]. Compared with GAN models, diffusion models are more stable during training and provide better distribution coverage. Diffusion models contain two processes: a forward diffusion process and a reverse generation process. The forward diffusion process gradually adds Gaussian noise to the data and eventually transforms it into noise. The reverse generation process aims to recover the data from the noise by a denoising-like technique. A well-trained diffusion model is capable of generating images with random noise input. Similar to GAN models, diffusion models can achieve adversarial attacks by incorporating adversarial objectives [4–6].

GAN-based unrestricted adversarial attacks often exhibit poor performance on high-quality datasets, particularly in terms of visual quality, because they directly add the PGD perturbations to the GAN latents without theoretic supports. These attacks tend to generate low-quality adversarial examples compared to benign GAN examples [35]. Therefore, these attacks are not imperceptible among GAN synthetic data. Diffusion models, however, offer state-of-the-art generation performance [12] on challenging datasets like LSUN [41] and ImageNet [10]. The conditional diffusion models can generate images based on specific conditions by sampling from a perturbed conditional Gaussian noise, which can be carefully modified with adversarial objectives. These properties make diffusion models more suitable for conducting unrestricted adversarial attacks. Nevertheless, existing adversarial attack methods using diffusion models [4–6] adopt similar PGD perturbations to the sample in each reverse generation process, making them generate relatively low-quality adversarial examples.

In this paper, we propose a novel and interpretable unrestricted adversarial attack method called AdvDiff that utilizes diffusion models for adversarial examples generation, as shown in Figure 1. Specifically, AdvDiff uses a trained conditional diffusion model to conduct adversarial attacks with two new adversarial guidance techniques. 1) During the reverse generation process, we gradually add *adversarial guidance* by increasing the likelihood of the target attack label. 2) We perform the reverse generation process multiple times, adding adversarial prior knowledge to the initial noise with the *noise sampling guidance*.

Our theoretical analysis indicates that these adversarial guidance techniques can effectively craft adversarial examples by the reverse generation process with adversarial conditional sampling. Furthermore, the sampling of AdvDiff benefits from stable and high sample quality of the diffusion models sampling, which leads to the generation of realistic unrestricted adversarial examples. Through extensive experiments conducted on two datasets, i.e., the high-quality dataset ImageNet, and the small, robust dataset MNIST, we have observed a significant improvement in the attack performance using AdvDiff with diffusion models.

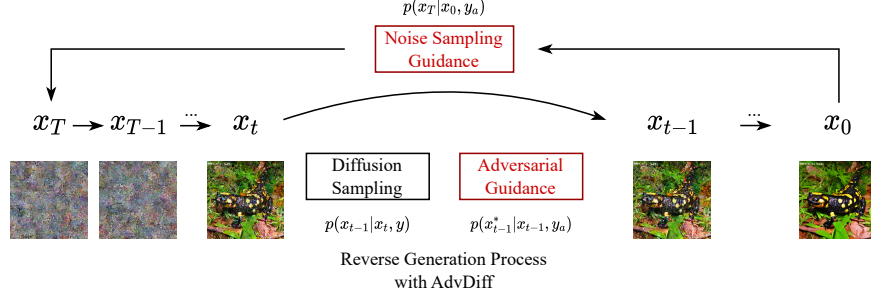


Fig.1: The two new guidance techniques in our AdvDiff to generate unrestricted adversarial examples. During the reverse generation process, the *adversarial guidance* is added at timestep x_t , which injects the adversarial objective y_a into the diffusion process. The *noise sampling guidance* modifies the original noise by increasing the conditional likelihood of y_a .

These results prove that our proposed AdvDiff is more effective than previous unrestricted adversarial attack methods in conducting unrestricted adversarial attacks to generate high-fidelity and diverse examples without decreasing the generation quality.

Our contributions can be summarized as follows:

- We propose AdvDiff, the new form unrestricted adversarial attack method that utilizes the reverse generation process of diffusion models to generate realistic adversarial examples.
- We design two new effective adversarial guidance techniques to the sampling process that incorporate adversarial objectives to the diffusion model without re-training the model. Theoretical analysis reveals that AdvDiff can generate unrestricted adversarial examples while preserving the high-quality and stable sampling of the conditional diffusion models.
- We perform extensive experiments to demonstrate that AdvDiff achieves an overwhelmingly better performance than GAN models on unrestricted adversarial example generation.

2 Preliminaries

In this section, we introduce the diffusion model and the classifier guidance for constructing our adversarial diffusion model.

