AT-GAN: AN ADVERSARIAL GENERATIVE MODEL FOR NON-CONSTRAINED ADVERSARIAL EXAMPLES

Anonymous authors
Paper under double-blind review

ABSTRACT

With the rapid development of adversarial machine learning, numerous adversarial attack methods have been proposed. Typical attacks are based on a search in the neighborhood of input image to generate a perturbed adversarial example. Since 2017, generative models are adopted for adversarial attacks, and most of them focus on generating adversarial perturbations from input noise or input image. Thus the output is restricted by input for these works. A recent work targets "unrestricted adversarial example" using generative model but their method is based on a search in the neighborhood of input noise, so actually their output is still constrained by input. In this work, we propose AT-GAN (Adversarial Transfer on Generative Adversarial Net) to *train an adversarial generative model that can directly produce adversarial examples*. Different from previous works, we aim to *learn the distribution of adversarial examples* so as to generate semantically meaningful adversaries. AT-GAN achieves this goal by first learning a generative model for real data, followed by transfer learning to obtain the desired generative model. Once trained and transferred, AT-GAN could generate adversarial examples directly for any input noise, denoted as non-constrained adversarial examples. Extensive experiments and visualizations show that AT-GAN can efficiently generate diverse adversarial examples that are realistic to human perception, and yields higher attack success rates against adversarially trained models.

1 INTRODUCTION

In recent years, Deep Neural Networks (DNNs) have been found vulnerable to adversarial examples (Szegedy et al., 2014), which are well-crafted samples with tiny perturbations imperceptible to humans but can fool the learning models. Despite the great success of the deep learning empowered applications, many of them are safety-critical, for example under the scenario of self-driving cars (Eykholt et al., 2018; Cao et al., 2019), raising serious concerns in academy and industry.

Numerous works of adversarial examples have been developed on adversarial attacks (Goodfellow et al., 2015; Carlini & Wagner, 2017; Madry et al., 2018), adversarial defenses (Goodfellow et al., 2015; Kurakin et al., 2017; Song et al., 2019) and exploring the property of adversarial examples (He et al., 2018; Shamir et al., 2019). For adversarial attacks, most studies focus on the *perturbation-based adversarial examples* constrained by the input images, which is also the generally accepted conception of adversarial examples. Generative models are also adopted recently to generate adversarial perturbations from an input noise (Reddy Mopuri et al., 2018; Omid et al., 2018) or from a given image (Xiao et al., 2018; Bai et al., 2020), and such perturbations are added to the original image to craft adversarial examples. Song et al. (2018) propose to search a neighborhood noise around the input noise of a Generative Adversarial Net (GAN) (Goodfellow et al., 2014) such that the output is an adversarial example, which they denoted as *unrestricted adversarial example* as there is no original image in their method. However, their output is still constrained by the input noise, and their search is time-consuming.

In this work, we propose a generative attack framework called AT-GAN (Adversarial Transfer on Generative Adversarial Net), which aims to learn the distribution of adversarial examples. Unlike previous works that constrain the adversaries in the neighborhood of input image or input noise, or the unrestricted adversarial attack (Song et al., 2018) that searches adversarial noise around the input noise, AT-GAN is an adversarial generative model that could produce semantically meaningful

adversarial examples directly from any input noise, and we call such examples *non-constrained adversarial examples*.

Specifically, we first develop a normal GAN to learn the distribution of benign data so that it can produce plausible images that the classifier and a human oracle will classify in the same way. Then we transfer the pre-trained GAN into an adversarial GAN called AT-GAN that can fool the target classifier while being still well recognized by the human oracle. AT-GAN is a conditional GAN that has learned to estimate the distribution of adversarial examples for the target classifier, so AT-GAN can directly generate adversarial examples from any random noise, leading to high diversity and efficiency.

To evaluate the effectiveness and efficiency of our approach, we develop AT-GAN on three benchmark datasets, namely MNIST, Fashion-MNIST and CelebA, and apply typical defense methods to compare AT-GAN with existing search-based attacks. Empirical results show that the non-constrained adversarial examples generated by AT-GAN yield higher attack success rates, and state-of-the-art adversarially trained models exhibit little robustness against AT-GAN, indicating the high diversity of our adversaries. In addition, AT-GAN, as a generation-based adversarial attack, is more efficient than the search-based adversarial attacks.

2 PRELIMINARIES

In this section, we provide definitions on several types of adversarial examples and adversarial attacks, and give a brief overview of adversarial attacks using GAN. Other related works on typical adversarial attacks and defenses (Goodfellow et al., 2015; Madry et al., 2018; Tramèr et al., 2018), as well as some typical GANs (Goodfellow et al., 2014; Radford et al., 2016; Odena et al., 2017; Arjovsky et al., 2017; Gulrajani et al., 2017) are introduced in Appendix A.

