



# Lipschitz constrained GANs via boundedness and continuity

Kanglin Liu<sup>1,2,3</sup> · Guoping Qiu<sup>1,2,3,4</sup>

Received: 10 December 2019 / Accepted: 17 April 2020 / Published online: 24 May 2020  
© The Author(s) 2020

## Abstract

One of the challenges in the study of generative adversarial networks (GANs) is the difficulty of its performance control. Lipschitz constraint is essential in guaranteeing training stability for GANs. Although heuristic methods such as weight clipping, gradient penalty and spectral normalization have been proposed to enforce Lipschitz constraint, it is still difficult to achieve a solution that is both practically effective and theoretically provably satisfying a Lipschitz constraint. In this paper, we introduce the boundedness and continuity (BC) conditions to enforce the Lipschitz constraint on the discriminator functions of GANs. We prove theoretically that GANs with discriminators meeting the BC conditions satisfy the Lipschitz constraint. We present a practically very effective implementation of a GAN based on a convolutional neural network (CNN) by forcing the CNN to satisfy the BC conditions (BC-GAN). We show that as compared to recent techniques including gradient penalty and spectral normalization, BC-GANs have not only better performances but also lower computational complexity.

**Keywords** Generative adversarial networks · Lipschitz constraint · Boundedness · Continuity

## 1 Introduction

Generative adversarial networks (GANs) [5] are hailed as one of the most significant developments in machine learning research of the past decade. Since its first introduction, GANs have been applied to a wide range of problems and numerous papers have been published. In a nutshell, GANs are constructed around two functions [4, 11]: the generator  $G$ , which maps a sample  $z$  to the data distribution, and the discriminator  $D$ , which is trained to distinguish real samples of a dataset from fake samples produced by the generator. With the goal of reducing the

difference between the distributions of fake and real samples, a GAN training algorithm trains  $G$  and  $D$  in tandem.

A major challenge of GANs is that controlling the performance of the discriminator is particularly difficult. Kullback–Leibler (KL) divergence was originally used as the loss function of the discriminator to determine the difference between the model and target distributions [16]. However, KL divergence is potentially noncontinuous with respect to the parameters of  $G$ , leading to the difficulty in training [2, 23]. Specifically, when the support of the model distribution and the support of the target distribution are disjoint, there exists a discriminator that can perfectly distinguish the model distribution from that of the target. Once such a discriminator is found, zero gradients would be backpropagated to  $G$  and the training of  $G$  would come to a complete stop before obtaining the optimal results. Such a phenomenon is referred to as the vanishing gradient problem.

The conventional form of Lipschitz constraint is given by:  $\|f(x_1) - f(x_2)\| \leq k \cdot \|x_1 - x_2\|$ . It is obvious that Lipschitz constraint requires the continuity of the constrained function and guarantees the boundedness of the gradient norm. Besides, it has been found that enforcing Lipschitz constraint can provide provable robustness against adversarial examples [21], improve generalization

---

✉ Guoping Qiu  
Guoping.Qiu@nottingham.ac.uk

Kanglin Liu  
max\_liu@szu.edu.cn

<sup>1</sup> Shenzhen University, Shenzhen, China

<sup>2</sup> Guangdong Key Laboratory of Intelligent Information Processing, Shenzhen, China

<sup>3</sup> Shenzhen Institute of Artificial Intelligence and Robotics for Society, Shenzhen, China

<sup>4</sup> University of Nottingham, Nottingham, UK

bounds [19], enable Wasserstein distance estimation [6], and also alleviate the training difficulty in GANs. Thus, a number of works have advocated the Lipschitz constraint. To be specific, weight clipping was first introduced to enforce the Lipschitz constraint [2]. However, it has been found that weight clipping may lead to the capacity underuse problem where training favors a discriminator that uses only a few features [6]. To overcome the weakness of weight clipping, regularization terms like gradient penalty are added to the loss function to enforce Lipschitz constraint on  $D$  [6, 12, 15]. More recently, Miyato *et al.* [13] introduce spectral normalization to control the Lipschitz constraint of  $D$  by normalizing the weight matrix of the layers, which is regarded as an improvement on orthonormal regularization [18]. Using gradient penalty or spectral normalization can stabilize the training and gain-improved performance. However, it has been found that gradient penalty suffers from the problem of not being able to regularize the function at the points outside of the support of the current generative distribution [13]. In addition, spectral normalization has been found to suffer from the problem of gradient norm attenuation [1, 10], *i.e.*, a layer with a Lipschitz bound of 1 can reduce the norm of the gradient during backpropagation, and each step of backprop gradually attenuates the gradient norm, resulting in a much smaller Jacobian for the network's function than is theoretically allowed. Also as we will show in Sects. 3 and 4.3, these new methods have the capacity underuse problem (see Proposition 1 and Fig. 1). Therefore, despite recent progress, it remains challenging to achieve practical success as well as provably satisfying a Lipschitz constraint.

In this paper, we introduce the boundedness and continuity ( $BC$ ) conditions to enforce the Lipschitz constraint and introduce a CNN-based implementation of GANs with discriminators satisfying the  $BC$  conditions. We make the following contributions:

- (a) We prove that SN-GANs, one of the latest GAN training algorithms that use spectral normalization, will prevent the discriminator functions from obtaining the optimal solution when applying Wasserstein distance as the loss metric even though the Lipschitz constraint is satisfied.
- (b) We present  $BC$  conditions to enforce the Lipschitz constraint for the GANs' discriminator functions and introduce a CNN-based implementation of GANs by enforcing the  $BC$  conditions (BC-GANs). We show that the performances of BC-GANs are competitive to state-of-the-art algorithms such as SN-GAN and WGAN-GP but having lower computational complexity.

