TAIGen: Training-Free Adversarial Image Generation via Diffusion Models

Susim Roy[1,2], Anubhooti Jain[2], Mayank Vatsa[2], Richa Singh[2]
[1]University at Buffalo, [2]IIT Jodhpur
susimmuk@buffalo.edu, {jain.44, mvatsa, richa}@iitj.ac.in

Abstract

Adversarial attacks from generative models often produce low-quality images and require substantial computational resources. Diffusion models, though capable of high-quality generation, typically need hundreds of sampling steps for adversarial generation. This paper introduces TAIGen, a training-free black-box method for efficient adversarial image generation. TAIGen produces adversarial examples using only 3-20 sampling steps from unconditional diffusion models. Our key finding is that perturbations injected during the mixing step interval achieve comparable attack effectiveness without processing all timesteps. We develop a selective RGB channel strategy that applies attention maps to the red channel while using GradCAM-guided perturbations on green and blue channels. This design preserves image structure while maximizing misclassification in target models. TAIGen maintains visual quality with PSNR above 30 dB across all tested datasets. On ImageNet with VGGNet as source, TAIGen achieves 70.6% success against ResNet, 80.8% against MNASNet, and 97.8% against ShuffleNet. The method generates adversarial examples 10× faster than existing diffusion-based attacks. Our method achieves the lowest robust accuracy, indicating it is the most impactful attack as the defense mechanism is least successful in purifying the images generated by TAIGen.

1. Introduction

Diffusion models have quickly become a cornerstone in generative AI and their performance has surpassed the previous benchmarks in many tasks [6, 12, 15]. Diffusion models have shown to be successful in overcoming the issues faced by previous models such as training instabilities by GANs or the limited expressiveness by the autoregressive models. Additionally, the semantic space of diffusion models stays consistent over the timesteps[31] which can be used for image manipulation without re-training. They have been utilized to generate adversarial examples, which can lead deep learning models to fail at their intended tasks.

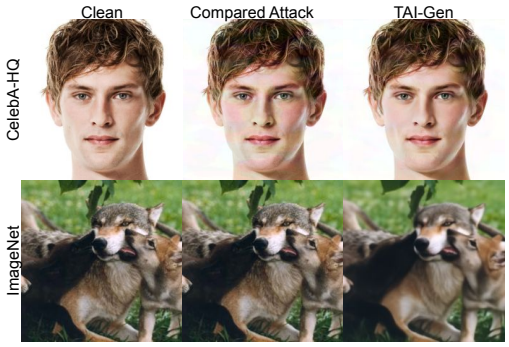

Figure 1. Images with their adversarially attacked counterparts. The first column is the clean images. The second is adversarial images generated using BPDA+EOT attack for the CelebA-HQ and AutoAttack for the ImageNet datasets. The third column is the adversarial images generated using our proposed TAIGen attack.

Some works have explored generating adversarial samples using diffusion models [27, 33, 56]. As can be seen in Figure 1, the different attack methods add unwanted artifacts to certain regions in faces(top row) and in multi-subject images(bottom row) affecting it's quality. However, most of them are unrestricted in nature and utilize diffusion models in a white-box setting. They also utilize almost all the timesteps involved in the diffusion generation process which is inefficient. Finally, those methods have not fully utilized the underlying variation of the training dataset stored in the model which undermines the generative capability.

We propose a training-free method, TAIGen, that utilizes only a handful of timesteps to generate adversarial examples in a quick and robust manner by selectively modifying the RGB channels of the image. We first find the significant timesteps by using key timestep [60] for convergence of Gaussian distribution to the data distribution. We further take advantage of attention maps as well as GradCAMs [44] to perturb the images. The RGB channels are modified selectively with this information as discussed later in detail. We showcase the performance of the proposed algorithm on three datasets – CIFAR-10, CelebA-HQ, and ImageNet

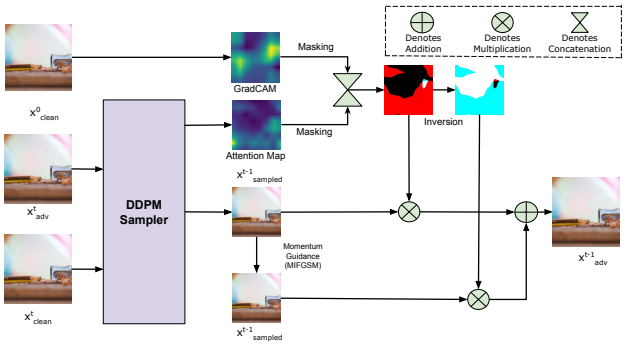

Figure 2. The proposed TAIGen adversarial attack where masked GradCAM and Attention Map are utilized for N time steps in an iterative manner. We also use MI-FGSM [16] during the iteration to introduce momentum guidance. The DDPM Sampler denotes the noise prediction model and the backward process scheduler as a single unit.

datasets along with GradCAM analysis for the resulting adversarial images.

2. Related Work

In this section, we review the current landscape for different adversarial attacks and purification methods.

Adversarial Attacks: Adversarial Attacks continue to mislead deep learning models. Current research has shown the impressive success of adversarial attacks on a variety of models and modalities. The attacks can be categorized as white, gray, and black-box attacks based on the information available to the attacker about the source model. Further, they can be categorized based on their formulation as gradient-based, patch-based, and semantic-based methods. The adversarial perturbation can be solved as a constrained optimization problem, like FGSM [18], PGD [36], or CW [7] attacks, they are restricted by l_p-norm and utilize gradients of the source model. These attacks have been further improved by using parameters like momentum in MI-FGSM [16] or AdaMSI-FGM [34] attacks. Another way is to target semantic space to create adversarial examples for more natural samples like ColorFool [45]. These adversarial samples can also be generated, as has been done using VAEs, GANs [29, 48, 55], and more recently, diffusion models. With GANs, there are attacks like the Latent-HSJA attack [38], which generates unrestricted adversarial examples in a black-box setting, or the GMAA attack [32], which creates a whole adversarial manifold over a point-wise generation.

Owing to the poor quality of generated samples, especially on high-quality datasets, Diffusion models have been used instead of GANs to generate the adversarial samples. The models can utilize gradients as in AdvDiff [13], a training-free method, where PGD gradients are injected into the sampling of generative models giving two novel guidance techniques called adversarial and noise sampling guidance for adversarial example generation. AdvDiffuser [9] also generates unrestricted adversarial examples and utilizes GradCAM [44] to retain important regions, only perturbing the less important ones. Even for face recognition systems, diffusion models have been used for generation. Adv-Diffusion [33] is one such attack that protects face identity by masking them using a reference image via a latent diffusion model. Another approach is Diff-AM [50], which is a diffused-based adversarial makeup transfer method used for face protection by first having a text-guided makeup removal module, followed by image-guided adversarial makeup transfer. An ensemble of face recognition models is used as the attack strategy within this method. DiffAttack [27] is an attack that breaks diffusion-based adversarial purification methods by proposing a novel deviated reconstruction loss. These diffusion-based attacks are able to generate high-quality adversarial samples, however, with our work, we propose a training-free, less-time-consuming diffusion-based generation attack.