2.1 Diffusion Model

The Denoising Diffusion Probabilistic Model (DDPM) [17] is utilized to learn the Gaussian transitions from $p(x_t) = \mathcal{N}(x_t; 0, \mathbf{I})$ to recover the data $x_0 \sim q(x_0)$

with a Markov chain. We call this denoising-like process the *reverse generation process*. Following the pre-defined T time steps, DDPM obtains a sequence of noisy data $\{x_{T-1}, \dots, x_1\}$ and finally recovers the data x_0 . It is defined as:

$$p_\theta(x_{t-1}|x_t) := \mathcal{N}(x_{t-1} : \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)) \quad (1)$$

Conversely, the *forward diffusion process* leverages a fixed Markov chain to iteratively add Gaussian noise to the sampled data $x_0 \sim q(x_0)$ according to the scheduling function β_1, \dots, β_N . Specifically, it is defined as:

$$q(x_t|x_{t-1}) := \mathcal{N}(x_t : \sqrt{1 - \beta_t}x_{t-1}, \beta_t \mathbf{I}) \quad (2)$$

Training of the DDPM requires accurate inference of the mean value μ_θ with a deep learning model in the reverse generation process. The objective is to learn the variational lower-bound (VLB) on $\log p_\theta(x_0)$. In order to complete the training objective, we train the model ϵ_θ to predict the added Gaussian noise in the forward diffusion process. The standard mean-squared error loss is adopted:

$$\mathcal{L}_{DDPM} := E_{t \sim [1, T], \epsilon \sim \mathcal{N}(0, \mathbf{I})} \|\epsilon - \epsilon_\theta(x_t, t)\|^2 \quad (3)$$

Song et al. [34] proposed DDIM, which provides an alternative noising process without restricting to a Markov chain. DDIM can achieve a much faster sampling than DDPM with the same training as DDPM. We perform experiments on both DDPM and DDIM to demonstrate the usability of our AdvDiff. DDPM is adopted to introduce our method for simplicity.

2.2 Classifier-Guided Guidance

Dhariwal et al. [12] achieved conditional diffusion sampling by adopting a trained classifier. The conditional information is injected into the diffusion model by modifying the mean value $\mu_\theta(x_t, t)$ of the samples according to the gradient of the prediction of the target class y by the trained classifier. They adopted log probability to calculate the gradient, and the mean value is given by:

$$\hat{\mu}_\theta(x_t, t) = \mu_\theta(x_t, t) + s \cdot \nabla_{x_t} \log p_\phi(y|x_t) \quad (4)$$

where s is the guidance scale.

2.3 Classifier-Free Guidance

Ho et al. [18] recently proposed a new conditional diffusion model using classifier-free guidance that injects class information without adopting an additional classifier. The classifier-free guidance utilizes a conditional diffusion model $p_\theta(x|y)$ for image synthesis with given labels. For effective training, they jointly train the unconditional diffusion model $p_\theta(x|\emptyset)$ and the conditional diffusion model $p_\theta(x|y)$, where the unconditional diffusion model is simply replacing the label

information with \emptyset . Sampling is performed by pushing the model towards the latent space of $p_\theta(x|y)$ and away from $p_\theta(x|\emptyset)$:

$$\hat{\epsilon}_\theta(x_t|y) = \epsilon_\theta(x_t|\emptyset) + w \cdot (\epsilon_\theta(x_t|y) - \epsilon_\theta(x_t|\emptyset)) \quad (5)$$

where w is the weight parameter for class guidance and \emptyset is the empty set.

The idea of classifier-free guidance is inspired by the gradient of an implicit classifier $p^i(y|x) \propto p(x|y)/p(x)$, the gradient of the classifier would be:

$$\begin{aligned} \nabla_x \log p^i(y|x) &\propto \nabla_x \log p(x|y) - \nabla_x \log p(x) \\ &\propto \epsilon_\theta(x_t|y) - \epsilon_\theta(x_t|\emptyset) \end{aligned} \quad (6)$$

The classifier-free guidance has a good capability of generating high-quality conditional images, which is critical for performing adversarial attacks. The generation of these images does not rely on a classification model and thus can better fit the conditional distribution of the data.

3 Adversarial Diffusion Sampling

3.1 Rethinking Unrestricted Adversarial Examples

Song et al. [35] presented a new form of adversarial examples called unrestricted adversarial examples (UAEs). These adversarial examples are not generated by adding perturbations over the clean data but are directly generated by any generative model. UAEs can be viewed as false negative errors in the classification tasks, and they can also bring server security problems to deep learning models. These generative-based UAEs can be formulated as:

$$A_{\text{UAE}} \triangleq \{x \in \mathcal{G}(z_{\text{adv}}, y) | y \neq f(x)\} \quad (7)$$

where $f(\cdot)$ is the target model for unrestricted adversarial attacks. The unrestricted adversarial attacks aim to generate UAEs that fool the target model while still can be visually perceived as the image from ground truth label y .

Previous UAE works adopt GAN models for the generation of UAEs, and these works perturb the GAN latents by maximizing the cross-entropy loss of the target model, i.e., $\max_{z_{\text{adv}}} \mathcal{L}(f(\mathcal{G}(z_{\text{adv}}, y)), y)$. Ideally, the generated UAEs should guarantee similar generation quality to the samples crafted by standard z because successful adversarial examples should be imperceptible to humans. In other words, UAEs should not be identified among the samples with adversarial latents and standard latents.