## 2 Related work

### 2.1 Generative adversarial networks (GANs)

Generative adversarial networks (GANs) are a special generative model to learn a generator  $G$  to capture the data distribution via an adversarial process. Specifically, a discriminator  $D$  is introduced to distinguish the generated images from the real ones, while the generator  $G$  is updated to confuse the discriminator. The adversarial process is formulated as a minimax game as:

$$\min_G \max_D V(G, D) \quad (1)$$

where min and max of  $G$  and  $D$  are taken over the set of the generator and discriminator functions, respectively.  $V(G, D)$  is to evaluate the difference in the two distributions of  $q_x$  and  $q_g$ , where  $q_x$  is the data distribution, and  $q_g$  is the generated distribution. The conventional form of  $V(G, D)$  is given by Kullback–Leibler (KL) divergence:  $E_{x \sim q_x} [\log D(x)] + E_{x' \sim q_g} [\log(1 - D(x'))]$  [16].

### 2.2 Methods to enforce Lipschitz constraint

Applying KL divergence as the implementation of  $V(G, D)$  could lead to the training difficulty, *e.g.*, the vanishing gradient problem. Thus, numerous methods have been introduced to solve this problem by enforcing the Lipschitz constraint, including weight clipping [2], gradient penalty [4] and spectral normalization [13].

Weight clipping was introduced by Wasserstein GAN (WGAN) [2], which used Wasserstein distance to measure the differences between real and fake distributions instead of KL divergence.

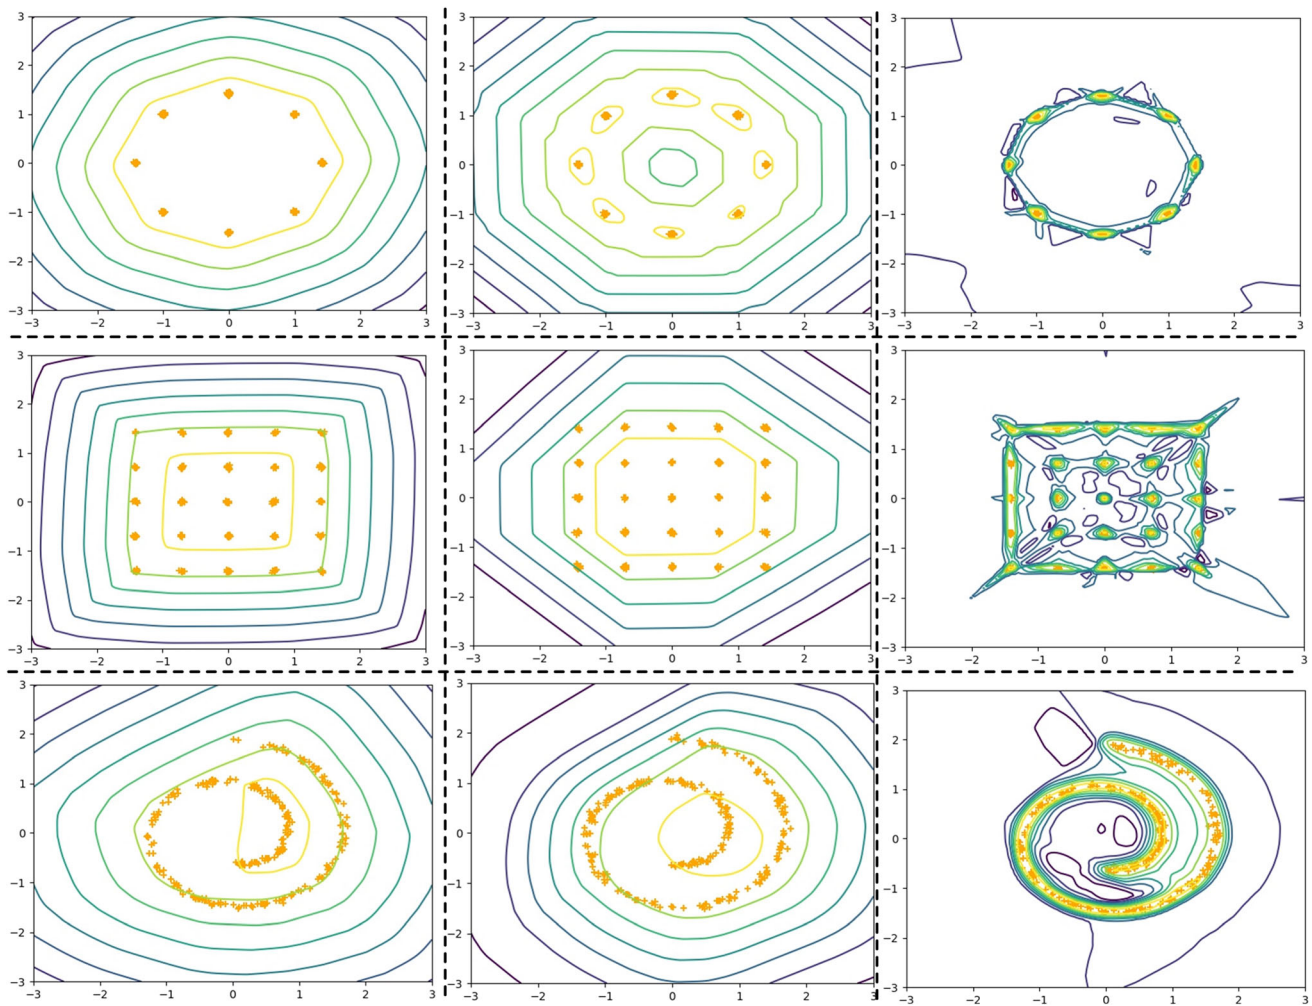
$$W(P_r, P_g) = \sup_{f \in Lip1} E_{x \sim P_r} [f(x)] - E_{x \sim P_g} [f(x)] \quad (2)$$

where  $W(P_r, P_g)$  represents the Wasserstein distance,  $P_r$  and  $P_g$  are the real and fake distributions, respectively. Weight clipping enforces the Lipschitz constraint by truncating each element of the weight matrices. Wasserstein distance shows superiority over KL divergence, because it can effectively avoid the vanishing gradient problem brought by KL divergence. In contrast to weight clipping, gradient penalty [6] penalizes the gradient at sample points to enforce Lipschitz constraint:

$$L_D = E[f(G(z))] - E[f(x)] + \underbrace{\alpha E[(\|\nabla f(x)\| - 1)^2]}_{\text{gradient penalty}} \quad (3)$$

where  $L_D$  is the loss objective for the discriminator, and  $\alpha$  is a hyperparameter.