Adversarial Defense and Purification: With attack, several defense strategies have also been proposed in the literature, and similar to adversarial attacks, a variety of architectures can be utilized to defend against adversarial samples. Adversarial training still remains one of the most effective defense techniques. A self-supervised adversarial training method, CAFD [59], was proposed to remove adversarial noise in the class activation space. Another adversarial training method [41] was proposed using generative models with a disentangled latent space to generate low, mid, and high-level changes in the samples. Several works have also studied adversarial robustness and detection in face recognition, including efforts on universal perturbation detection [1], attack-agnostic perturbation detectors [2], and robustness analysis of deep models against adversarial threats [19]. More recently diffusion models have been utilized for adversarial purification. DiffPure [39] is one such method where a clean image is generated after diffusing the adversarial sample with a small amount of noise in the diffusion process. MimicDiffusion [47] approximates the generation process for a clean image and compares it with that of the adversarial sample to purify the image given two guidance techniques.

3. TAIGen - The Proposed Attack

Our work, Figure 2, is based primarily on the Denoising Diffusion Probabilistic Model (DDPM); we therefore expand on it first in this section, followed by the problem formulation, the choice of a few sampling steps used in our al-

gorithm, the selective information passed to different channels of the intermediate image, and the final methodology.

3.1. Background

In this section, we introduce the Denoising Diffusion Probabilistic Model (DDPM) [22] which works via the inversion and sampling processes. Let us denote both the processes with the timestep index t with T as the diffusion length and each sample as $\{x_t\}_{t=0}^{t=T}$. DDPM constructs a discrete Markov chain $\{x_t\}_{t=0}^{t=T}$ with discrete time variables t following $p(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1-\beta_t}x_{t-1}, \beta_t\mathbf{I})$, where β_t is the sequence of positive noise scales(e.g. linear, cosine scheduling) during inversion. Considering $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{s=1}^{t} \alpha_s$, the sampling process in which the model learns to remove the Gaussian noise and predict x_{t-1} from a given latent variable x_t can be formulated as Equation 1.

$$x_{t-1} = \frac{1}{\sqrt{\alpha_t}}(x_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}}\epsilon_\theta(x_t, t)) + \sigma_t \mathbf{z} \quad (1)$$

Here $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$, $\sigma_t = \eta\sqrt{\beta_t(1-\bar{\alpha}_{t-1})/(1-\bar{\alpha}_t)}$ is the standard deviation. η is the level of stochasticity which is set to 1 here. ϵ_θ denotes the diffusion model parameterized by θ which is trained by minimizing the variational lower bound $\mathbb{E}_{q(x_{0:T})}\log(q(x_{1:T}|x_0) \,/\, p_\theta(x_{0:T}))$ via the density gradient loss $\mathcal{L}_d$ as shown in Equation 2.

$$\mathcal{L}_d = \mathbb{E}_{t,\epsilon}\left[\frac{\beta_t^2}{2\sigma_t^2\alpha_t(1-\bar{\alpha}_t)}||\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}x_0 + \sqrt{1-\bar{\alpha}_t}\epsilon, t)||_2^2\right] \quad (2)$$

3.2. Timestep Selection for Semantic Control

With TAIGen, we want to create a robust attack that will be visually similar to the original image whilst being computationally less expensive to create. Nonetheless, because the diffusion model evaluations are inherently deterministic, there is a fixed mapping from these noises to the final samples. Recent studies such as [3] has shown that some noises are better than others in sampling. We thus choose a small interval($t << T$) of time steps during the sampling process instead of the entire T steps. Additionally, this also increases the robustness when compared to a single timestep attack.

Small Interval over single-step: As pointed out by [60], a mixing step in the reverse process is defined as a key step(t_{mixing}) in which the Gaussian distribution of x_T converges to the final distribution of the input data in the denoising process under some pre-defined distance measures. Instead of resorting to manual checking of the quality of the final denoised images, we compute the radius of the latent encodings following the generative trajectory. Specifically, given a d-dimensional Gaussian centered at the origin with variance σ^2, for a point $x = (x_1, x_2, \ldots, x_d)$ chosen at random from Gaussian, the radius is the square root of the expected squared length of x:

$$\sqrt{E(x_1^2 + x_2^2 + \cdots + x_d^2)} = \sqrt{dE(x_1^2)} = \sqrt{d}\sigma \quad (3)$$

The earliest step with the largest radius shift $\Delta r \approx 4$ is an approximation of the mixing step. By manipulating the latent space in this step, we can gain significant control over the distribution of the reconstructed sample. We utilize this mixing step to reduce the number of required sampling steps during adversarial image generation. Instead of using all the sampling steps during the reverse process, we only perform our proposed algorithm on the mixing step. However, empirically, we find that instead of focusing on a single step, a small interval of steps generates a more robust adversarial image. The generated image has the same image quality as compared to a single-step attack. Therefore, we choose the interval(N) to be a hyperparameter in our experiments with $N \ll T$ with T being the total number of timesteps in the sampling process.

From all steps to small interval: Figure 4 can help understand the reduction of sampling steps to a small interval. The point where the red line diverges from the blue line till the maxima is the size N. Contrary to previous beliefs [31], the asymmetric nature of the stochastic denoising and deterministic denoising process has a big advantage, i.e.,

Algorithm 1 Adversarial Image Sampling

for $t \in [T, T-1, \ldots, 1]$ **do**
  $\epsilon_t, w_{t-1} \leftarrow \epsilon_\theta(\hat{x}_t, t)$
  $\hat{x}_{t-1} = \frac{1}{\sqrt{\bar{\alpha}_t}}\left(x_t - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}}\epsilon_t\right) + \sigma_t z$
  if $(t \leq t_{start})$ & $(t \geq t_{end})$ **then**
    $\epsilon_t \leftarrow \epsilon_\theta(\hat{x}_t^{adv}, t)$
    $z_0 \leftarrow \frac{1}{\sqrt{\bar{\alpha}_t}}\left(x_t - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}}\epsilon_t\right) + \sigma_t z$
    $G_{t-1} \leftarrow 1(M > \Omega)$
    $W_{t-1} \leftarrow 1(w_{t-1} > \Phi)$
    $C_{t-1} \leftarrow W_{t-1} \oplus G_{t-1} \oplus G_{t-1}$
    for $i \in [0, 1, 2, ...I-1]$ **do**
      $g_{i+1} = \mu g_i + \frac{\nabla_z J(z_i, y)}{\|\nabla_z J(z_i, y)\|_1}$
      $z_{i+1} = z_i + \alpha \cdot \text{Sign}(g_{i+1})$
    end for
    return z_I
  end if
  $\hat{x}_{t-1}^{adv} \leftarrow C_{t-1} \odot \hat{x}_{t-1}^{adv} + (1 - C_{t-1}) \odot z_I$
  if $t == t_{end}$ **then**
    $\hat{x}_t = \hat{x}_t^{adv}$
  end if
end for
return $\hat{x}_0$