2.1 DEFINITIONS ON ADVERSARIES

Let $\mathcal{X}$ be the set of all digital images under consideration for a learning task, $\mathcal{Y} \in \mathbb{R}$ be the output label space and $p_z \in \mathbb{R}^m$ be an arbitrary probability distribution (*e.g.* Gaussian distribution) where m is the dimension of p_z. A deep learning classifier $f : \mathcal{X} \to \mathcal{Y}$ takes an image $x \in \mathcal{X}$ and predicts its label $f(x)$. Suppose p_x and p_{adv} are the distributions of benign images and adversarial examples, respectively. Assume we have an oracle classifier $o : \mathcal{X} \to \mathcal{Y}$, which could always predict the correct label for any image $x \in \mathcal{X}$, we define several types of adversarial examples as follows.

For perturbation-based adversarial examples (Szegedy et al., 2014; Goodfellow et al., 2015; Moosavi-Dezfooli et al., 2016), tiny perturbations are added to the input images, which are imperceptible to humans but can cause the target classifier to make wrong predictions.

Definition 1. ***Perturbation-based Adversarial Examples.*** Given a subset (trainset or testset) images $\mathcal{T} \subset \mathcal{X}$ and a small constant $\epsilon > 0$, the perturbation-based adversarial examples can be defined as: $\mathcal{A}_p = \{x_{adv} \in \mathcal{X} | \exists x \in \mathcal{T}, \|x - x_{adv}\|_p < \epsilon \wedge f(x_{adv}) \neq o(x_{adv}) = f(x) = o(x)\}$.

Song et al. (2018) define a new type of adversarial examples called unrestricted adversarial examples, which is not related to the subset (trainset or testset) images, by adding adversarial perturbation to the input noise of a mapping, such as GAN, so that the output of the perturbed noise is an adversary to the target classifier.

Definition 2. ***Unrestricted Adversarial Examples.*** Given a mapping G from $z \sim p_z$ to $G(z, y) \sim p_\theta$, where p_θ is an approximated distribution of p_x, and a small constant $\epsilon > 0$, the unrestricted adversarial examples can be defined as: $\mathcal{A}_u = \{G(z^*, y_s) \in \mathcal{X} | \exists z \sim p_z, z^* \sim p_z, \|z - z^*\|_p < \epsilon \wedge f(G(z^*, y_s)) \neq o(G(z^*, y_s)) = f(G(z, y_s)) = o(G(z, y_s)) = y_s\}$ where y_s is the source label.

In this work, we train a conditional GAN to learn the distribution of adversarial examples and output the corresponding adversary directly from any input noise. To clarify the difference with Song et al. (2018), we call our generated adversaries the non-constrained adversarial examples.

Definition 3. ***Non-constrained Adversarial Examples.*** If there is a mapping G^* from $z \sim p_z$ to $G^*(z, y) \sim q_\theta$, where q_θ is an approximated distribution of p_{adv}, the non-constrained adversarial examples can be defined as $\mathcal{A}_n = \{G^*(z, y_s) \in \mathcal{X} | f(G^*(z, y_s)) \neq o(G^*(z, y_s)) = y_s\}$ where y_s is the source label.

Here we need to find a mapping G^*, *e.g.* a generative model, such that for $z \sim p_z$, $G^*(z, y)$ is an image in $\mathcal{X}$ and the output distribution is an approximated distribution of p_{adv}, for example using the Kullback-Leibler divergence (Kullback & Leibler, 1951), $KL(q_\theta || p_{adv}) < \epsilon$ for a small constant ϵ.

In summary, *perturbation-based adversarial examples* are based on perturbing an image $x \in \mathcal{X}$, and *unrestricted adversarial examples* (Song et al., 2018) perturbs an input noise $z \sim p_z$ for an existing mapping G. Most perturbation-based adversarial attacks and Song et al. (2018) fall into the *search-based adversarial attack*.

Definition 4. ***Search-based Adversarial Attack.*** Given an input vector $v \in \mathcal{V}$ (either benign image x or random vector z), the search-based adversarial attack searches a vector $v' : \|v - v'\|_p < \epsilon$ where v' leads to an adversarial example for the target classifier.

In contrast, *non-constrained adversarial examples* are more generalized so that we need to learn a mapping G^* such that for any input noise sampled from distribution p_z, the output is an adversarial image. Such a mapping to be learned is called an Adversarial Generative Model (AGM), and our method falls into the *generation-based adversarial attack*.

Definition 5. ***Generation-based Adversarial Attack.*** Given an input vector $v \in \mathcal{V}$ (either benign image x or random vector z), the generation-based adversarial attack generates adversarial perturbation or adversarial example directly from v, usually adopting generative models.

2.2 Generative Models for Adversarial Attack

Generative models have been adopted for adversarial attack in recent works (Baluja & Fischer, 2017). Reddy Mopuri et al. (2018) propose a Network for Adversary Generation (NAG) that models the distribution of adversarial perturbations for a target classifier so that their NAG can craft adversarial perturbations from any given random vector, which will be added to the original image to fool the target classifier. Omid et al. (2018) propose to generate universal or image-dependent adversarial perturbations using U-Net (Ronneberger et al., 2015) or ResNet Generator (He et al., 2016) from any given random noise. Xiao et al. (2018) propose to train AdvGAN that takes an original image as the input and generate adversarial perturbation for the input to craft an adversarial example. Bai et al. (2020) further propose AI-GAN that adopts projected gradient descent (PGD) (Madry et al., 2018) in the training stage to train a GAN to generate target adversarial perturbation for the input image and target class. The above attack methods all fall into the generation-based adversarial attack, and their crafted examples fall into the perturbation-based adversarial examples.