Spectral normalization is a weight normalization method, which controls the Lipschitz constraint of the



**Fig. 1** Value surface of the discriminators trained to optimality on toy datasets. The yellow dots are data points, the lines are the value surfaces of the discriminators. Left column: spectral normalization. Middle column: gradient penalty. Right column: the proposed method. The upper, middle and lower rows are trained on 8-Gaussian,

25-Gaussian and the Swiss roll distributions, respectively. The generator is held fixed at real data plus unit-variance Gaussian noise. It is seen that discriminators trained with gradient penalty as well as spectral normalization have failed to capture the high moments of the data distribution (color figure online)

discriminator function by literally constraining the spectral norm of each layer. The implementation of the spectral normalization can be expressed as:

$$W_{SN}(W) := W/\sigma(W) \tag{4}$$

where  $W$  represents the weight matrix in each network layer,  $\sigma(W)$  is the spectral norm of matrix  $W$ , which equals to the largest singular value of the matrix  $W$ , and  $W_{SN}(W)$  represents the normalized weight matrix. To a certain extent, spectral normalization has succeeded in facilitating stable training and improving performance.

### 3 Existing problems

Although heuristic methods have been proposed to enforce Lipschitz constraint, it is still difficult to achieve a solution that is both practically effective and theoretically provably satisfying the Lipschitz constraint. To be specific, weight clipping was proven to be unsatisfactory in [4], and it can lead to the capacity underuse problem where training favors a discriminator that uses only a few features [6]. In addition, gradient penalty suffers from the obvious problem of not being able to regularize the function at the points outside of the support of the current generative distribution. In fact, the generative distribution and its support gradually change in the course of the training, and this can destabilize the effect of the regularization itself [13]. Moreover, it has been found that spectral normalization suffers from the

gradient norm attenuation problem [1, 10]. Furthermore, we have found that applying spectral normalization prevents the discriminator functions from obtaining the optimal solutions when using Wasserstein distance as the loss metric. To provide an explanation to this problem, we present Proposition 1.

Let  $P_r$  and  $P_g$  be the distributions of real images and generated images in  $X$ , a compact metric space. The discriminator function  $f$  is constructed based on a neural network of the following form with input  $x$ :

$$f(x, \theta) = W^{L+1} a_L(W^L(a_{L-1}(\dots a_1(W^1 x)))) \tag{5}$$

where  $\theta := \{W^1, W^2, \dots, W^{L+1}\}$  is the learning parameter set, and  $a_l$  is an element-wise nonlinear activation function. Spectral normalization is applied on  $f$  to guarantee the Lipschitz constraint.

**Proposition 1** *When using Wasserstein distance as the loss metric of  $f$ , the optimal solution to  $f$  is unreachable.*

### 4 Enforcing boundedness and continuity in CNN-based GANs

Finding a proper way to enforce the Lipschitz constraint remains an open problem. Motivated by this, we search for a better way to enforce the Lipschitz constraint.

#### 4.1 BC Conditions

The purpose is to find the discriminator from the set of  $k$ -Lipschitz continuous functions [7], which obeys the following condition:

$$|f(x_1) - f(x_2)| \leq k|x_1 - x_2| \tag{6}$$

Equation (6) is referred to as the Lipschitz continuity or Lipschitz constraint. If the discriminator function  $f$  satisfies the following conditions, it is guaranteed to meet the condition of Eq. (6):

- (a) Boundedness:  $f$  is a bounded function.
- (b) Continuity:  $f$  is a continuous function, and the number of points where  $f$  is continuous but not differentiable is finite. Besides, if  $f$  is differentiable at point  $x$ , its derivative is finite.

Conditions (a) and (b) are referred to as the boundedness and continuity (BC) conditions. A discriminator satisfying the BC conditions is referred as a bounded discriminator, and a GAN model with BC conditions enforced is referred to as BC-GAN. Following Theorems 1 and 2 guarantee that meeting the BC conditions is sufficient to enforce the Lipschitz constraint of Eq. (6) (see Proofs in “Appendix”)

**Theorem 1** *Let  $\Psi$  be the set of all  $f : X \rightarrow R$ , where  $f$  is a continuous function. In addition, the number of points where  $f$  is continuous but not differentiable is finite. Besides, if  $f$  is differentiable at point  $x$ , its derivative is finite. Then,  $f$  in  $\Psi$  satisfies Lipschitz constraint.*

**Theorem 2** *Let  $P_r$  and  $P_g$  be the distributions of real images and generated images in  $X$ , a compact metric space. Let  $\Omega$  be the set of all  $f : X \rightarrow R$ , where  $f$  is a continuous and bounded function. And, the number of points where  $f$  is continuous but not differentiable is finite. Besides, if  $f$  is differentiable at point  $x$ , its derivative is finite. The set  $\Omega$  can be expressed as:*

$$\Omega : \left\{ f \mid \|f(x)\| \leq m, \text{ if } \frac{\partial f(x)}{\partial x} \text{ exists, } \left\| \frac{\partial f(x)}{\partial x} \right\| < \infty \right\} \tag{7}$$

where  $m$  represents the bound. Then, there must exist a  $k$ , and we have a computable  $k \cdot W(P_r, P_g)$ :

$$k \cdot W(P_r, P_g) = \sup_{f \in \Omega} E_{x \sim P_r} [f(x)] - E_{x \sim P_g} [f(x)] \tag{8}$$

where  $W(P_r, P_g)$  represents the Wasserstein distance between  $P_r$  and  $P_g$  [5, 23].