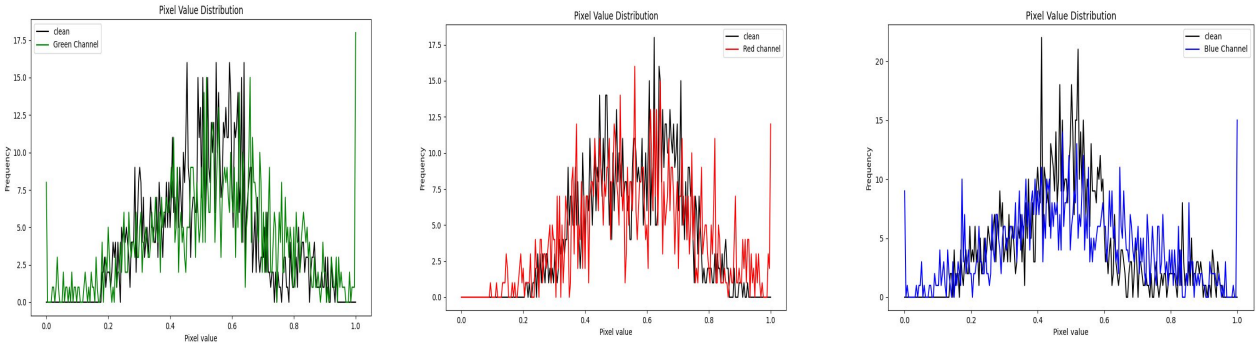

Figure 3. Distribution of pixel values when attention maps are used in (a) Left: *Green* channel, (b) Middle: *Red* channel, (c) Right: *Blue* channel. The black line indicates the distribution for the clean image, while the colored one is for the attacked image. Compared to others, variation from clean to attacked pixels is the least in the red channel.

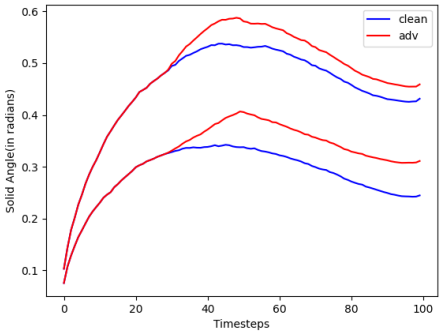

Figure 4. Two samples from the CIFAR-10 dataset have been used. The Y-axis represents the solid angle between the latent variable x_t and $\hat{x}_t$ during the forward and backward process respectively measured using the arcos function. The blue and red lines represent the convergence of original and adversarial reconstructions, respectively, highlighting the asymmetry and mixing step approximation.

control of the latent variable x_t in this interval via the ϵ-space of the model can significantly control the semantics of the reconstructed image. We can also think of it as a perturbing latent variable near the region where it differs the most (hence maxima) from the actual inverted latent variable at that timestep. Therefore, we utilize this property to reduce the time needed to generate a sample, unlike methods that use almost all T timesteps in the worst-case scenario [9, 33]. Note that we find t_{start} and t_{end} empirically and these could change with the change in T. In any scenario however, t_{start} can be found by calculating the radius of the latent variable. We provide a more detailed formulation of Figure 4 in the supplementary material section.

3.3. Attention Maps: Leveraging h-space of DDPM Sampler

The role of h-space has been shown to be critical in determining the semantics of the resultant image [31] along with proving that it is solely responsible for controlling the semantic space [26] of the images along with the skip connections. To effectively use this high-dimensional information, we extract the attention map: $W \in \mathcal{R}^{B\times HW\times HW}$, HW being the number of tokens, from the self-attention module in the h-space or the middle-block of the U-Net architecture, which produces $\epsilon_\theta(x_t, t)$ as it's output. We observe that if we consider all the attention maps from the encoder and decoder, our visual quality remains almost the same whilst the robust accuracy increases. Moreover, [23] showed that self-attention guidance can improve the stability and enhance the quality of the reconstructed image. Hence, we choose only to use the attention map from the middle block as it suffices for the majority of the latent space information.

3.4. Adversarial Classifier Guidance

The neurons in the convolutional layers look for high-level semantics and detailed spatial information. Grad-CAM uses the gradient information flowing into the last convolutional layer of the CNN to assign importance values to each neuron for a particular decision of interest. In order to obtain the class-discriminative localization map Grad-CAM, $M \in \mathbb{R}^{h\times w}$, we first compute the gradient of the score for class c, y^c (before the softmax), with respect to feature map activations A^k of a convolutional layer, i.e. $\frac{\delta y^c}{\delta A^k}$. Using this, we calculate the neuron importance weight α_k^c via:

$$\alpha_k^c = \underbrace{\frac{1}{Z}\sum_i\sum_j}_{} \underbrace{\frac{\partial y^c}{\partial A_{ij}^k}}_{} \tag{4}$$

Model	Attack	Resnet	EfficientNet	GoogleLeNet	MNASNet	MobileNet	SuffleNet	SqueezeNet	VGG
ResNet	PGD	**100.0**	36.6	40.4	43.4	40.2	50.6	52.6	46.6
	AutoAttack	**100.0**	34.6	40.2	42.8	40.8	50.0	54.0	48.4
	AdaMSI-FGM	**100.0**	51.2	55.4	54.6	46.8	55.2	65.0	62.4
	TAIGen (Ours)	68.8	**62.0**	**60.6**	**87.4**	**61.2**	**94.2**	**96.41**	**91.6**
GoogleLeNet	PGD	38.8	33.8	**100.0**	39.6	41.2	48.6	49.8	43.2
	AutoAttack	38.8	33.6	**100.0**	41.6	41.2	49.2	53.8	44.8
	AdaMSI-FGM	50.0	42.8	**100.0**	49.0	45.4	53.6	60.2	54.2
	TAIGen (Ours)	**58.4**	**49.6**	50.0	**70.21**	**51.2**	**90.4**	**91.8**	**82.8**
MobileNet	PGD	35.6	36.6	35.6	53.4	**100.0**	50.4	51.8	42.0
	AutoAttack	33.6	32.6	35.0	51.6	**100.0**	50.8	49.8	40.8
	AdaMSI-FGM	43.0	48.0	43.4	67.0	**100.0**	56.4	59.4	51.0
	TAIGen (Ours)	**64.2**	**62.4**	**57.0**	**81.8**	73.4	**92.99**	**94.6**	**89.0**
VGG	PGD	43.8	38.4	39.2	46.8	43.0	49.4	59.0	**100.0**
	AutoAttack	42.4	37.0	39.8	46.0	41.6	48.6	58.6	**100.0**
	AdaMSI-FGM	58.2	53.2	52.2	58.6	48.6	55.2	70.0	**100.0**
	TAIGen (Ours)	**70.6**	**66.4**	**61.6**	**80.8**	**65.8**	**95.0**	**97.81**	94.8

Table 1. ASR (%) of adversarial attacks against eight models on the ImageNet dataset. The row depicts the source models, while the column depicts the target models.

Attribute	Method	Clean Acc	ASR
Eyeglasses	BPDA+EOT	100	91.6
	TAIGen (Ours)	100	**100**
Smiling	BPDA+EOT	97.46	99.41
	TAIGen (Ours)	97.46	**100**
Gender	BPDA+EOT	100	**100**
	TAIGen (Ours)	100	**100**

Table 2. Comparison of BPDA+EOT attack on CelebA-HQ under l_∞ bound and $\epsilon = 8/255$.