However, due to GAN's poor interpretability, there's no theoretical support on z_{adv} that can craft UAEs with normally trained GANs. The generator of GAN is not trained with $z_{\text{adv}} = z + \nabla \mathcal{L}$ but only $z \sim \mathcal{N}(0, \mathbf{I})$. Therefore, GAN-based UAEs encounter a significant decrease in generation quality because samples with z_{adv} are not well-trained compared with samples with $z \sim \mathcal{N}(0, \mathbf{I})$. Moreover, the GAN latents are sampled from low dimensional latent spaces. Therefore, GANs are extremely sensitive to the latent z [23, 33]. If we inject




Fig. 2: Unrestricted adversarial examples generated by the diffusion model. The generated adversarial examples should be visually indistinguishable from clean data with label y but wrongly classified by the target classifier f .

gradients of the classification results into GAN latents, GAN-based methods are more likely to generate flipped-label UAEs (images corresponding to the targeted attack label y_a instead of the conditional generation label y) and distorted UAEs. However, these generation issues are hard to address only by attack success rate (ASR). In other words, even with a high ASR, some of the successful UAEs with GAN-based methods should be identified as failure cases for poor visual quality. However, such cases can not be reflected by ASR but can be evaluated by generation quality. All these problems may indicate that GAN models are not suitable for generative-based adversarial attacks.

Diffusion models have shown better performance on image generation than GAN models [12]. They are log-likelihood models with interpretable generation processes. In this paper, we aim to generate UAEs by injecting the adversarial loss with theoretical proof and without sabotaging the benign generation process, where we increase the conditional likelihood on the target attack label by following the diffusion process. The perturbations are gradually injected with the backward generation process of the diffusion model by the same sample procedure. As shown in Figure 2, the diffusion model can sample images from the conditional distribution $p(x|y)$. The samples from $p(x|y, f(x) \neq y)$ are the adversarial examples that are misclassified by $f(\cdot)$. These examples also follow the data distribution $p(x|y)$ but on the other side of the label y 's decision boundary of $f(\cdot)$. Moreover, the diffusion model's generation process takes multiple sampling steps. Thus, we don't need one strong perturbation to the latent like GAN-based methods. The AdvDiff perturbations at each step are unnoticeable, and perturbations are added to the high dimensional sampled data rather than low dimensional latents. Therefore, AdvDiff with diffusion models can preserve the generation quality and barely generates flipped-label or distorted UAEs.

3.2 Adversarial Diffusion Sampling with Theoretical Support

There are several existing adversarial attack methods [4–6] that adopt diffusion models to generate adversarial examples. However, these methods still adopt

PGD or I-FGSM gradients to perturb the diffusion process for constructing adversarial examples. As discussed earlier, the generation process of diffusion models is a specially designed sampling process from given distributions. Such adversarial gradients change the original generation process and can harm the generation quality of the diffusion model. Additionally, these methods fail to give a comprehensive discussion of the adversarial guidance with theoretical analysis. Therefore, we aim to design a **general** and **interpretable** method to generate adversarial examples using diffusion models **without** affecting the benign diffusion process.

3.3 Adversarial Guidance

Inspired by Dhariwal’s work [12] that achieves the conditional image generation by classifier gradient guidance $\nabla_{x_t} \log p_\phi(y|x_t)$, we generate our UAEs with adversarial gradient guidance over the reverse generation process. Our attack aims at utilizing a conditional diffusion model $\epsilon_\theta(x_t, y)$ to generate x_0 fits the ground truth label y while deceiving the target classifier with $p_f(x_0) \neq y$. These generated samples are the false negative results in p_f ’s classification results.

Normally, we will obtain the images with label y by following the standard reverse generation process with classifier-free guidance:

$$x_{t-1} = \mu(x_t, y) + \sigma_t \varepsilon \quad (8)$$

where $\mu(x_t, y)$ is the conditional mean value and ε is sampled from $\varepsilon \sim \mathcal{N}(0, \mathbf{I})$.

Sampling by Equation 8, we obtain the samples with the generation process $p(x_{t-1}|x_t, y)$. Following the above-mentioned definition of UAEs, we can get our adversarial examples by adding adversarial guidance to the standard reverse process, which is performing another sampling with the adversarial generation process $p(x_{t-1}^*|x_{t-1}, f(x) \neq y)$. We find that specifying a target label for the adversarial generation process is more effective during experiments. Suggest the target label y_a is the target for the adversarial attacks, the adversarial example is sampled by the following steps:

$$x_{t-1}^* = x_{t-1} + \sigma_t^2 s \nabla_{x_{t-1}} \log p_f(y_a|x_{t-1}) \quad (9)$$

where s is the adversarial guidance scale. The derivation of Equation 9 is given in Appendix A. Intuitively, the adversarial guidance encourages the generation of samples with a higher likelihood of the target label.

In practice, we utilize the classifier-free guidance to train a conditional diffusion model $\epsilon_\theta(\cdot)$ as our basic generation model.