According to Theorems 1 and 2, it is obvious that the BC conditions are sufficient to enforce the Lipschitz constraint. Furthermore,  $k \cdot W(P_r, P_g)$  is bounded and computable and can be obtained as:

$$k \cdot W(P_r, P_g) = \max_{f \in \Omega} E_{x \sim P_r} [f(x)] - E_{z \sim p(z)} [f(G(z))] \tag{9}$$

Then,  $k \cdot W(P_r, P_g)$  can be applied as a new loss metric to guide the training of  $D$ . Logically, the new objective for  $D$  is:

$$L_D = \min_{f \in \Omega} E_{z \sim p(z)} [f(G(z))] - E_{x \sim P_r} [f(x)] \tag{10}$$

Theorem 3 in [2] tells us that

$$\nabla_{\theta} k W(P_r, P_g) = -E_{z \sim p(z)} [\nabla_{\theta} f(G(z))] \tag{11}$$

where  $\theta$  is the parameters of  $G$ . Equation (11) indicates that using gradient descent to update the parameters in  $G$  is a principled method to train the network of  $G$ . Finally, the new objective for  $G$  can be obtained:

$$L_G = \min_{\theta} -E_{z \sim p(z)} [f(G(z))] \tag{12}$$

#### 4.2 Implementation of BC conditions

In this paper, we introduce a simple but efficient implementation of BC conditions. When applying the BC conditions to  $D$ , the training of  $D$  can be equivalently regarded

as a conditional (constrained) optimization process. Then, Eq. (10) can be updated as:

$$\min_{f \in \Omega} \{E_{z \sim p(z)}[f(G(z))] - E_{x \sim P_r}[f(x)]\} \tag{13}$$

$$s.t. \|f(x)\| \leq m, \quad \text{if } \frac{\partial f(x)}{\partial x} \text{ exists, } \quad \left\| \frac{\partial f(x)}{\partial x} \right\| < \infty$$

---



---

**Algorithm 1: BC-GAN**

---

Require:  
 the number of  $D$  iteration per  $G$  iteration  $n_{critic}$ ,  
 the batch size  $n$ , the bound  $m$ ,  
 initial critic parameter  $w_0$ ,  
 initial generator parameters  $\theta_0$   
 1: **while**  $\theta$  has not converged **do**  
 2: Sample  $\{x^{(i)}\}_{i=1}^n \sim P_r$   
 3: Sample  $\{z^{(i)}\}_{i=1}^n \sim P_z$   
 4: **for**  $t=1, 2, \dots, n_{critic}$  **do**  
 5:  $\frac{1}{n} \sum_{i=1}^n f(x^{(i)}) \rightarrow L_r$   
 6:  $\frac{1}{n} \sum_{i=1}^n f(g(z^{(i)})) \rightarrow L_g$   
 7:  $[L_g - L_r + \beta \cdot \max(\|f(x)\| - m, 0)] \rightarrow L_D$   
 8: Adam( $\nabla_w L_D$ )  $\rightarrow w$   
 9: **end for**  
 10: Adam( $\nabla_{\theta} [-\frac{1}{n} \sum_{i=1}^n f(g(z^{(i)})])$ )  $\rightarrow \theta$   
 11: **end while**

---

In this paper, the discriminator function  $f$  is implemented by a deep neural network, which applies a series of convolutional and nonlinear operations. Both convolutional and nonlinear functions are continuous, which means that  $D$  is a continuous function. Moreover, the gradients of the output of  $D$  with respect to the input are always finite. As a result, condition (b) is satisfied naturally. To guarantee condition (a), the Lagrange multiplier method can be applied here; then, the objective of  $D$  can be written as the following equation:

$$L_D = \min_f \{E_{z \sim p(z)}[f(G(z))] - E_{x \sim P_r}[f(x)]\} + \beta \cdot \max(\|f(x)\| - m, 0) \tag{14}$$

where  $\beta$  is the hyperparameter and  $m$  represents the bound. The term  $\max(\|f(x)\| - m, 0)$  plays the role of forcing  $D$  to be a bounded function, while  $E_{z \sim p(z)}[f(G(z))] - E_{x \sim p(x)}[f(x)]$  is used to determine  $k \cdot W(P_r, P_g)$ . The procedure of training the BC-GAN is described in Algorithm 1.

### 4.3 Validity

In order to verify the validity of proposed BC conditions, we use synthetic datasets as those presented in [15] to test discriminator’s performance. Specifically, discriminators are trained to distinguish the fake distribution from the real one. The toy distributions hold the fake distribution  $P_g$  as the real distribution  $P_r$  plus unit-variance Gaussian noise. Theoretically, discriminator with good performance is

more likely to learn the high moments of the data distributions and model the real distribution. Figure 1 illustrates the value surfaces of the discriminator. It is clearly seen that discriminator enforced by BC conditions has a good performance on discriminating the real samples from the fake ones, demonstrating the validity of proposed method.

### 4.4 Comparison with spectral normalization and gradient penalty

Gradient penalty, spectral normalization and our proposed method are inspired by different motivations to enforce the Lipschitz constraint on  $D$ . Therefore, they differ in the way of implementation and in principle. The first difference is the way of implementation. Gradient penalty and our method operate on the loss function directly, while spectral normalization constrains the weight matrix instead of the loss metric.

Secondly, they differ in principle. For BC-GAN,  $k \cdot W(P_r, P_g)$  is applied to evaluate the difference between the fake and real distributions instead of  $W(P_r, P_g)$ , which is used in WGAN-GP and WGAN. Moreover, WGAN-GP and SN-GAN strictly constrain the Lipschitz constant to be 1 or a known constant, while BC-GAN eases the restriction on the Lipschitz constant, and  $k$  is an unknown scalar parameter which will have no influence on the training of the network. Therefore,  $k \cdot W(P_r, P_g)$  can be employed as a new loss metric to guide the training of  $D$ .

To visualize the differences, we still use the synthetic datasets to test discriminators’ performance. Figure 1 illustrates the value surfaces of the discriminators. It is obvious that discriminators trained with gradient penalty as well as spectral normalization have pathological value surfaces even when optimization has completed, and they have failed to capture the high moments of the data distributions and instead model very simple approximations to the optimal functions. In contrast, BC-GANs have successfully learned the higher moments of the data distributions, and the discriminator can distinguish the real distribution from the fake one much better.