Furthermore, we perform a weighted combination of forward activation maps, and follow it by a ReLU to obtain:

$$M = \text{ReLU}\left(\sum_k \alpha_k^c A^k\right) \tag{5}$$

Inspired by [9], we use the gradient weighted class activation mapping (GradCAM) from the last convolution layer of each source classifier f corresponding to an input sample x_0. However, instead of producing the GradCAM(M) corresponding to the true label y relating x_0, we randomly choose another label $\hat{y}$ and localize the specific regions relating to that class in the feature space of the GradCAM. Another method of doing the same could be to negate the importance weight α_k^c. However, this could lead to confusion for the classifier about which regions are contextually important to a class and which are just negative space. This lead us to choose a random label $\hat{y}$. We also want to add adversarial noise to the latent image due to which the perspective of the momentum-boosted iterative method [16], which helps to escape local minima and stabilize the update direction, was introduced. This is done by approximating an adversarial example x_I, starting from a reference image x_0 as formulated in Equation 6

$$\begin{aligned} g_{t+1} &= \mu \cdot g_t + \frac{(\nabla_x J(f(x_t^{adv}), y)}{||\nabla_x J(f(x_t^{adv}), y)||_1} \\ x_{t+1}^{adv} &= x_t^{adv} + \alpha \cdot sign(g_{t+1}) \end{aligned} \tag{6}$$

Here, J denotes the cross-entropy loss used for C-class classification problems, μ is the momentum factor and $\|\cdot\|_1$ is the l_1 norm. In our experiments, we make our adversarial examples satisfy the l_∞ bound by setting $\alpha = \epsilon/I$ with I being the number of iterations. Momentum, as observed empirically in our experiments, helped escape any local minima and reach the saddle point solution for minimizing the negative log-likelihood $\log p_\theta(\hat{x}_t)$ whilst maximizing the adversarial loss $J(f(x_t^{adv}), y)$. However, as mentioned by [49] it might escape saddle points too. We therefore use a combination of masked GradCAM (G_t) and masked attention maps (W_t) to counter it using Equation 7

$$C_t = W_t \oplus G_t \oplus G_t \tag{7}$$

where $\oplus$ denotes the concatenation operation along the RGB channel. By doing so, the range of values of $(z_I - \hat{x}_{t-1}^{adv})$ are scaled by a factor of $(1 - C_{t-1})$. This provides two advantages: (1) $(1 - W_t)$ helps increase adversarial loss $J(f(x_t^{adv}), y)$ due to it's inherent masking of relevant pixels and also pushes the sample away from clean sample convergence route during sampling and (2)$(1 - G_t)$ provides the correct gradients of the classifer relating to that sample for minimizing $\log p_\theta(\hat{x}_t)$. Moreover, to make the $W_t \in \mathbb{R}^{B\times HW\times HW}$ match in dimensions with the GradCAM($G_t \in \mathbb{R}^{B\times h\times w}$), which is the same as the input sample, we used linear interpolation on W_t to match it's dimensions with G_t. As discussed in earlier sections,

Attack Variant	Interval	ASR(↑)	FID(↓)	PSNR(↑)	SSIM(↑)
Clean	-	-	0.09	32.20	88.66
$l_\infty(\epsilon = 4/255)$	27-22	**99.7**	0.62	28.44	82.61
$l_\infty(\epsilon = 4/255)$	21-17	99.6	0.55	29.19	83.93
$l_\infty(\epsilon = 4/255)$	20-18	77.8	0.20	30.83	87
$l_\infty(\epsilon = 8/255)$	19	37.8	**0.15**	**31.53**	**87.56**

Table 3. Comparison of image quality and attack success rate with different time step intervals on the CelebA-HQ dataset with $T = 100$. The results on the clean images is also reported for understanding the efficacy of our approach.

Model	Classifier Acc.	Clean Acc.	Epsilon	ASR (w/o early stopping)	ASR (w early stopping)
WideResNet-28-10	95.53	84.76	4/255	35.28	**65.44**
			8/255	63.55	**97.47**
WideResNet-70-16	95.53	84.76	4/255	31.09	**57.81**
			8/255	61.48	**88.21**

Table 4. Comparison of different attack strengths against variations of WideResNet classifiers on the CIFAR-10 dataset. All the experiments are in the white-box setting for both with and without early stopping.

TAIGen utilizes attention maps and GradCAM during the generation of adversarial images. This utilization is done carefully by selectively modifying the RGB channels. Modifying all channels with GradCAM reduced the image quality and modifying all of them with the attention map loses the randomization introduced using the GradCAM. As observed from Figure 3, the combination of the two has two benefits. We can see that the variation of the pixel values in the red channel is the least from the target pixel values. Thus, we use W_t on the Red channel as it preserves the quality. The saturation rate is better in this case, which implies that the l_∞ norm is well-satisfied, leading to higher subtlety and transferability across different black-box models. Hence, the artifacts are removed and the semantic organization is maintained in $\hat{x_t}$.

3.5. Combing the Components: The Proposed Algorithm

As discussed in earlier sections, TAIGen utilizes attention maps and GradCAM during the generation of adversarial images over the chosen N interval of time steps. Hence, we have an iterative TAIGen attack as seen in Algorithm 1. We use the GradCAM M of a random non-objective class $\hat{y}$ to put more weight to the attacked regions which are present in z_i and the attention maps to help the model remember the actual regions of interest for proper reconstruction. Thus, we use the G_t twice and W_t once. Figure 2 shows a visual demonstration of our algorithm and Algorithm 1 gives an overview of it.

4. Experiments and Results

We conduct experiments across multiple datasets and models to show the performance of our method. The experimental settings, results, and observations are detailed as follows.

4.1. Experimental Setting

4.1.1. Datasets

To evaluate the proposed TAIGen algorithm, we evaluate it on the CelebA-HQ [28], CIFAR-10 [30] and ImageNet [14] datasets. CIFAR-10 is a widely known image dataset that contains $60,000$ images of resolution 32×32. We consider 512 images from the validation set ($10,000$ images) over 3 different seeds following the protocol maintained by [27]. CelebA-HQ is a high-quality image dataset that contains $30,000$ images of resolution 1024×1024 which was made based on the CelebA dataset. We randomly select 512 images from this dataset. ImageNet is another widely known dataset that contains images of resolution 224×224. Following the protocol maintained by recent attacks [34], we only consider randomly chosen 512 images or 1000 images from the validation set of ILRVSC 2012 [42] for experiments.

4.1.2. Models

We use different models for our experiments. For experiments with ImageNet, we use ResNet-34 [20], EfficientNet-b0 [52], GoogleLeNet [51], MobileNet-small [24], and VGG-11 [46] as source classifiers. While MNASNet-0-5 [53], ShuffleNet-v2-x0-5 [35], and SqueezeNet-1-1 [25] are the target models. For CIFAR-10. we use the variants of the WideResNet [58] model.

4.1.3. Hyperparameters and Implementation Details

For CIFAR-10 and ImageNet, we consider $T = 100$ and $N = 20$. For CelebA-HQ, we consider $T = 50$ and $N = 3$. We use ϵ values as $4/255$, $8/255$, and $16/255$ to compare

Method	Robust Acc.
SPSA*	81.29
Square Attack*	81.68
Joint Attack (Full)*	76.26
Diff-BPDA*	75.00
AutoAttack*	70.64
TAIGen (Ours)	**65.10**

Table 5. DDPM-based purification on CIFAR-10 dataset when compared with other attack methods on WideResNet-28-10 under l_∞ norm and $\epsilon = 8/255$.

between different attack strengths. Additional details are available in the supplementary material. All the experiments were carried out on a single 32 GB NVIDIA-V100 GPU.