3.4 Noise Sampling Guidance

We can improve the reverse process by adding an adversarial label prior to the noise data x_T . The UAEs are a subset of the dataset labeled with y . They can be viewed as the conditional probability distribution with $p(x|y, f(x) = y_a)$ during




Fig. 3: Adversarial examples on the MNIST dataset. Perturbation-based attack methods generate noise patterns to conduct attacks, while unrestricted adversarial attacks (U-GAN and AdvDiff) are imperceptible to the clean data.

sampling, and y_a is the target label for the adversarial attack. Therefore, we can add the adversarial label prior to x_t with Bayes’ theorem:

$$\begin{aligned} p(x_T|y_a) &= \frac{p(y_a|x_T)p(x_T)}{p(y_a)} = \frac{p(y_a|x_T, x_0)p(x_T|x_0)}{p(y_a|x_0)} \\ &= p(x_T|x_0)e^{\log p(y_a|x_T) - \log p(y_a|x_0)} \end{aligned} \quad (10)$$

We can infer the x_t with the adversarial prior by Equation 10, i.e.,

$$x_T = (\mu(x_0, y) + \sigma_t \varepsilon) + \bar{\sigma}_T^2 a \nabla_{x_0} \log p_f(y_a|x_0) \quad (11)$$

where a is the noise sampling guidance scale. See Appendix B for detailed proof.

Equation 11 is similar to Equation 9 as they both add adversarial guidance to the reverse generation process. However, the noise sampling guidance is added to x_T according to the final classification gradient $\nabla_{x_0} \log p_f(y_a|x_0)$, which provides a strong adversarial guidance signal directly to the initial input of the generative model. The gradient of Equation 11 is effective as it reflects the eventual classification result of the target classifier.

3.5 Training-Free Adversarial Attack

The proposed adversarial attack does not require additional modification on the training of the diffusion model. The adversarial examples are sampled by using Algorithm 1 over the trained classifier-free diffusion model $\epsilon_\theta(\cdot)$. We give the AdvDiff algorithm on DDIM in the Appendix.

4 Experiments

Datasets and target models. We use two datasets for major evaluation: MNIST [11] and ImageNet [10]. MNIST is a 10-classes dataset consisting of handwritten numbers from 0 to 9. We adopt the MNIST dataset to evaluate our method for low-quality robust image generation. ImageNet is a large visual

Algorithm 1 DDPM Adversarial Diffusion Sampling

Require: y_a : target label for adversarial attack
Require: y : ground truth class label
Require: s, a : adversarial guidance scale
Require: w : classification guidance scale
Require: N : noise sampling guidance steps
Require: T : reverse generation process timestep

- 1: $x_T \sim \mathcal{N}(0, \mathbf{I})$
- 2: $x_{adv} = \emptyset$
- 3: **for** $i = 1 \dots N$ **do**
- 4: **for** $t = T, \dots, 1$ **do**
- 5: $\tilde{\epsilon}_t = (1 + w)\epsilon_\theta(x_t, y) - w\epsilon_\theta(x_t)$
- 6: Classifier-free sampling x_{t-1} with $\tilde{\epsilon}_t$.
- 7: Input x_{t-1} to target model and get the gradient $\log p_f(y_a|x_{t-1})$
- 8: $x_{t-1}^* = x_{t-1} + \sigma_t^2 s \nabla_{x_{t-1}} \log p_f(y_a|x_{t-1})$
- 9: **end for**
- 10: Obtain classification result from $f(x_0)$
- 11: Compute the gradient with $\log p_f(y_a|x_0)$
- 12: Update x_T by $x_T = x_T + \bar{\sigma}_T^2 a \nabla_{x_0} \log p_f(y_a|x_0)$
- 13: $x_{adv} \leftarrow x_0$ if $f(x_0) = y_a$
- 14: **end for**
- 15: **return** x_{adv}

database with 1000 object classes and is used for the high-quality generation task. For target classifiers, we adopt simple LeNet5 [20], and ResNet18 [15] for the MNIST dataset, and the widely-used ResNet50 [15] and WideResNet50-2 [42] for the ImageNet dataset.

Comparisons. It is not applicable to give a clear comparison between perturbation attacks and unrestricted attacks because perturbation attacks have the corresponding ground truth while unrestricted attacks do not. We mainly compare our method with the unrestricted adversarial attack U-GAN [35] and give the discussion with the AutoAttack [8], PGD [25], BIM [13], and C&W [3] perturbation-based attacks under norm $\ell_{\text{inf}} = 8/255$. For U-GAN, We adopt AC-GAN [29] for the MNIST dataset, and SAGAN [43] and BigGAN [2] for the ImageNet dataset, as AC-GAN has shown poor performance on ImageNet. We use the official code from DiffAttack [4] and implement AdvDiffuser by ourselves [5] for comparisons. We do not compare with Chen et al. [6], because they use a similar method as DiffAttack and without official code. Because existing diffusion model attacks are all untargeted attacks, we include the untargeted version of AdvDiff for a clear comparison, which is represented by “AdvDiff-Untargeted”.