Different from the above methods that generate adversarial perturbations, Song et al. (2018) propose to search a random vector z^* around the input noise vector z of AC-GAN (Odena et al., 2017) such that the corresponding output of AC-GAN is an adversarial example for the target classifier. Their method falls into the search-based adversarial attack, and their crafted examples fall into the unrestricted adversarial examples as there is no original image in their method.

3 AT-GAN: An Adversarial Generative Model

Here we first introduce the estimation on the distribution of adversarial examples, then propose the AT-GAN framework, a *generation-based adversarial attack* for crafting *non-constrained adversarial examples*. Further analysis is provided that AT-GAN could learn the adversary distribution.

3.1 Estimating the Adversarial Distribution

In order to generate non-constrained adversarial examples, we need to estimate the distribution of adversarial examples $p_{adv}(x_{adv}|y_{true})$ where y_{true} is the true label. Given the parameterized estimated distribution of adversarial examples $q_\theta(x|y_{true})$, we can define the estimation problem as:

$$q_{\theta^*}(x_{adv}|y_{true}) = \arg\min_{\theta \in \Omega} KL(q_\theta(x_{adv}|y_{true}) \| p_{adv}(x_{adv}|y_{true})), \tag{1}$$

where θ indicates trainable parameters and Ω is the parameter space.

It is hard to calculate equation 1 directly as $p_{adv}(x_{adv}|y_{true})$ is unknown. Inspired by the perturbation-based adversarial examples, as shown in Figure 1, we postulate that for each adversarial example x_{adv}, there exists some benign examples x where $\|x - x_{adv}\|_p < \epsilon$. In other words, $p_{adv}(x_{adv}|y_{true})$

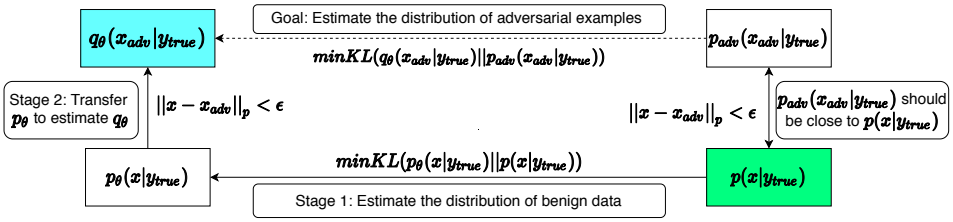

Figure 1: Estimating the distribution of adversarial examples q_θ in two stages: 1) estimate the distribution of benign data p_θ. 2) transfer p_θ to estimate q_θ.

is close to $p(x|y_{true})$ to some extent and we can obtain $p(x|y_{true})$ by Bayes' theorem, $p(x|y_{true}) = \frac{p(y_{true}|x)\cdot p(x)}{p(y_{true})}$, where $p(y_{true}|x)$, $p(x)$ and $p(y_{true})$ can be obtained directly from the trainset. Thus, we can approximately solve equation 1 in two stages: 1) Fit the distribution of benign data p_θ. 2) Transfer p_θ to estimate the distribution of adversarial examples q_θ.

Specifically, we propose an adversarial generative model called AT-GAN to learn the distribution of adversarial examples. The overall architecture of AT-GAN is illustrated in Figure 2. Corresponding to the above two stages, we first train the GAN model called AC-WGAN_GP, which combines AC-GAN (Odena et al., 2017) and WGAN_GP (Gulrajani et al., 2017) to get a generator $G_{original}$, to learn p_θ (See Appendix B). Then we transfer $G_{original}$ to attack the target classifier f for the learning of q_θ.

3.2 Transferring the Generator for Attack

After the original generator $G_{original}$ is trained, we transfer the generator $G_{original}$ to learn the distribution of adversarial examples in order to attack the target model. As illustrated in Figure 2 (b), there are three neural networks, including the original generator $G_{original}$, the attack generator G_{attack} to be transferred that is initialized by the weights of $G_{original}$, and the classifier f to be attacked. The goal of the second stage can be described as:

$$G^*_{attack} = \arg\min_{G_{attack}} ||G_{original}(z, y_s) - G_{attack}(z, y_s)||_p \quad s.\,t.\, f(G(z, y_s)) = y_t \neq y_s, \tag{2}$$

where y_t denotes the target label, $\|\cdot\|_p$ denotes the ℓ_p norm and we focus on $p = 2$ in this work.