4.1.4. Evaluation Metrics

In the case of CIFAR-10, we use the Robust Accuracy metric against the Diffpure [39] purification technique. For CelebA-HQ and ImageNet, we use the Attack Success Rate (ASR) metric and report the top-1 accuracy. To evaluate the image quality of our generated samples, we also use the Peak-Signal-to-Noise Ratio (PSNR), Structural Similarity (SSIM), and Frechet Inception Distance (FID) metrics.

4.2. Results and Analysis

Transferability with ImageNet: We evaluate results for TAIGen along with other attack algorithms including PGD [36], AutoAttack [11], and AdaMSI-FGM [34]. We report the ASR in Table 1. We observed that our method has a higher ASR rate than the other considered methods in almost all the cases for the black-box setting. While for the white-box setting, we do not perform as well. We conjecture that this is due to the fact that our algorithm balances all the aspects equally. Therefore, although we are using the gradients of the source classifier in the interval N, they are being integrated in a non-trivial manner in x_t such that it can successfully confuse the other target classifiers into misclassification while in case of same target classifer as the source, some semantics are preserved in the GradCAM which guides the classifer. We also note that with the increase in the parameter space of the model, like in VGG-11, TAIGen performs very well. This is because small variations in the parameters of these models can significantly affect the prediction.

Imperceptibility with CelebA-HQ: To evaluate our attack on face recognition systems, we use the CelebA-HQ dataset and BPDA+EOT attack [21] to compare our algorithm against as reported in Table 2. We use the attribute prediction task to evaluate our results with target attributes selected as *Eyeglasses*, *Smiling*, and *Gender*. We note that TAIGen performs better than the considered attack method. Further, we evaluate the image quality of our attack over different time intervals on the attribute-wise robust target classifier proposed by [8]. We report ASR along with the image quality metrics in Table 3. We use only the gender-attribute prediction task for this experiment. As shown by the ASR metric, $t_{mixing} = 19$ is an effective mixing step and our comparison with a single-step version of our attack. It sufficiently brings down the accuracy while maintaining the semantics of the image to almost near perfection. As we increase the interval size around the mixing step, we see the accuracy reduce even further with almost 100% success rate with just $N = 5$ steps around the mixing step. The second row shows us that as long as we are around the mixing step, we will still have a high success rate. Note that the mixing step is approximated via the method explained earlier and may vary a bit for different batches of samples. For intrusion in biometric systems, this procedure can be extended to similar datasets.

Attack	Time(in seconds)
ACA*	125.33
TAIGen (Ours)(w/o early stopping)	12
TAIGen (Ours)(w early stopping)	**6.6**

Table 6. Comparison of ACA attack which creates adversarial images using the ImageNet dataset with MobileNet-V2 as the source and target classifier. $*$ indicates that the values have been taken from their respective papers in the reported tables.

Robustness to Purification techniques: Purification methods have been used to evaluate the robustness of several attack methods. To compare the robustness of TAIGen, we evaluate it on the DiffPure [39] purification method. We compare it against SPSA [54], Square attack [4], Diff-BPDA[5], Joint attack[57], and AutoAttack [11] l_∞ threat model. SPSA attack approximates the gradients by randomly sampling from a predefined distribution and using the finite-difference method. Square attack heuristically searches for adversarial examples in a low-dimensional space with the constraints of perturbation patterns. Joint attack leverages the classifier gradients and the difference between the input and the purified samples. The results are reported in Table 5, where our method outperforms all the considered methods in terms of robust accuracy.

4.3. Ablation and Analysis

Early Stopping: Using early stopping with our method significantly reduces generation time, as shown in Table 6, outperforming the ACA attack [10]. On the CIFAR-10 dataset, with WideResNet-28-10 and WideResNet-70-16 as

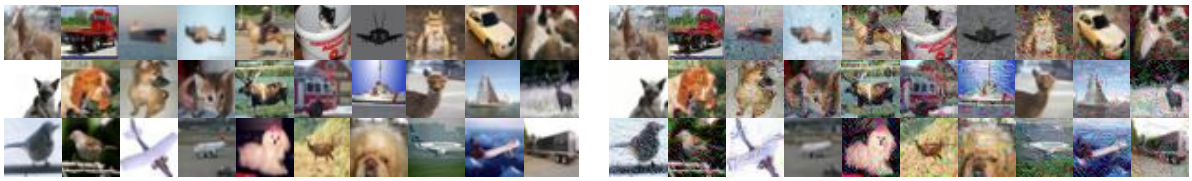

Figure 5. A set of few original (left) and attacked images (right) of the CIFAR-10 dataset attacked using TAIGen.

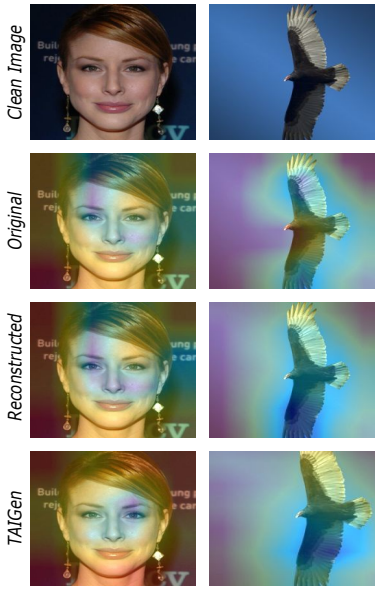

Figure 6. Heatmaps produced by the classifiers on the CelebA-HQ (first column) and ImageNet (second column) datasets. The second row shows the GradCAM of the original image produced by the classifier, while the third and fourth rows show GradCAM of the DDPM reconstructed image and the TAIGen-based DDPM reconstructed image.

robust classifiers, early stopping increases ASR by $\sim 30\%$ (Table 4). This demonstrates that early stopping effectively reduces high-frequency noise due to less cumulative scaling of the latent variable x_t whilst speeding up generation. Further details can be found in the supplementary material.

GradCAM Analysis: We pull the weights from the last convolution layer of the attribute classifier in the case of CelebA-HQ and Resnet-34 in the case of ImageNet to create the GradCAMs. We observe that TAIGen avoids focusing on regions sensitive to gender classification and localizes areas completely irrelevant to the classification, as seen in Figure 6. A major reason for the shift in semantic localization can be attributed to the masking operation we are performing on the GradCAM and attention maps in our algorithm. This is because we are attending to each of the channels disproportionately which leads to improper scaling by the coefficient of $\epsilon_\theta(x_t, t)$ across the timesteps. Another important point is that in the white-box setting, our method can easily displace points of interest when there are fewer important features. However, when the complexity of the feature of an image increases, as in the case of ImageNet, W_t masks many regions, which prevents the addition of perturbations to them.

Visualization: Figures 5 and 6 show the visualizations of our attacked images. In the case of CIFAR-10, we keep $l_\infty = 8/255$ to provide the worst-case scenario of our method in very low-resolution images. However, most biometric datasets such as face and iris are of high resolution and we notice that we obtain high-quality adversarial images in high-resolution images, thereby increasing applicability. Additionally, our method does not add any unnecessary features, such as a change in shape or texture, to the original image. This restricts the diffusion models to add unwanted features to the reconstructed images by maintaining the sanctity of the h-space through the attention maps.