Implementation details. Because our adversarial diffusion sampling does not require additional training to the original diffusion model, we use the pre-trained diffusion model in our experiment. We adopt DDPM [17] with classifier-free guidance for the MNIST dataset and Latent Diffusion Model (LDM) [31] with DDIM sampler for the ImageNet dataset. For MNIST dataset, we use $N =$

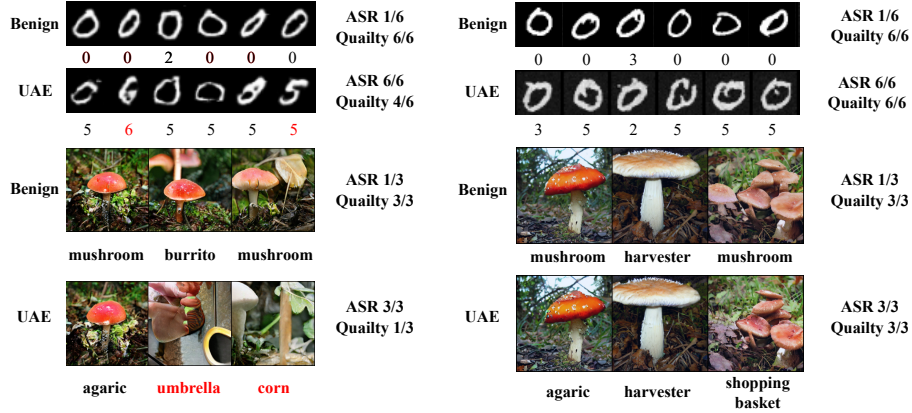


Fig. 4: Comparisons of unrestricted adversarial attacks between GANs and diffusion models on two datasets.. Left: generated samples from U-GAN (BigGAN for ImageNet dataset). Right: generated samples from AdvDiff. We generate unrestricted adversarial examples on the MNIST “0” label and ImageNet “mushroom” label. U-GAN is more likely to generate adversarial examples with the target label, i.e., examples with red font. However, AdvDiff tends to generate the “false negative” samples by the target classifier by combining features from the target label.

10, $s = 0.5$, and $a = 1.0$, And $N = 5$, $s = 0.7$, and $a = 0.5$ for ImageNet dataset. More details and experiments are given in the Appendix.

Evaluation metrics. We utilize the top-1 classification result to evaluate the Attack Success Rate (ASR) on different attack methods under untargeted attack settings. As discussed earlier, GAN-based UAEs often encounter severe generation quality drops compared to benign GAN samples. Therefore, we give comparisons of generation performance on ImageNet to evaluate the attack performance of different UAEs in imperceptibly. The results are averaged with five runs. We use ResNet50 as the target model for default settings.

4.1 Attack Performance

MNIST We show the attack success rate against the normally trained model and adversarially trained model [25] in the MNIST dataset. All the selected adversarial attacks achieve over 90% attack success rate against the normally trained model. The adversarially trained model can effectively defend against perturbation-based adversarial attacks for their noise-like perturbation generation patterns, as reported in Table 1. However, the UAEs obviously perform better with their non-perturbed image generation. Despite the fact that the unrestricted attack can break through the adversarial defenses, the crafted adversarial examples should also be imperceptible to humans for a reasonably successful attack. The visualized adversarial examples in Figure 3 show that the

Table 1: The attack success rate on MNIST dataset.

Method	ASR(%)			
	LeNet5		ResNet18	
	Clean	PGD-AT	Clean	PGD-AT
PGD	99.8	25.6	99.3	20.8
BIM	99.6	34.6	100	31.5
C&W	100	68.6	100	64.5
U-GAN	88.5	<i>79.4</i>	85.6	<i>75.1</i>
AdvDiff	94.2	88.6	92.1	86.5

Table 2: The attack success rate on ImageNet dataset. U-SAGAN and U-BigGAN represent the base GAN models for U-GAN are SAGAN and BigGAN, respectively.

Method	ASR(%)						Time (s)
	ResNet50			WideResNet50-2			
	Clean	Diff	Pure PGD-AT	Clean	Diff	Pure PGD-AT	
AutoAttack	95.1	22.2	56.2	94.9	20.6	55.4	0.5
U-SAGAN	99.3	30.5	80.6	98.9	28.6	70.1	10.4
U-BigGAN	96.8	40.1	81.5	96.5	35.5	78.4	11.2
AdvDiffuser	95.4	28.9	90.6	94.6	26.5	88.9	38.6
DiffAttack	92.8	30.6	88.4	90.6	27.6	85.3	28.2
AdvDiff	99.8	41.6	92.4	99.9	38.5	90.6	9.2
AdvDiff-Untargeted	99.5	75.2	94.5	99.4	70.5	92.6	9.6

perturbation-based adversarial attacks tend to blur the original images while U-GAN can generate mislabeled adversarial examples.