To optimize equation 2, we construct the loss function by L_1 and L_2, where L_1 aims to assure that f yields the target label y_t that is fixed for target attack for each category:

$$L_1 = \mathbb{E}_{z\sim p_z}[H(f(G_{attack}(z, y_s)), y_t)]. \tag{3}$$

Here $H(\cdot,\cdot)$ denotes the cross entropy between the two terms and y_s is sampled from $\mathcal{Y}$. L_2 aims to assure that the adversarial generator G_{attack} generates realistic examples:

$$L_2 = \mathbb{E}_{z\sim p_z}[||G_{original}(z, y_s) + \rho - G_{attack}(z, y_s)||_p]. \tag{4}$$

Here ρ is a small uniform random noise constrained by both l_0 and l_∞ norm. We add ρ to constrain $G_{attack}(z, y_s)$ to be in the neighborhood of $G_{original}(z, y_s)$ rather than be exactly the same as $G_{original}(z, y_s)$.

The objective function for transferring $G_{original}$ to G_{attack} can be formulated as $L = \langle \alpha L_1, \beta L_2 \rangle$, where α and β are hyperparameters to control the training process. Note that in the case that $\alpha = 1$ and $\beta \to \infty$, the objective function is similar to that of the perturbation-based attacks (Goodfellow et al., 2015; Tramèr et al., 2018; Madry et al., 2018). For the untargeted attack, we can replace y_t in L_a with the maximum confidence of prediction label y except for y_s, $\max_{y\neq y_s} f(y|G_{attack}(z, y_s))$.

3.3 Theoretical Analysis on AT-GAN

This subsection provides theoretical analysis on why AT-GAN can generate as realistic and diverse non-constrained adversarial examples as real data. We will prove that under ideal condition, AT-GAN can estimate the distribution of adversarial examples, which is close to that of real data.

Suppose p_{data} is the distribution of real data, p_g and p_a are the distribution learned by the generator of AC-WGAN_GP and AT-GAN respectively. For the optimization of equation 4, L_2 aims to constrain the image generated by G_{attack} in the ϵ-neighborhood of $G_{original}$. We prove that under the ideal

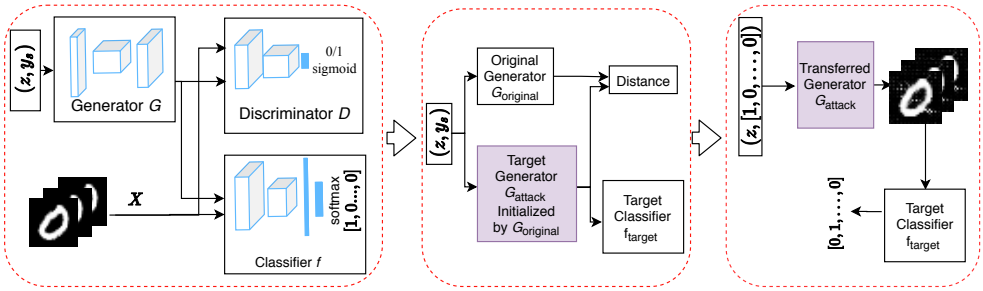

(a) AC-WGAN_GP (b) Transfer G for the Attack (c) Generate Adversarial Examples

Figure 2: The architecture of AT-GAN. The first training stage of AT-GAN is same as AC-WGAN_GP. Once trained, we regard G as the original model $G_{original}$ and transfer $G_{original}$ to attack target classifier to obtain G_{attack}. Once finished, AT-GAN can generate adversarial examples by G_{attack}.

condition that L_2 guarantees $G_{attack}(z, y_s)$ to be close enough to $G_{original}(z, y_s)$ for any input noise z, the distribution of AT-GAN almost coincides the distribution of AC-WGAN_GP. Formally, we state our result for the two distributions as follows.

Theorem 1. Suppose $\max_{z,y} L_2 < \epsilon$, we have $KL(p_a \| p_g) \to 0$ when $\epsilon \to 0$.

The proof of Theorem 1 is in Appendix C. Samangouei et al. (2018) prove that the global optimum of WGAN is $p_g = p_{data}$ and we show that the optimum of AC-WGAN_GP has the same property. We formalize the property as follows.

Theorem 2. The global minimum of the virtual training of AC-WGAN_GP is achieved if and only if $p_g = p_{data}$.

The proof of Theorem 2 is in Appendix C. According to Theorem 1 and 2, under the ideal condition, we conclude $p_a \approx p_g = p_{data}$, which indicates that the distribution of non-constrained adversarial examples learned by AT-GAN is very close to that of real data as discussed in Section 3.1, so that the non-constrained adversarial instances are as realistic and diverse as the real data.

4 Experiments

To validate the effectiveness and efficiency of the proposed approach, we empirically evaluate AT-GAN on various datasets, and demonstrate that AT-GAN yields higher attack success rates against adversarially trained models with higher efficiency. Besides, AT-GAN can learn a distribution of adversarial examples which is close to the real data distribution, and generate realistic and diverse adversarial examples.

4.1 Experimental Setup

Datasets. We consider three standard datasets, namely MNIST (LeCun et al., 1989), Fashion-MNIST (Xiao et al., 2017) and CelebA (Liu et al., 2015). MNIST is a dataset of hand written digits from 0 to 9. Fashion-MNIST is similar to MNIST with 10 categories of fashion clothes. CelebA contains more than $200,000$ celebrity faces. We group them according to female/male and focus on gender classification as in Song et al. (2018). For all datasets, we normalize the pixel values into range $[0, 1]$.