### 4.5 Convergence measure

One advantage of using Wasserstein distance as the metric over KL divergence is the meaningful loss. The Wasserstein distance  $W(P_r, P_g)$  shows the property of convergence [6]. If it stops decreasing, then the training of the network can be terminated. This property is useful as one does not have to stare at the generated samples to figure out the failure modes. To obtain the convergence measure in the proposed BC-GAN, a corresponding indicator of the training stage is introduced:

$$I_{GD} = \frac{1}{\|\nabla_x f(x)\|_2} \quad (15)$$

To prove that proposed indicator  $I_{GD}$  is capable of convergence measure, Theorem 3 is introduced.

**Theorem 3** Let  $P_r$  and  $P_g$  be the distributions of real and generated images,  $x$  is the image located in  $P_r$  and  $P_g$ , and  $f$  is the discriminator function, bounded by the BC conditions.  $I_{GD}$  in Eq. 15 is proportional to  $W(P_r, P_g)$ .

## 5 Experiments

### 5.1 Experimental setup

In order to assess the performance of BC-GAN, image generation experiments are conducted on CIFAR-10 [20], STL-10 [8] and CELEBA [25] datasets. Two widely used GAN architectures, including the standard CNN and ResNet-based CNN [6], are applied for image generation task. For the architecture details, please see “Appendix”. Equations (12) and (14) are used as the loss metric of  $D$  and  $G$ , respectively.  $I_{GD}$  in Eq. (15) acts as the role of measuring convergence.  $m$  and  $\beta$  in Eq. (14) are set as 0.5 and 2, respectively. For optimization, the Adam [9] is utilized in all the experiments with  $\alpha = 0.0002$ ,  $\beta_1 = 0$ ,  $\beta_2 = 0.9$ .  $D$  updates 5 times per  $G$  update. To keep it identical to previous GANs, we set the batch size as 64. Inception score [17] and Fréchet inception distance [8] are utilized for quantitative assessment of generated examples.

Although inception score and Fréchet inception distance are widely used as an evaluation metric for GANs, Barratt [3] suggests that it should be more systematic and careful when evaluating and comparing generative models, because inception score may not correlate well with the image quality strictly. Recently, Catherine [14] proposes a new method to evaluate the generative models, called skill rating. Skill rating evaluates models by carrying out tournaments between the discriminators and generators. For better evaluation, results assessed by skill rating are also presented.

### 5.2 Results on image generation

Image generation tasks are carried out on the CIFAR-10 and STL-10 datasets. Based on the ResNet-based CNN architecture, we obtain the average inception score of 8.40 and 9.15 for image generation on CIFAR-10 and STL-10, respectively. We compare our algorithm against multiple benchmark methods. In Table 1, we show the inception score and Fréchet inception distance of different methods with their corresponding optimal settings on CIFAR-10 and

**Table 1** IS and FID of unsupervised image generation on CIFAR-10 and STL-10. IS is the inception score, and FID represents Fréchet inception distance. For IS, higher is better, while lower is better for FID

Method	CIFAR-10		STL-10	
	IS	FID	IS	FID
Real data	11.24 ± .12	7.8	26.08±.26	7.9
<i>Standard CNN</i>				
DCGAN	6.64 ± .14		7.84±.07	
WGAN-GP	6.53 ± .08	40.2	8.42±.13	55.1
SN-GAN	7.42 ± .06	29.3	8.28±.09	53.1
BC-GAN	<b>7.48 ± .06</b>	<b>28.9</b>	8.30 ±.12	54.5
<i>ResNet</i>				
WGAN-GP	7.86 ± .13			
SN-GAN	8.22 ± .05	21.7	9.10±.04	40.1
BC-GAN	<b>8.40 ± .10</b>	<b>20.8</b>	<b>9.15±.17</b>	<b>39.9</b>
LR-GAN [24]	7.17 ± .07			
DFM [22]	7.72 ± .13		8.51±.12	
Orthonormal [13]	7.40 ± .04	29	8.56±.09	46.7

Bold means our proposed model

STL-10 datasets. As illustrated in Table 1, BC-GAN has comparable performances with the state-of-the-art GANs. We also conduct image generation on CELEBA [25] dataset. Examples of generated images are shown in Fig. 2 and 3.

Skill rating [14] is recently introduced to judge the GAN model by matches between  $G$  and  $D$ . To determine the outcome of a match between  $G$  and  $D$ ,  $D$  judges two batches: one batch of samples from  $G$  and one batch of real data. Every sample  $x$  that is not judged correctly by  $D$  (e.g.,  $D(x) > 0.5$  for the generated data or  $D(x) < 0.5$  for the real data) counts as a win for  $G$  and is used to compute its win rate. Win rate tests the performance between  $D$  and  $G$  dynamically in the training process and judges whether  $D$  or  $G$  dominates, while the other stops updating. If  $D$  dominates and  $G$  stops updating, win rate for  $G$  decreases dramatically. We make some modifications, because we use Wasserstein distance to determine the difference between fake and real data instead of probability. As a result, we show the loss of  $D$  instead of the win rate in Fig. 4. When  $D$  in the latter iteration is used to distinguish the generated images in the early iteration from real images, it outputs a large loss, meaning that  $D$  can easily distinguish the generated images (fake images) from real images. And the images generated in the latter iteration can also easily fool  $D$  in the early iteration. Therefore, there is a healthy training, and the performance of  $D$  and  $G$  is continuously improved in the training process.