ImageNet It is reported that deep learning models on ImageNet are extremely vulnerable to adversarial attacks. However, the state-of-the-art adversarial defense DiffPure [28] and adversarial training [25] can still defend against the perturbation-based attacks, as reported in Table 2. More UAs evade the current defenses, but the generation quality of U-GAN is relatively poor compared to our adversarial examples. This phenomenon also shows that the performance of UAs is heavily affected by the generation quality of the generation model. The adversarial examples generated by AdvDiff are more aggressive and stealthy than U-GAN’s. Meanwhile, the generation speed of AdvDiff is the best among all the unrestricted adversarial attack methods. Note that we adopt the clean images generated by LDM to achieve DiffAttack and AutoAttack for a fair comparison.

4.2 Generation Quality: True ASR for UAs

We witness similar ASR with U-GAN and AdvDiff. However, imperceptibility is also critical for a successful unrestricted adversarial attack, so we adopt the evaluation metrics in [12] to compare the generation quality with and without performing unrestricted attacks. Table 3 shows that the AdvDiff achieves an overwhelming better IS score and similar FID score on the large-scale ImageNet dataset, where FID [16] and IS [32] scores are commonly adopted for evaluating the quality of a generative model. Because the generation of UAs does not modify the data distribution of the generated images, the Precision score can be inferred as generation quality, while the Recall score indicates the flipped-label problems. We witness the frequent generation of flipped-label UAs and low-quality UAs from GAN-based methods, which is reflected by the decrease in the Precision score and the increase in the Recall score. Figure 4 illustrates this problem with some examples. It can be further proved that U-BigGAN achieves much higher image quality on non-reference metrics than reference metrics, as shown in Table 4.

We find the IS score is heavily affected by the transferability of adversarial examples due to the calculation method. Therefore, we further compare the im-

Table 3: The generation performance on the ImagetNet dataset.

Method	FID (\downarrow)	sFID (\downarrow)	IS (\uparrow)	Precision (\uparrow)	Recall (\uparrow)
SAGAN	41.9	50.2	26.7	0.50	0.51
BigGAN	19.3	45.7	250.3	0.95	0.21
LDM	12.3	25.4	385.5	0.94	0.73
U-SAGAN	52.8/+26%	52.2/+4%	12.5/-53%	0.58	0.57
U-BigGAN	25.4/+31%	52.1/+14%	129.4/-48%	0.81	0.35
AdvDiffuser	26.8/+117%	38.6/+51%	206.8/-46%	0.70	0.75
DiffAttack	20.5/+66%	40.2/+58%	264.3/-31%	0.83	0.73
AdvDiff	16.2 /+31%	30.4 /+20%	343.8 /-10%	0.90	0.75
AdvDiff-Untargeted	22.8/+85%	33.4/+28%	220.8/-45%	0.85	0.76

Table 4: The image quality on the ImagetNet dataset.

Method	FID (\downarrow)	LPIPS (\downarrow)	SSIM (\uparrow)	BRISQUE [26] (\downarrow)	TRES (\uparrow)
AutoAttack	26.5	0.72	0.21	34.4	69.8
U-BigGAN	25.4	0.50	0.32	19.4	80.3
AdvDiffuser	26.8	0.21	0.84	18.9	75.6
DiffAttack	20.5	0.15	0.75	22.6	67.8
AdvDiff	16.2	0.03	0.96	18.1	82.1
AdvDiff-Untargeted	22.8	0.14	0.85	16.2	76.8

age quality of adversarial examples by commonly used metrics in Table 4. The results show that AdvDiff (average 5 out of 5) and AdvDiff-Untargeted (average 4 out of 5) outperform existing adversarial attack methods using diffusion models. The perturbation-based adversarial attacks, i.e., AutoAttack, achieve much worse image quality compared with UAEs.

4.3 UAEs against Defenses and Black-box Models

Current defenses assume the adversarial examples are based on perturbations over data from the training dataset, i.e., $x_{adv} = x + \nabla \mathcal{L}, x \in D$. However, UAEs are synthetic data generated by the generative model. Because of different data sources, current defenses are hard to defend UAEs, which brings severe security concerns to deep learning applications. The proposed AdvDiff achieves an average of 36.8% ASR against various defenses, while AutoAttack only achieves 30.7% ASR with significantly lower image quality. We also test the attack transferability of AdvDiff and the results show that the untargeted version of AdvDiff achieves the best performance against black-box models. Experiment results are given in Table 5.

4.4 Better Adversarial Diffusion Sampling

We present detailed comparisons with DiffAttack and AdvDiffuser. The results show that the proposed adversarial guidance achieves significantly higher gener-

Table 5: The attack success rates (%) of ResNet50 examples for transfer attack and attack against defenses on the ImagetNet dataset.