Baselines. We compare AT-GAN with the search-based attack methods, including Song's (Song et al., 2018) for unrestricted adversarial examples, as well as FGSM (Goodfellow et al., 2015), PGD (Madry et al., 2018) and R+FGSM (Tramèr et al., 2018) for perturbation-based adversarial examples. Note that although the perturbation-based results are not directly comparable to ours as they are limited to small perturbations on real images, they can provide a good sense on the model robustness.

Models. For MNIST and Fashion-MNIST, we adopt four models used in Tramèr et al. (2018), denoted as Model A to D. For CelebA, we consider three models, namely CNN, VGG16 and ResNet. More details about the architectures are provided in Appendix D.

Evaluation Setup. We consider normal training and existing strong defenses, namely adversarial training (Goodfellow et al., 2015), ensemble adversarial training (Tramèr et al., 2018) and iterative adversarial training (Madry et al., 2018). All experiments are conducted on a single Titan X GPU and the hyper-parameters are described in Appendix D.

Table 1: Attack success rate (ASR, %) of adversarial examples generated by AT-GAN and the baseline attacks against models by normal training and various adversarial training methods. For each model, the highest ASR is highlighted in **bold**. Notation: Nor. — Normal training, Adv. — Adversarial training, Ens. — Ensemble adversarial training, Iter. Adv. — Iterative adversarial training.

(a) Comparison of AT-GAN and other attacks on MNIST.

Attack	Model A				Model B				Model C				Model D			
	Nor.	Adv.	Ens.	Iter.	Nor.	Adv.	Ens.	Iter.	Nor.	Adv.	Ens.	Iter.	Nor.	Adv.	Ens.	Iter.
FGSM	68.8	2.9	8.7	4.0	88.0	9.5	18.0	9.2	70.7	4.3	7.8	4.7	89.6	23.5	34.6	26.7
PGD	**100.0**	92.6	85.1	5.9	**100.0**	42.4	98.2	36.2	**100.0**	76.9	96.4	9.6	91.7	96.8	**99.5**	81.1
R+FGSM	77.9	28.4	20.3	2.6	96.5	7.1	42.1	4.6	81.6	7.2	19.9	2.9	93.8	76.2	51.6	25.3
Song's	82.0	70.5	75.0	84.6	76.9	65.0	72.0	80.7	74.2	75.6	72.6	87.8	67.7	43.6	56.3	44.5
AT-GAN	98.7	**97.5**	**96.7**	**91.4**	99.5	**97.7**	**99.3**	**95.6**	99.3	**95.8**	**96.9**	**90.0**	**99.9**	**99.9**	**99.5**	**99.7**

(b) Comparison of AT-GAN and other attacks on Fashion-MNIST.

Attack	Model A				Model B				Model C				Model D			
	Nor.	Adv.	Ens.	Iter.	Nor.	Adv.	Ens.	Iter.	Nor.	Adv.	Ens.	Iter.	Nor.	Adv.	Ens.	Iter.
FGSM	68.8	2.9	8.7	4.0	88.0	9.5	18.0	9.2	70.7	4.3	7.8	4.7	89.6	23.5	34.6	26.7
PGD	**100.0**	92.6	85.1	5.9	**100.0**	42.4	98.2	36.2	**100.0**	76.9	96.4	9.6	91.7	96.8	**99.5**	81.1
R+FGSM	77.9	28.4	20.3	2.6	96.5	7.1	42.1	4.6	81.6	7.2	19.9	2.9	93.8	76.2	51.6	25.3
Song's	82.0	70.5	75.0	84.6	76.9	65.0	72.0	80.7	74.2	75.6	72.6	87.8	67.7	43.6	56.3	44.5
AT-GAN	98.7	**97.5**	**96.7**	**91.4**	99.5	**97.7**	**99.3**	**95.6**	99.3	**95.8**	**96.9**	**90.0**	**99.9**	**99.9**	**99.5**	**99.7**

(c) Comparison of AT-GAN and other attacks on CelebA.

Attack	CNN				VGG				ResNet			
	Nor.	Adv.	Ens.	Iter.	Nor.	Adv.	Ens.	Iter.	Nor.	Adv.	Ens.	Iter.
FGSM	81.2	11.8	14.7	9.5	76.7	10.6	16.2	8.9	98.7	9.8	12.3	9.5
PGD	97.3	16.4	22.6	11.4	87.9	14.6	26.3	10.7	**100.0**	10.9	15.1	10.5
R+FGSM	68.7	9.5	11.3	7.9	68.4	8.7	13.2	7.3	97.5	7.9	9.5	7.8
Song's	90.6	83.4	85.7	89.8	**98.7**	87.5	95.7	81.6	99.2	93.4	91.0	90.6
AT-GAN	**97.5**	**98.9**	**95.9**	**99.6**	97.1	**96.7**	**95.8**	**97.8**	98.4	**98.5**	**97.3**	**98.5**

4.2 Evaluation Results

4.2.1 Comparison on Attack Success Rate

To validate the attack effectiveness, we compare AT-GAN with the baselines under white-box setting. Since Athalye et al. (2018) have shown that the currently most effective defense method is adversarial training, we consider adversarially trained models as the defense models. The attack success rates are reported in Table 1.