When applying KL divergence as the loss metric of  $D$ , the training of GANs suffers from the vanishing gradient

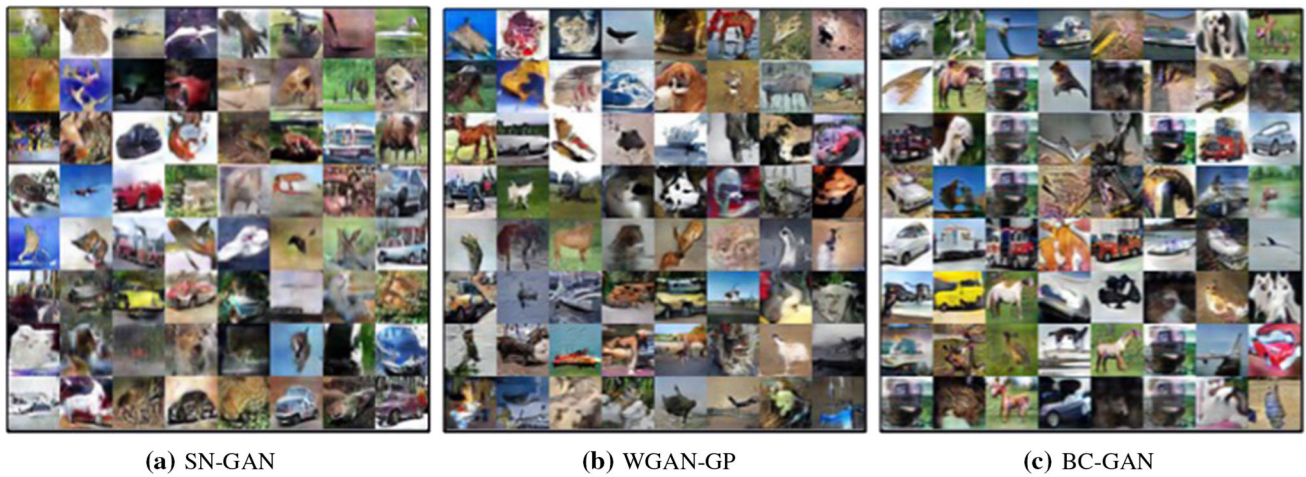


Fig. 2 Image generation on CIFAR-10 dataset using a SN-GAN, b WGAN-GP and c BC-GAN

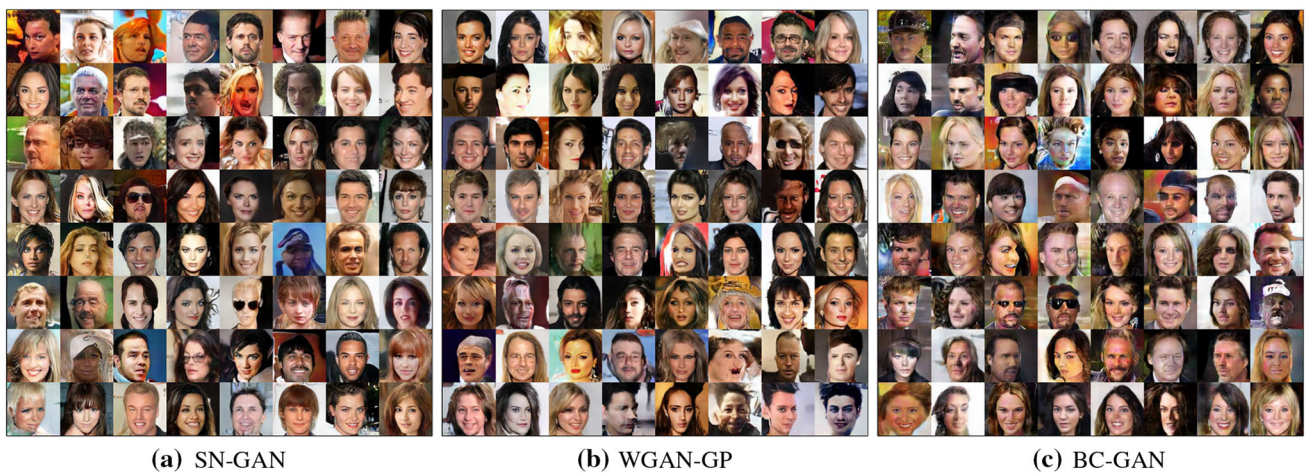


Fig. 3 Image generation on CELEBA dataset using a SN-GAN, b WGAN-GP and c BC-GAN

problem, *i.e.*, zero gradient would backpropagate to  $G$ , and the training would completely stop. As a comparison, Fig. 4 shows a healthy training during the entire iterations, further indicating the effectiveness of BC-GANs.

## 6 Analysis

### 6.1 Bound $m$

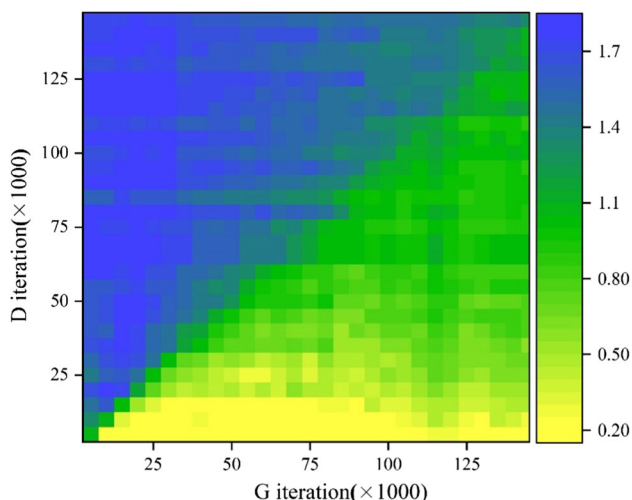
The parameter  $m$  in Eq. (14) represents the bound of  $D$ , and it actually controls the gradient  $\partial L_D / \partial x$ , where  $L_D$  is the loss of  $D$ ,  $x$  is the image and  $\partial L_D / \partial x$  is the gradient backpropagated from  $D$  to  $G$ , which indeed affects the training of  $G$  and further influences the model performance. Explanation is as followed. The discriminator  $f$  is a bounded function. Given enough iterations,  $f_{x \sim P_r}(x)$  would always converge to  $m$  and  $f_{x \sim P_g}(x)$  would converge to  $-m$ .