5. Conclusion

Our proposed TAIGen method is a black-box, few-step, iterative diffusion-based attack method that selectively modifies the channels in the target image. Through experiments on different datasets and models, we show the robust performance of our attack in terms of performance, imperceptibility, and transferability. We observed that our method performs well in a black-box setting and thus has better transferability. The method generates high-quality adversarial images, specifically in biometric datasets. Although our method has some limitations, such as low performance in a white-box setting, we are able to generate a robust adversarial image that performs better than many state-of-the-art methods and showcases the effectiveness of diffusion models. In future work, we hope to improve the attack to withstand stronger purification techniques and address the issue of deterioration of quality in low-resolution images.

6. Acknowledgement

The authors acknowledge the support of IndiaAI and Meta through the Srijan: Centre of Excellence for Generative AI.

References

[1] Akshay Agarwal, Richa Singh, Mayank Vatsa, and Nalini Ratha. Are image-agnostic universal adversarial perturbations for face recognition difficult to detect? In *2018 IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pages 1–7, 2018. 2

[2] Akshay Agarwal, Gaurav Goswami, Mayank Vatsa, Richa Singh, and Nalini K. Ratha. Damad: Database, attack, and model agnostic adversarial perturbation detector. *IEEE Transactions on Neural Networks and Learning Systems*, 33 (8):3277–3289, 2022. 2

[3] Donghoon Ahn, Jiwon Kang, Sanghyun Lee, Jaewon Min, Minjae Kim, Wooseok Jang, Hyoungwon Cho, Sayak Paul, Seonhwa Kim, Eunju Cha, Kyong Hwan Jin, and Seungryong Kim. A noise is worth diffusion guidance. *ArXiv*, abs/2412.03895, 2024. 3

[4] Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion, and Matthias Hein. Square attack: A query-efficient black-box adversarial attack via random search. In *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII*, page 484–501, Berlin, Heidelberg, 2020. Springer-Verlag. 7

[5] Tsachi Blau, Roy Ganz, Bahjat Kawar, Alex M. Bronstein, and Michael Elad. Threat model-agnostic adversarial defense using diffusion models. *ArXiv*, abs/2207.08089, 2022. 7

[6] Hanqun Cao, Cheng Tan, Zhangyang Gao, Yilun Xu, Guangyong Chen, Pheng-Ann Heng, and Stan Z. Li. A survey on generative diffusion models. *IEEE Trans. Knowl. Data Eng.*, 36(7):2814–2830, 2024. 1

[7] Nicholas Carlini and David A. Wagner. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy, SP*, pages 39–57, 2017. 2

[8] Lucy Chai, Jun-Yan Zhu, Eli Shechtman, Phillip Isola, and Richard Zhang. Ensembling with deep generative views. In *CVPR*, 2021. 7

[9] Xinquan Chen, Xitong Gao, Juanjuan Zhao, Kejiang Ye, and Cheng-Zhong Xu. Advdiffuser: Natural adversarial example synthesis with diffusion models. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4539–4549, 2023. 2, 4, 5

[10] Zhaoyu Chen, Bo Li, Shuang Wu, Kaixun Jiang, Shouhong Ding, and Wenqiang Zhang. Content-based unrestricted adversarial attack. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2023. Curran Associates Inc. 7

[11] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *Proceedings of the 37th International Conference on Machine Learning*. JMLR.org, 2020. 7

[12] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(9):10850–10869, 2023. 1

[13] Xuelong Dai, Kaisheng Liang, and Bin Xiao. Advdiff: Generating unrestricted adversarial examples using diffusion models. In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part XLVI*, page 93–109, Berlin, Heidelberg, 2024. Springer-Verlag. 2

[14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 6

[15] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, pages 8780–8794, 2021. 1, 12

[16] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 9185–9193, 2018. 2, 5

[17] Logan Engstrom, Andrew Ilyas, Hadi Salman, Shibani Santurkar, and Dimitris Tsipras. Robustness (python library), 2019. 13

[18] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *3rd International Conference on Learning Representations, ICLR*, 2015. 2

[19] Gaurav Goswami, Nalini Ratha, Akshay Agarwal, Richa Singh, and Mayank Vatsa. Unravelling robustness of deep learning based face recognition against adversarial attacks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), 2018. 2

[20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 6, 13

[21] Mitch Hill, Jonathan Craig Mitchell, and Song-Chun Zhu. Stochastic security: Adversarial defense using long-run dynamics of energy-based models. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. 7

[22] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2020. Curran Associates Inc. 3, 12

[23] Susung Hong, Gyuseong Lee, Wooseok Jang, and Seungryong Kim. Improving sample quality of diffusion models using self-attention guidance. In *Proceedings - 2023 IEEE/CVF International Conference on Computer Vision, ICCV 2023*, pages 7428–7437. Institute of Electrical and Electronics Engineers Inc., 2023. Publisher Copyright: © 2023 IEEE.; 2023 IEEE/CVF International Conference on Computer Vision, ICCV 2023 ; Conference date: 02-10-2023 Through 06-10-2023. 4

[24] Andrew G. Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, Quoc V. Le, and Hartwig Adam. Searching for mobilenetv3. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1314–1324, 2019. 6

[25] Forrest N. Iandola, Song Han, Matthew W. Moskewicz, Khalid Ashraf, William J. Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <0.5mb model size. *arXiv:1602.07360*, 2016. 6

[26] Jaeseok Jeong, Mingi Kwon, and Youngjung Uh. Training-free content injection using h-space in diffusion models. *2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 5139–5149, 2023. 4

[27] Mintong Kang, Dawn Song, and Bo Li. Diffattack: Evasion attacks against diffusion-based adversarial purification. In *NeurIPS*, 2023. 1, 2, 6

[28] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *CoRR*, abs/1710.10196, 2017. 6

[29] Jernej Kos, Ian Fischer, and Dawn Song. Adversarial examples for generative models. In *2018 IEEE Security and Privacy Workshops, SP Workshops*, pages 36–42, 2018. 2

[30] Alex Krizhevsky. Learning multiple layers of features from tiny images. pages 32–33, 2009. 6

[31] Mingi Kwon, Jaeseok Jeong, and Youngjung Uh. Diffusion models already have a semantic latent space. In *The Eleventh International Conference on Learning Representations*, 2023. 1, 3, 4

[32] Qian Li, Yuxiao Hu, Ye Liu, Dongxiao Zhang, Xin Jin, and Yuntian Chen. Discrete point-wise attack is not enough: Generalized manifold adversarial attack for face recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR*, pages 20575–20584, 2023. 2

[33] Decheng Liu, Xijun Wang, Chunlei Peng, Nannan Wang, Ruimin Hu, and Xinbo Gao. Adv-diffusion: Imperceptible adversarial face identity attack via latent diffusion model. In *AAAI, IAAI, EAAI*, pages 3585–3593, 2024. 1, 2, 4

[34] Sheng Long, Wei Tao, Shuohao Li, Jun Lei, and Jun Zhang. On the convergence of an adaptive momentum method for adversarial attacks. In *AAAI*, pages 14132–14140, 2024. 2, 6, 7

[35] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Computer Vision – ECCV 2018: 15th European Conference, Munich, Germany, September 8–14, 2018, Proceedings, Part XIV*, page 122–138, Berlin, Heidelberg, 2018. Springer-Verlag. 6

[36] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018. 2, 7

[37] Yichuan Mo, Dongxian Wu, Yifei Wang, Yiwen Guo, and Yisen Wang. When adversarial training meets vision transformers: Recipes from training to architecture. In *NeurIPS*, 2022. 13