On MNIST, AT-GAN achieves the highest Attack Success Rate (ASR) against the baselines on all defense models. As for normal training, AT-GAN achieves the highest ASR on Model D, and the second highest ASR of over 98% on the other models. On Fashion-MNIST, AT-GAN achieves the highest ASR on average. On CelebA, AT-GAN achieves the highest ASR almost on all models, with two exceptions under normal training but the results of AT-GAN are close to the highest.

In general, we see that AT-GAN achieves the highest attack performance above 90% on all the defense models. As AT-GAN aims to estimate the distribution of adversarial examples, adversarial training with some specific attacks has little robustness against AT-GAN, raising a new security issue for the development of more generalized adversarial training models.

4.2.2 Comparison on Attack Efficiency

There are many scenarios where one needs a large number of adversarial examples, such as adversarial training or exploring the property of adversarial examples. Thus, the efficiency for generating adversarial examples is important, but ignored in most previous works.

As an adversarial generative model, once trained, AT-GAN can generate adversarial examples very quickly. Here we evaluate the efficiency of each attack method for Model A on MNIST. The average time of generating/searching 1000 adversarial examples is summarized in Table 2. Among the five attack methods, AT-GAN is the fastest as it could craft adversarial examples without target classifier and gradient calculation. Note that Song's needs much longer time than others as it needs multiple searches and queries to generate one adversarial example. Here we only focus on the efficiency of generating adversarial examples after AT-GAN is transferred, *i.e.* we have already found the generator G^*, as in such case we could generate as many adversarial examples as we need.

Table 2: Comparison on the average example generation time, measured by generating 1000 adversarial instances using Model A on MNIST.

	FGSM	PGD	R+FGSM	Song's	AT-GAN
Time	0.3s	1.8s	0.4s	$\geq 2.5min$	0.2s

(a) Adversarial examples for Model A on MNIST

(b) Adversarial examples for Model A on Fashion-MNIST

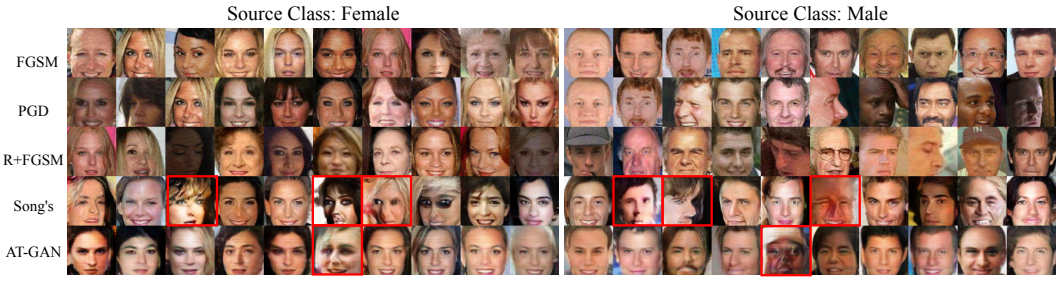

(c) Adversarial examples for CNN on CelebA

Figure 3: Adversarial examples generated by various attacks on three datasets (Zoom in for details). The red borders indicate unrealistic adversarial examples generated by Song's method or AT-GAN.

4.2.3 Visualization of Adversarial Examples

Since the goal of adversarial examples is to fool target neural networks but not to fool human oracle, in Figure 3 we illustrate some adversarial examples generated by different attacks for Modle A on MNIST and Fashion-MNIST, and CNN on CelebA.

On MNIST, AT-GAN generates slightly more realistic images than Song's, *e.g.* "0" and "3". On Fashion-MNIST and CelebA, some adversarial examples generated by Song's method are not as realistic as AT-GAN to human perception, for example "t-shirt/top (0) ", "sandal (5)" and some facial details. Note that Song's tends to distort the foreground that makes the images on MNIST more clean but some images are not realistic while AT-GAN tends to distort the background. As for perturbation-based attacks, their adversarial examples are not clear enough, especially on MNIST and Fashion-MNIST, due to the adversarial perturbations. There are also some unnatural samples generated by AT-GAN due to the limitation of GAN and we hope some better generative models can solve such issue. See more examples crafted by AT-GAN for target attack in Appendix D.

In general, AT-GAN can generate realistic and diverse adversarial examples as equation 1 forces the generated non-constrained adversarial examples to be close to the benign examples generated by the original generator.