And considering that  $f$  satisfies  $k$ -Lipschitz constraint, the following condition is satisfied:

$$\|f_{x_r \sim P_r}(x_r) - f_{x_g \sim P_g}(x_g)\| \approx 2m \leq k \|x_r - x_g\| \tag{16}$$

$$\frac{2m}{\|x_r - x_g\|} \leq k \tag{17}$$

$k$  determines the upper bound of the gradient backpropagated from  $D$  to  $G$  and is directly proportional to  $D$ . Increasing  $m$  enhances the upper bound of the gradients  $\partial L_D / \partial x$ . This is verified by the experiment shown in Fig. 5a. Moreover, the gradients are used to guide the training of the generator and naturally affect the performance of the model. Increasing  $m$  from 0.5 to 2 leads to decreased performance (inception score drops from 8.40 to 7.56). Therefore, properly controlling the gradient is important for improving the performance of GAN models. And the bound  $m$  provides such a mechanism for controlling the gradient.  $m$  is recommended to be taken as 0.5 for



**Fig. 4** Matches between  $D$  and  $G$ . Wasserstein distance is utilized to indicate the results instead of the win rate. With larger value of the Wasserstein distance,  $D$  is more likely to distinguish the real images from the fake ones. Lower value of the Wasserstein distance indicates that  $G$  is more likely to fool  $D$

image generation task on CIFAR-10. One possible explanation why a smaller  $m$  (hence smaller gradients back-propagated) in the training leads to better performances is that the error surfaces are highly nonlinear, the backpropagation is a gradient descent and greedy algorithm, small gradients may help the optimization lead to a deeper local minimum or indeed the global minimum of the error surface.

We also monitor the variation of the gradient on WGAN-GP and SN-GAN. It is found that the behavior of the gradient variation varies on different models. The gradient penalty term in WGAN-GP forces the gradient of the output of  $D$  with respect to the input to be a fixed number. Therefore, as shown in Fig. 5b, the gradient is around 1 in the whole training process. For SN-GAN and


our BC-GAN in Fig. 5c, the variation of the gradient is similar. With training process going on, the gradient tends to increase until convergence is reached. The difference is that the amplitude of the gradient in SN-GAN is larger than that in BC-GAN. As mentioned above, the amplitude of the gradient indeed affects the training of the generator. However, SN-GAN provides no mechanism for controlling the gradient, while the bound  $m$  in BC-GAN acts as the role of controlling the gradient. Thus, at least in this perspective, BC-GAN has a better performance control over SN-GAN.

### 6.2 Meaningful training stage indicator $I_{GD}$

We introduce a new indicator  $I_{GD}$  for monitoring the training stage. Figure 6a shows the correlation of  $-I_{GD}$  with inception score during the training process. Because  $I_{GD}$  decreases with the iteration, we use  $-I_{GD}$  instead. As we can see,  $-I_{GD}$  has a positive correlation with the inception score. As it is easier to visualize the correlation between  $I_{GD}$  and image quality in higher-resolution images, we perform image generation task on CELEBA [25] dataset and show the variation of  $I_{GD}$  with iterations in Fig. 6b. It is clearly seen that  $I_{GD}$  correlates well with image quality during the training process.


### 6.3 Training time

It is worth noting that BC-GAN is computationally efficient. We list the computational time for 100 generator updates in Fig. 7. WGAN-GP requires more computational time because it needs to calculate the gradient of the gradient norm  $\|\nabla_x D\|_2$ , which needs one whole round of forward and backward propagation. And spectral normalization needs to calculate the largest singular value of the matrices in each layer. What is worse, for gradient penalty

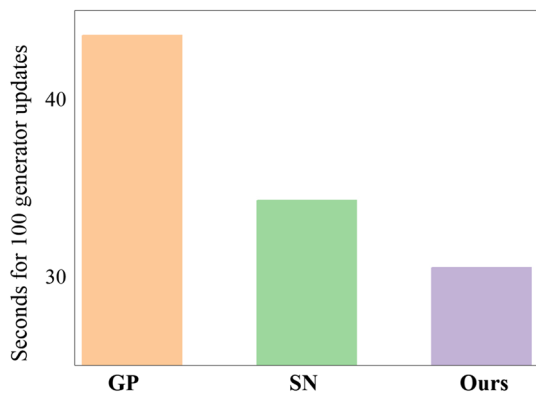


**Fig. 5** **a** Variation of the gradient  $\partial L_D / \partial x$  with iterations in BC-GAN. Larger  $m$  leads to higher gradients. **b** Variation of the gradient with iterations in WGAN-GP. **c** Comparison of the gradient variation of SN-GAN and BC-GAN, where SN represents SN-GAN, and BC is BC-GAN





**Fig. 6** **a** Correlation of  $-I_{GD}$  with inception score on CIFAR10. **b** Variation of  $I_{GD}$  with iteration for the training on CELEBA database.  $I_{GD}$  correlates well with the image quality, indicating that  $I_{GD}$  can be regarded as the indicator of the training stage



**Fig. 7** Computation time for 100 generator updates. GP for WGAN-GP and SN for SN-GAN. We use standard CNN as the architecture. Tests are based on Nvidia 1080Ti

and spectral normalization, the extra computational costs increase with the increase in layers. As for BC-GAN, there is no matrix operation or gradient calculation in the backpropagation. As a result, it has lower computational cost.

## 7 Concluding remarks

In this paper, we have introduced a new generative adversarial network training technique called BC-GAN which utilizes bounded discriminator to enforce Lipschitz constraint. In addition to provide theoretical background, we have also presented practical implementation procedures for training BC-GAN. Experiments on synthetical as well as real data show that the new BC-GAN performs better and has lower computational complexity than recent techniques such as spectral normalization GAN (SN-GAN) and Wasserstein GAN with gradient penalty (WGAN-GP). We have also introduced a new training convergence measure which correlates directly with the image quality of the generator output and can be conveniently used to

monitor training progress and to decide when training is completed.