[38] Dongbin Na, Sangwoo Ji, and Jong Kim. Unrestricted black-box adversarial attack using GAN with limited queries. In *ECCV Workshops*, pages 467–482, 2022. 2

[39] Weili Nie, Brandon Guo, Yujia Huang, Chaowei Xiao, Arash Vahdat, and Animashree Anandkumar. Diffusion models for adversarial purification. In *ICML*, pages 16805–16827. PMLR, 2022. 2, 7

[40] Weili Nie, Brandon Guo, Yujia Huang, Chaowei Xiao, Arash Vahdat, and Anima Anandkumar. Diffusion models for adversarial purification. In *International Conference on Machine Learning (ICML)*, 2022. 12

[41] Omid Poursaeed, Tianxing Jiang, Harry Yang, Serge J. Belongie, and Ser-Nam Lim. Robustness and generalization via generative adversarial training. In *IEEE/CVF ICCV*, pages 15691–15700, 2021. 2

[42] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115 (3):211–252, 2015. 6

[43] Hadi Salman, Andrew Ilyas, Logan Engstrom, Ashish Kapoor, and Aleksander Madry. Do adversarially robust imagenet models transfer better? *CoRR*, abs/2007.08489, 2020. 12, 13

[44] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *IEEE ICCV*, pages 618–626, 2017. 1, 2

[45] Ali Shahin Shamsabadi, Ricardo Sánchez-Matilla, and Andrea Cavallaro. Colorfool: Semantic adversarial colorization. In *IEEE/CVF CVPR*, pages 1148–1157, 2020. 2

[46] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. 6

[47] Kaiyu Song, Hanjiang Lai, Yan Pan, and Jian Yin. Mimicdiffusion: Purifying adversarial perturbation via mimicking clean diffusion model. In *IEEE/CVF (CVPR)*, pages 24665–24674, 2024. 2

[48] Yang Song, Rui Shu, Nate Kushman, and Stefano Ermon. Constructing unrestricted adversarial examples with generative models. In *NeurIPS*, pages 8322–8333, 2018. 2

[49] Tao Sun, Dongsheng Li, Zhe Quan, Hao Jiang, Shengguo Li, and Yong Dou. Heavy-ball algorithms always escape saddle points. In *International Joint Conference on Artificial Intelligence*, 2019. 5

[50] Yuhao Sun, Lingyun Yu, Hongtao Xie, Jiaming Li, and Yongdong Zhang. Diffam: Diffusion-based adversarial makeup transfer for facial privacy protection. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 24584–24594, 2024. 2

[51] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott E. Reed, Dragomir Anguelov, D. Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, 2014. 6

[52] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *ArXiv*, abs/1905.11946, 2019. 6

[53] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, and Quoc V. Le. Mnasnet: Platform-aware neural architecture search for mobile. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2815–2823, 2018. 6

[54] Jonathan Uesato, Brendan O'Donoghue, Pushmeet Kohli, and Aäron van den Oord. Adversarial risk and the dangers of evaluating against weak attacks. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, pages 5032–5041. PMLR, 2018. 7

[55] Xiaosen Wang, Kun He, and John E. Hopcroft. AT-GAN: A generative attack model for adversarial transferring on generative adversarial nets. *CoRR*, abs/1904.07793, 2019. 2

[56] Haotian Xue, Alexandre Araujo, Bin Hu, and Yongxin Chen. Diffusion-based adversarial sample generation for improved stealthiness and controllability. In *NeurIPS*, 2023. 1

[57] Jongmin Yoon, {Sung Ju} Hwang, and Juho Lee. Adversarial purification with score-based generative models. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021*, pages 12062–12072. ML Research Press, 2021. Publisher Copyright: Copyright © 2021 by the author(s); 38th International Conference on Machine Learning, ICML 2021 ; Conference date: 18-07-2021 Through 24-07-2021. 7

[58] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016. 6

[59] Dawei Zhou, Nannan Wang, Chunlei Peng, Xinbo Gao, Xiaoyu Wang, Jun Yu, and Tongliang Liu. Removing adversarial noise in class activation feature space. In *IEEE/CVF ICCV*, pages 7858–7867, 2021. 2

[60] Ye Zhu, Yu Wu, Zhiwei Deng, Olga Russakovsky, and Yan Yan. Boundary guided learning-free semantic control with diffusion models. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2023. 1, 3, 13

TAIGen: Training-Free Adversarial Image Generation via Diffusion Models

Supplementary Material

7. Experimental Setting

In addition to the details presented in the main paper, the additional experimental settings are detailed below.

7.1. Hyperparameter Details:

The values of Ω and Φ have been set to 0.9. Additionally, we used linear beta time scheduling (β_t) with the standard DDPM [22] configurations in all the cases.

ImageNet: We used the Guided Diffusion model [15] from OpenAI on this dataset. The batch size was set to 4. We used 512 images, each of size 224×224. The values of t_{start} and t_{end} here are 80 and 60 respectively. We also kept $\epsilon = 4/255$ when comparing against PGD, AutoAttack, and AdaMSI-FGM. While testing against ACA, we used the MobileNet-V2 as the source and target model and 1000 images were chosen at random. The reported time is in seconds/image. The momentum factor was set to 1.2, and I was fixed at 20 iterations in both of these experiments. We perform a 10-class classification and randomly choose a class from the top-5 classes next in line to the true label class for creating the Grad-CAM.

CelebA-HQ: Whilst testing against BPDA+EOT, we kept $\epsilon = 8/255$, $t_{start} = 20$ and $t_{end} = 18$ and batch size as 4. Since there are only two classes in this case, we choose the class that does not correspond to the true label in this case for creating the Grad-CAM.

CIFAR-10: We set batch size as 8, $I = 50$, 512 as the number of images, $t_{start} = 80$ and $t_{end} = 55$. We also fixed $\epsilon = 4/255$ when comparing the robust accuracy against SPSA, Square Attack, Joint Attack (Full), Diff-BPDA and Auto-Attack. The momentum factor is 1.8. We perform a 1000-class classification and choose a class randomly from the top-500 classes, which are next in line to the true label class for creating the Grad-CAM.

7.2. Additional Experimental Results:

To compare the image quality against the two variations (that is, with and without early stopping) of our algorithm, we randomly choose 512 images from the validation set of CIFAR-10 and report the PSNR and SSIM values in Table 7. We also showcase our algorithm with early stopping in Algorithm 2. We observe that image quality is retained more with early stopping. However, one limitation of using this is that the images become more susceptible to purification algorithms like Diffpure [40]. This is because, with early stopping, the images are not as robust as those that have undergone more iterations of adding adversarial noise. Thus, they are easier to purify. Since we present a black-box attack, we use the WideResNet-50-2 classifier [43] as the

	PSNR	SSIM
Clean	24.16	81.55
w/o early stopping	22.42	75.39
w early stopping	**23.28**	**77.12**

Table 7. PSNR and SSIM values on the CIFAR-10 dataset. The model used is the standard WideResNet-28-10 with $\epsilon = 4/255$.