4.3 Visualization of Adversarial Distribution

As discussed in Section 3.3, we provide a brief analysis that AT-GAN can learn a distribution of adversarial examples close to the distribution of real image data. To identify it empirically, we randomly choose $5,000$ benign images and $5,000$ adversarial examples generated by different attack methods, and merge these images according to their real label for MNIST and Fashion-MNIST. Then we use t-SNE (Maaten & Hinton, 2008) on these images to illustrate the distributions in 2 dimensions. t-SNE models each high-dimensional object in such a way that similar objects are modeled by nearby

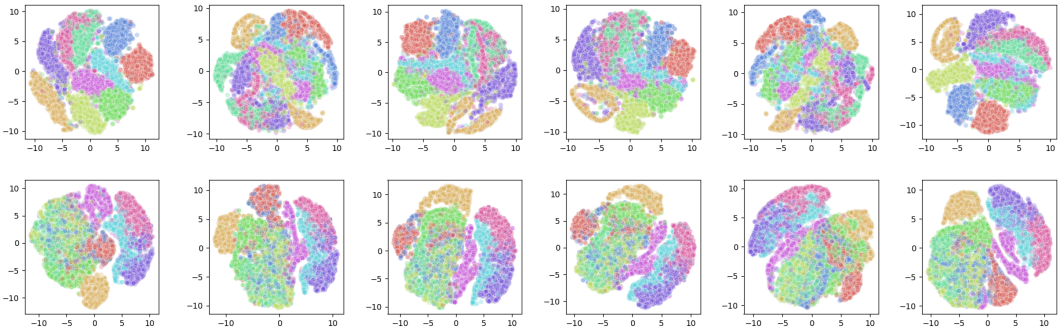

(a) Testset (b) FGSM (c) PGD (d) R+FGSM (e) Song (f) AT-GAN

Figure 4: T-SNE visualization for the combination of testset and adversarial examples generated by various attacks on MNIST (top) and Fashion-MNIST (bottom). For (a), we use total testset with 10,000 images. For (b) to (f), we use 5,000 sampled images in testset and 5,000 adversarial examples generated by various attacks. The position of each class is random due to the property of t-SNE.

points and dissimilar objects are modeled by distant points with high probability. It indicates that, if the adversarial examples have different distribution to the benign data, t-SNE could not deal with them well and the points with different categories will overlap with each other after the dimension reduction, *i.e*. the results will be in chaos.

The results are illustrated in Figure 4. For AT-GAN, different categories are separated as that of the test set while those of other methods are mixed with each other, especially on MNIST (top). It indicates the distribution AT-GAN learned is indeed very close to the distribution of real data.

To further validate that AT-GAN learns a different distribution from the original GAN rather than just adding some constant universal perturbation vector. In Appendix E, we illustrate some instances generated by the original generator and AT-GAN for the same input. We find that for different inputs, the original generator outputs different images and the difference between the instances generated by the original generator and AT-GAN is also different, indicating that AT-GAN indeed learns a different distribution from the original GAN.

5 Conclusion

In this work, we propose a generation-based adversarial attack method, called AT-GAN (Adversarial Transfer on Generative Adversarial Net), that aims to learn the distribution of adversarial examples for the target classifier. The generated adversaries are "non-constrained" as we do no search at all in the neighborhood of the input, and once trained AT-GAN can output adversarial examples directly for any input noise drawn from arbitrary distribution (e.g. Gaussian distribution). Extensive experiments and visualizations show that AT-GAN achieves highest attack success rates against adversarially trained models and can generate diverse and realistic adversarial examples efficiently.

Our work also suggests that adversarial training, a popular defense method based on perturbation-based adversarial examples, could not guarantee robustness against non-constrained adversarial examples. A possible reason is that AT-GAN learns a more complete version of the adversarial example distribution, which is much more diverse than that of the perturbation-based method.

Note that the generation quality of adversarial examples by AT-GAN depends on the generation quality of the base GAN model used for the transferring. Thus, if in the future the base GAN could generate more realistic images on complex dataset, AT-GAN will be able to generate higher quality examples drawn from the adversarial image distribution. Our method also suggests a new way of adversarial attack by designing an adversarial generative model directly. There are several other interesting questions related to our work that can be explored in future work. For instance, what is the distribution of adversarial examples really like? Is it a continuous or smooth manifold? How close could we learn such distribution through GAN? We hope our work could inspire more researches in this direction.

REFERENCES

Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein GAN. *arXiv Preprint arXiv:1701.07875*, 2017.

Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated Gradients Give a False Sense of Security: Circumventing Defenses to Adversarial Examples. *International Conference on Machine Learning*, 80:274–283, 2018.

Tao Bai, Jun Zhao, Jinlin Zhu, Shoudong Han, Jiefeng Chen, and Bo Li. AI-GAN: Attack-inspired generation of adversarial examples. *arXiv preprint arXiv:2002.02196*, 2020.

Shumeet Baluja and Ian Fischer. Adversarial transformation networks: Learning to generate adversarial examples. *arXiv preprint arXiv:1703.09387*, 2017.

Arjun Bhagoji, Warren He, Bo Li, and Dawn Song. Exploring the Space of Black-box Attacks on Deep Neural Networks. *arXiv Preprint arXiv:1712.09491*, 2017.

Jacob Buckman, Aurko Roy, Colin Raffel, and Ian Goodfellow. Thermometer Encoding: One Hot Way To Resist Adversarial Examples. *International Conference on Learning Representations*, 2018.