Method	ResNet-152 [15]	Inception v3 [36]	ViT-B [14]	BEiT [1]
AutoAttack	32.5	38.6	9.3	45.3
U-BigGAN	30.8	35.3	30.1	69.4
AdvDiffuser	18.3	20.0	18.5	79.4
DiffAttack	21.1	43.9	17.4	78.0
AdvDiff	20.5	14.9	17.8	78.8
AdvDiff-Untargeted	52.0	42.7	36.0	81.5
Method	Adv-Inception [25]	AdvProp [38]	DiffPure [28]	HGD [24]
AutoAttack	14.6	69.6	22.2	20.5
U-BigGAN	40.6	75.2	40.1	22.6
AdvDiffuser	24.4	84.0	30.5	10.8
DiffAttack	30.9	85.1	30.6	20.5
AdvDiff	19.4	89.7	41.6	17.8
AdvDiff-Untargeted	60.1	95.3	75.2	53.8
Method	R&P [39]	RS [7]	NRP [27]	Bit-Red [40]
AutoAttack	20.6	38.9	39.4	19.8
U-BigGAN	14.2	34.5	30.9	13.1
AdvDiffuser	15.4	38.4	40.5	11.4
DiffAttack	23.7	40.8	38.5	20.1
AdvDiff	17.4	47.6	45.2	15.8
AdvDiff-Untargeted	56.8	82.8	74.2	52.6

ation quality than PGD-based adversarial guidance. With PGD gradient guidance, the diffusion model generates images with a similar Recall score but a much lower Precision score, which indicates that the PGD gradient influences the benign generation process and causes the generation of low-quality images. The result proves that the adversarial guidance of diffusion models should be carefully designed without affecting the benign sampling process. Meanwhile, the generation speed of AdvDiff is the best among the existing diffusion attack methods. Note that AdvDiff (36.8%) slightly outperforms AdvDiffuser (32.0%) and DiffAttack (36.2%) against defenses. However, previous attacks achieve slightly better transfer attack performance than the original AdvDiff. The reason could be the gradient of the cross-entropy loss is shared among nearly all the deep learning models and is better at attack transferability against these models. Nevertheless, the untargeted version of AdvDiff achieves overwhelmingly better performance, which further demonstrates the effectiveness of the proposed adversarial sampling. But the generation quality is affected, we leave a better design in the future work.

4.5 Ablation Study

We discuss the impact of the parameters of AdvDiff in the subsection. Note that our proposed method does not require re-training the conditional diffusion models. The ablation study is performed only on the sampling process.




Fig. 5: Ablation study of the impact of parameters in AdvDiff. The results are generated from the ImageNet dataset against the ResNet50 model. We adopt the ASR and IS scores to show the impact of attack performance and generation quality.

Adversarial guidance scale s and a . The magnitudes of s and a greatly affect the ASR of AdvDiff, as shown in Figure 5. Noted that we witness the generation of unrealistic images when setting the adversarial guidance extremely large. See the Appendix for detailed discussions.

Noise sampling guidance steps N . Like the iteration times of GAN-based unrestricted adversarial attacks, larger steps N can effectively increase the attack performance against an accurate classifier, as shown in Figure 5. However, it can affect the initial noise distribution and hence decreases the generation quality. During experiments, we observe that adversarial guidance is already capable of generating adversarial examples with high ASR. Thus, we can set a small noise sampling guidance step N for better sample quality.

Adversarial guidance timestep t^* . The reverse diffusion process gradually denoises the input noise. Therefore we generally get noisy images at most timesteps. Because the target classifier is not able to classify the noisy input, the adversarial guidance is not effective in the early reverse diffusion process. Figure 5 shows our results, and we can improve the performance of adversarial guidance by training a separate classifier, which we leave for future work.

5 Conclusion

In this work, we propose a new method called AdvDiff, which can conduct unrestricted adversarial attacks using any pre-trained conditional diffusion model. We propose two novel adversarial guidance techniques in AdvDiff that lead diffusion models to obtain high-quality, realistic adversarial examples without disrupting the diffusion process. Experiments show that our AdvDiff vastly outperforms GAN-based and diffusion-based attacks in terms of attack success rate and image generation quality, especially in the ImageNet dataset. AdvDiff indicates that diffusion models have demonstrated effectiveness in adversarial attacks, and highlights the need for further research to enhance AI model robustness against unrestricted attacks.

Acknowledgements

This work was supported in part by HK RGC GRF under Grant PolyU 15201323.