Algorithm 2 Adversarial Image Generation with Early Stopping

function TAIGEN(Input Image x_0, Noise Schedule $\beta_{1:T}$, Target Classifier f, Grad-CAM g_{CAM}, Diffusion U-Net Model ϵ_θ, Ground Truth y, Adversarial Iterations I, Momentum Factor μ, Step Size α, Cross-Entropy loss J)
  Ensure: Adversarial Image x_{adv}
  Initialize: $x_T \leftarrow x_0 \cdot \sqrt{\bar{\alpha}_T} + \sqrt{1-\bar{\alpha}_T} \cdot \epsilon, \epsilon \sim \mathcal{N}(0, I)$
  $\hat{x}_t^{adv} \leftarrow \hat{x}_{t_{start}}$, $\sigma_t^2 \leftarrow \beta_t$, $z \sim \mathcal{N}(0, I)$, $N = t_{start} - t_{end}$
  for $t \in [T, T-1, \ldots, 1]$ **do**
    $\epsilon_t, w_{t-1} \leftarrow \epsilon_\theta(\hat{x}_t, t)$
    $\hat{x}_{t-1} = \frac{1}{\sqrt{\bar{\alpha}_t}} \left(x_t - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}} \epsilon_t \right) + \sigma_t z$
    $x_0 \leftarrow \frac{1}{\sqrt{\bar{\alpha}_t}} (x_t - \sqrt{1-\bar{\alpha}_t} \epsilon_t)$
    if $\arg\max f(x_0) \neq y$ **then**
      $\hat{x}_0 = x_0$
      break
    else if $(t \leq t_{start})$ & $(t \geq t_{end})$ **then**
      $\epsilon_t \leftarrow \epsilon_\theta(\hat{x}_t^{adv}, t)$
      $z_0 \leftarrow \frac{1}{\sqrt{\bar{\alpha}_t}} \left(x_t - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}} \epsilon_t \right) + \sigma_t z$
      $G_{t-1} \leftarrow 1(g_{CAM} > \Omega)$
      $W_{t-1} \leftarrow 1(w_{t-1} > \Phi)$
      $C_{t-1} \leftarrow (W_{t-1} \oplus G_{t-1} \oplus G_{t-1})$
      for $i \in [0, 1, 2, \ldots I-1]$ **do**
        $g_{i+1} = \mu g_i + \frac{\nabla_z J(z_i, y)}{\|\nabla_z J(z_i, y)\|_1}$
        $z_{i+1} = z_i + \alpha \cdot \text{Sign}(g_{i+1})$
      end for
      return z_I
    end if
    $\hat{x}_{t-1}^{adv} \leftarrow C_{t-1} \odot \hat{x}_{t-1}^{adv} + (1 - C_{t-1}) \odot z_I$
    if $t == t_{end}$ **then**
      $\hat{x}_t = \hat{x}_t^{adv}$
    end if
  end for
  return $\hat{x}_0$
end function

Model	Clean Accuracy (%)	ASR (%)
R50 [20]	76.52	96.10
R50 [43]	64.02	57.23
WRN-50-2 [43]	68.46	54.11
R50 [17]	62.56	57.43
ViT-B [37]	68.38	49.22

Table 8. Accuracy and ASR on the ImageNet dataset with $\epsilon = 4/255$ using the WideResNet-50-2 as the source model.

source model and test the efficacy on various adversarially trained target models, ResNet-50 [17], ResNet-50[43] and WideResNet-50-2 [43], ViT-B [37] and a standard ResNet-50 model [20]. The values for this experiment are reported in Table 8.

8. Finding Mixing Step

[60] has explained in detail about the mixing step. We give a brief explanation of the same. We first define the radius of a high-dimensional Gaussian space. Mathematically, it can defined as $r = \sigma\sqrt{d}$. Now let us take a point in this vector space, $x = (x_1, x_2, \ldots, x_d)$, chosen at random from Gaussian, the square root of the expected square length of x is formulated in Equation 8.

$$\sqrt{\mathrm{E}(x_1^2 + x_2^2 + \cdots + x_d^2)} = \sqrt{d\mathrm{E}(x_1^2)} = \sqrt{d}\sigma \quad (8)$$

Equation 8 is used to find the radius of our sampled latent variables at each time step. We also define the total variation distance which will be used in the further proofs in Equation 9.

$$||\mu - \tau||_{TV} = \frac{1}{2}\sum_{x\in\mathcal{X}} |\mu(x) - \tau(x)| \quad (9)$$

Here μ and τ are two probability distributions on $\mathcal{X}$. Next, we define the quantity $\sigma_t(x, y)$ for an irreducible transition matrix P with stationary distribution π in Equation 10 and $d^{(p)}$ distance in Equation 11.

$$\sigma_t(x, y) = \frac{P^t(x, y)}{\pi(y)} \quad (10)$$

$$d^p(t) := \max_{x\in\mathrm{X}} ||\sigma_t(x, .) - 1||_p \quad (11)$$

Replacing the above notations with the ones from a standard DDPM model, we get Equations 12 and 13.

$$d^1(t) := \max_{x\in\mathrm{X}} ||\sigma_t(x, .) - 1||_1 \quad (12)$$

and,

$$\sigma_t(x, y) = \frac{P^t(x, y)}{\pi(y)} = \frac{x \sim \mathcal{N}(x_t; \sqrt{1-\beta_t}x_{t-1}, \beta_t\mathrm{I})}{y \sim \mathcal{N}(0, \mathrm{I}_d)} \quad (13)$$

Mathematically, we can define the mixing time via Equation 14

$$t^{(1)}_{mix}(\epsilon) := \inf\{t \geq 0; d^1(t) \leq \epsilon\} \quad (14)$$

Taking $\epsilon = 1/2$, we get Equation 15

$$t^{(1)}_{mix}(\epsilon) := \inf\{t \geq 0; d^1(t) \leq \frac{1}{2}\} \quad (15)$$

Replacing Equation 15 with Equation 13, we get Equation 16.

$$\max_{x\in\mathcal{X}} || \frac{x \sim \mathcal{N}(x_t; \sqrt{1-\beta_t}x_{t-1}, \beta_t\mathrm{I})}{y \sim \mathcal{N}(0, \mathrm{I}_d)} - 1|| \leq \frac{1}{2} \quad (16)$$

Using Equation 9, we can substitute in Equation 16 which gives us the approximation in Equation 17.

$$||x \sim \mathcal{N}(x_t; \sqrt{1-\beta_t}x_{t-1}, \beta_t\mathrm{I})|| \leq 4 \quad (17)$$

Equation 17 searches for the mixing step where the Gaussian radius changes by an amount of **4** units.

9. Solid Angle using Quaternions

To find the similarity between the latent variable x_t during the forward process and $\hat{x}_t$ during the backward process, we use the cosine similarity between these two variables. Specifically, we consider the vector or quaternion q_1 and q_2 and the cosine similarity is found using Equation 18.

$$\Omega = \cos^{-1}\left(\frac{q_1}{||q_1||_2} \cdot \frac{q_2}{||q_2||_2}\right) \quad (18)$$

As observed from the graph (Figure 3 of the main paper), it forms a concave curve. Intuitively, we can say that while the Gaussian distribution doesn't converge with the stationary distribution, the solid angle or dissimilarity increases. At a certain point, in this case the maxima, they converge and then the similarity increases between the two vectors. Thus, via this analysis, it is evident that the inversion and the sampling processes are **not** symmetric. Now our intuition lies in utilizing this fact and perturbing the $\hat{x}_t$ around the time step corresponding to the maxima. If instead, we perturbed at any other time step, the dissimilarity would again increase till it reached a maximum (which would not be the mixing step) and then decrease. In the latter case, we empirically found the artifacts are visible in the reconstructed image.