Yulong Cao, Chaowei Xiao, Dawei Yang, Jing Fang, Ruigang Yang, Mingyan Liu, and Bo Li. Adversarial Objects Against LiDAR-Based Autonomous Driving Systems. *arXiv Preprint arXiv:1907.05418*, 2019.

Nicholas Carlini and David Wagner. Towards Evaluating the Robustness of Neural Networks. *IEEE Symposium on Security and Privacy*, pp. 39–57, 2017.

Xi Chen, Yan Duan, Rein Houthooft, John Schulman, Ilya Sutskever, and Pieter Abbeel. InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets. *Neural Information Processing Systems*, pp. 2172–2180, 2016.

Kevin Eykholt, Ivan Evtimov, Earlence Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. Robust Physical-World Attacks on Deep Learning Models. *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Neural Information Processing Systems*, pp. 2672–2680, 2014.

Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and Harnessing Adversarial Examples. *International Conference on Learning Representations*, 2015.

Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved Training of Wasserstein GANs. *Neural Information Processing Systems*, pp. 5767–5777, 2017.

Chuan Guo, Mayank Rana, Moustapha Cisse, and Laurens Maaten. Countering Adversarial Images using Input Transformations. *International Conference on Learning Representations*, 2018.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.

Warren He, Bo Li, and Dawn Song. Decision Boundary Analysis of Adversarial Examples. *International Conference on Learning Representations*, 2018.

S. Kullback and R. A. Leibler. On information and sufficiency. *Ann. Math. Statist.*, 22(1):79–86, 1951.

Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial Machine Learning at Scale. *International Conference on Learning Representations*, 2017.

Yann LeCun, Bernhard Boser, John Denker, Donnie Henderson, Richard Howard, Wayne Hubbard, and Lawrence Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.

Yandong Li, Lijun Li, Liqiang Wang, Tong Zhang, and Boqing Gong. Nattack: Learning the distributions of adversarial examples for an improved black-box attack on deep neural networks. *International Conference on Machine Learning (ICML)*, 2019.

Fangzhou Liao, Ming Liang, Yinpeng Dong, Tianyu Pang, Xiaolin Hu, and Jun Zhu. Defense against Adversarial Attacks Using High-Level Representation Guided Denoiser. *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1778–1787, 2018.

Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into Transferable Adversarial Examples and Black-box Attacks. *International Conference on Learning Representations*, 2017.

Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep Learning Face Attributes in the Wild. *International Conference on Computer Vision*, pp. 3730–3738, 2015.

Laurens Maaten and Geoffrey Hinton. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605, 2008.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards Deep Learning Models Resistant to Adversarial Attacks. *International Conference on Learning Representations*, 2018.

Jan Metzen, Hendrik, Tim Genewein, Volker Fischer, and Bastian Bischoff. On Detecting Adversarial Perturbations. *International Conference on Learning Representations*, 2017.

Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *IEEE conference on computer vision and pattern recognition*, pp. 2574–2582, 2016.

Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional Image Synthesis With Auxiliary Classifier GANs. *International Conference on Learning Representations*, 2017.

Poursaeed Omid, Katsman Isay, Gao Bicheng, and Serge Belongie. Generative adversarial perturbations. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *International Conference on Learning Representations*, 2016.

Konda Reddy Mopuri, Utkarsh Ojha, Utsav Garg, and R Venkatesh Babu. NAG: Network for adversary generation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 742–751, 2018.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2015.

Pouya Samangouei, Maya Kabkab, and Rama Chellappa. Defense-GAN: Protecting Classifiers Against Adversarial Attacks Using Generative Models. *International Conference on Learning Representations*, 2018.

Adi Shamir, Itay Safran, Eyal Ronen, and Orr Dunkelman. A Simple Explanation for the Existence of Adversarial Examples with Small Hamming Distance. *arXiv Preprint arXiv:1801.02774*, 2019.

Shiwei Shen, Guoqing Jin, Ke Gao, and Yongdong Zhang. APE-GAN: Adversarial Perturbation Elimination with GAN. *IEEE International Conference on Acoustics,Speech and SP*, pp. 3842–3846, 2017.

Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. *International Conference on Learning Representations*, 2015.

Chuanbiao Song, Kun He, Liwei Wang, and John Hopcroft. Improving the Generalization of Adversarial Training with Domain Adaptation. *International Conference on Learning Representations*, 2019.

Yang Song, Rui Shu, Nate Kushman, and Stefano Ermon. Constructing unrestricted adversarial examples with generative models. *Neural Information Processing Systems*, pp. 8322–8333, 2018.

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *International Conference on Learning Representations*, 2014.

Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. Ensemble Adversarial Training: Attacks and Defenses. *International Conference on Learning Representations*, 2018.

Chaowei Xiao, Bo Li, Jun-Yan Zhu, Warren He, Mingyan Liu, and Dawn Song. Generating Adversarial Examples with Adversarial Networks. *International Joint Conferences on Artificial Intelligence*, pp. 3905–3911, 2018.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms. *arXiv preprint arXive:1708.07747*, 2017.