References

1. Bao, H., Dong, L., Piao, S., Wei, F.: BEit: BERT pre-training of image transformers. In: International Conference on Learning Representations (2022)
2. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. In: International Conference on Learning Representations (2018)
3. Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. In: 2017 IEEE Symposium on Security and Privacy (SP). pp. 39–57. IEEE (2017)
4. Chen, J., Chen, H., Chen, K., Zhang, Y., Zou, Z., Shi, Z.: Diffusion models for imperceptible and transferable adversarial attack. arXiv preprint arXiv:2305.08192 (2023)
5. Chen, X., Gao, X., Zhao, J., Ye, K., Xu, C.Z.: Advdiffuser: Natural adversarial example synthesis with diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4562–4572 (2023)
6. Chen, Z., Li, B., Wu, S., Jiang, K., Ding, S., Zhang, W.: Content-based unrestricted adversarial attack. In: Proceedings of the 37th International Conference on Neural Information Processing Systems. pp. 51719–51733 (2023)
7. Cohen, J., Rosenfeld, E., Kolter, Z.: Certified adversarial robustness via randomized smoothing. In: international conference on machine learning. pp. 1310–1320. PMLR (2019)
8. Croce, F., Hein, M.: Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In: International conference on machine learning. pp. 2206–2216. PMLR (2020)
9. Deb, D., Zhang, J., Jain, A.K.: Advfaces: Adversarial face synthesis. In: 2020 IEEE International Joint Conference on Biometrics (IJCB). pp. 1–10. IEEE (2020)
10. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
11. Deng, L.: The mnist database of handwritten digit images for machine learning research [best of the web]. IEEE signal processing magazine **29**(6), 141–142 (2012)
12. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. Advances in Neural Information Processing Systems **34**, 8780–8794 (2021)
13. Dong, Y., Liao, F., Pang, T., Su, H., Zhu, J., Hu, X., Li, J.: Boosting adversarial attacks with momentum. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 9185–9193 (2018)
14. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2020)
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
16. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems **30** (2017)

17. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020)
18. Ho, J., Salimans, T.: Classifier-free diffusion guidance. In: *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications* (2021)
19. Kim, G., Kwon, T., Ye, J.C.: Diffusionclip: Text-guided diffusion models for robust image manipulation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2426–2435 (2022)
20. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
21. Li, Y., Li, Y., Dai, X., Guo, S., Xiao, B.: Physical-world optical adversarial attacks on 3d face recognition. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 24699–24708 (2023)
22. Liang, K., Xiao, B.: Styleless: boosting the transferability of adversarial examples. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8163–8172 (2023)
23. Liang, Y., Wu, J., Lai, Y.K., Qin, Y.: Exploring and exploiting hubness priors for high-quality gan latent sampling. In: *International Conference on Machine Learning*. pp. 13271–13284. PMLR (2022)
24. Liao, F., Liang, M., Dong, Y., Pang, T., Hu, X., Zhu, J.: Defense against adversarial attacks using high-level representation guided denoiser. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1778–1787 (2018)
25. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. In: *International Conference on Learning Representations* (2018)
26. Mittal, A., Moorthy, A.K., Bovik, A.C.: Blind/referenceless image spatial quality evaluator. In: *2011 conference record of the forty fifth asilomar conference on signals, systems and computers (ASILOMAR)*. pp. 723–727. IEEE (2011)
27. Naseer, M., Khan, S., Hayat, M., Khan, F.S., Porikli, F.: A self-supervised approach for adversarial robustness. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 262–271 (2020)
28. Nie, W., Guo, B., Huang, Y., Xiao, C., Vahdat, A., Anandkumar, A.: Diffusion models for adversarial purification. In: *International Conference on Machine Learning (ICML)* (2022)
29. Odena, A., Olah, C., Shlens, J.: Conditional image synthesis with auxiliary classifier gans. In: *Proceedings of the International conference on machine learning*. pp. 2642–2651 (2017)
30. Qiu, H., Xiao, C., Yang, L., Yan, X., Lee, H., Li, B.: Semanticadv: Generating adversarial examples via attribute-conditioned image editing. In: *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*. pp. 19–37. Springer (2020)
31. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10684–10695 (2022)
32. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. *Advances in neural information processing systems* **29** (2016)
33. Shen, Y., Gu, J., Tang, X., Zhou, B.: Interpreting the latent space of gans for semantic face editing. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9243–9252 (2020)

34. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations (2021)
35. Song, Y., Shu, R., Kushman, N., Ermon, S.: Constructing unrestricted adversarial examples with generative models. In: Proceedings of the Advances in Neural Information Processing Systems. pp. 8322–8333 (2018)
36. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2818–2826 (2016)
37. 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)
38. Xie, C., Tan, M., Gong, B., Wang, J., Yuille, A.L., Le, Q.V.: Adversarial examples improve image recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 819–828 (2020)
39. Xie, C., Wang, J., Zhang, Z., Ren, Z., Yuille, A.: Mitigating adversarial effects through randomization. In: International Conference on Learning Representations (2018)
40. Xu, W., Evans, D., Qi, Y.: Feature squeezing: Detecting adversarial examples in deep neural networks. arXiv preprint arXiv:1704.01155 (2017)
41. Yu, F., Seff, A., Zhang, Y., Song, S., Funkhouser, T., Xiao, J.: Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. arXiv preprint arXiv:1506.03365 (2015)
42. Zagoruyko, S., Komodakis, N.: Wide residual networks. In: Proceedings of the British Machine Vision Conference 2016. British Machine Vision Association (2016)
43. Zhang, H., Goodfellow, I., Metaxas, D., Odena, A.: Self-attention generative adversarial networks. In: International conference on machine learning. pp. 7354–7363. PMLR (2019)