## Compliance with ethical standards

**Conflict of interest** The authors declare that they have no conflict of interest. We declare that we do not have any commercial or associative interest that represents a conflict of interest in connection with the work submitted.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## Appendix 1

### Proofs

Let  $P_r$  and  $P_g$  be the distributions of real images and generated images in  $X$ , a compact metric space. The discriminator function  $f$  is constructed based on a neural network of the following form with input  $x$ :

$$f(x, \theta) = W^{L+1} a_L(W^L(a_{L-1}(\dots a_1(W^1 x)))) \quad (18)$$

where  $\theta := \{W^1, W^2, \dots, W^{L+1}\}$  is the learning parameter set, and  $a_i$  is an element-wise nonlinear activation function.

Spectral normalization is applied on  $f$  to guarantee the Lipschitz constraint.

**Proposition 1** *When using Wasserstein distance as the loss metric of  $f$ , the optimal solution to  $f$  is unreachable.*

**Proof** Corollary 1 in [6] has proven that the optimal discriminator  $f^*$  has gradient norm 1 almost everywhere under  $P_r$  and  $P_g$  when using Wasserstein distance as the loss metric.

Suppose  $x$  can be expressed as  $[x_1, x_2, \dots, x_n]$ , and  $W^T W$  has eigenvalues  $\lambda = [\lambda_1, \lambda_2, \dots, \lambda_n]$ :

$$\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n \geq 0 \tag{19}$$

The eigenvectors of  $W^T W$  can be expressed as  $V = [v_1, v_2, \dots, v_n]$ . Then, we have:

$$\|Wx\|^2 = x^T W^T W x = x^T V^T \lambda V x \tag{20}$$

Supposing the transformation  $Vx = [y_1, y_2, \dots, y_n]$ , and using the relationship  $V^T V = I$ , we can have

$$\begin{aligned} x^T V^T \lambda V x &= \lambda_1 y_1^2 + \lambda_2 y_2^2 + \dots + \lambda_n y_n^2 \\ &\leq \lambda_1 (y_1^2 + y_2^2 + \dots + y_n^2) \\ &= \lambda_1 (x^T V^T V x) = \lambda_1 \|x\|^2 \end{aligned} \tag{21}$$

When spectral normalization is applied,  $\lambda_1$  is normalized to 1. As a result:

$$\|Wx\|^2 \leq \|x\|^2 \tag{22}$$

We can see that applying spectral normalization can guarantee  $W$  satisfies the Lipschitz constraint. The discriminator function  $f$  is implemented by convolutional neural networks, which is a combination of convolutional and nonlinear operations (Eq. (18)). Therefore, the following inequality is applied to observe the bound on  $\|f\|_{Lip}$  [13]:

$$\|f\|_{Lip} \leq \prod_{l=1}^{L+1} \sigma(W^l) \cdot \prod_{l=1}^{L+1} \|a_l\|_{Lip} = 1 \tag{23}$$

where  $\sigma(W)$  is the spectral norm of  $W$ .

When applying Wasserstein distance as the loss metric, Corollary 1 in [6] has proven that the optimal solution to the Lipschitz constrained discriminator has gradient norm 1 almost everywhere under  $P_r$  and  $P_g$ , which means  $\|f\|_{Lip}$  needs to reach upper bound of 1. However, if  $\|f\|_{Lip}$  in Eq. (23) needs to obtain the upper bound 1, the discriminator function becomes a linear function. However, the discriminator function is implemented by the combination of convolutional operation and nonlinear operation. Taking the Relu function as a representation of the nonlinear operation  $a_l$ ,  $a_l(x) = x(x > 0)$  or  $a_l(x) = 0(x \leq 0)$ . In another word,  $\|a_l(x)\| = \|x\|(x > 0)$ , and

$\|a_l(x)\| = 0 < \|x\|(x \leq 0)$ . If the discriminator function needs to obtain the upper bound of the Lipschitz constraint, all the nonlinear operations need to reach the upper bound as well:  $\|a_l(x)\| = \|x\|$ . Then, all the nonlinear functions are linear functions, and the discriminator function turns to a linear function. Obviously, a linear discriminator is not the optimal solution. Therefore, with the existence of nonlinear operation, applying spectral normalization prevents the discriminator functions from the optimal solution when applying Wasserstein distance as the loss metric.

**Theorem 1** *Let  $\Psi$  be the set of all  $f : X \rightarrow R$ , where  $f$  is a continuous function. In addition, the number of points where  $f$  is continuous but not differentiable is finite. Besides, if  $f$  is differentiable at point  $x$ , its derivative is finite. Then,  $f$  in  $\Psi$  satisfies Lipschitz constraint.*

**Proof** (i) Considering that  $f$  is derivable. According to Lagrange’s mean value theorem,

$$\frac{f(x_1) - f(x_2)}{x_1 - x_2} = \frac{\partial f(x_0)}{\partial x_0} (x_0 \in [x_1, x_2]) \tag{24}$$

Because  $\frac{\partial f}{\partial x}$  is finite:

$$\frac{f(x_1) - f(x_2)}{x_1 - x_2} = \frac{\partial f(x_0)}{\partial x_0} \leq k \tag{25}$$

where  $k$  is finite.

Moreover, we have:

$$\|f(x_1) - f(x_2)\| \leq k \|x_1 - x_2\| \tag{26}$$

Then,  $f$  satisfies Lipschitz constraint.

(ii) Considering that  $f$  is not derivable,  $f$  is a continuous function; then, there must be at least one point  $x_0$ , at which  $f$  is continuous but not derivable. We only consider that there is only one such point. For multiple points, the conclusion is the same. For any  $x_1$  and  $x_2$  ( $x_1, x_2 < x_0$  or  $x_1, x_2 > x_0$ ),  $f$  should satisfy the following:

$$\|f(x_1) - f(x_2)\| \leq k \|x_1 - x_2\| \tag{27}$$

because  $f$  is continuous and derivable in  $[x_1, x_2]$ .

For  $x_1$  and  $x_2$  ( $x_1 < x_0 < x_2$ ), we have

$$\begin{aligned} \|f(x_1) - f(x_2)\| &= \|f(x_1) - f(x_0) + f(x_0) - f(x_2)\| \\ &\leq \|f(x_1) - f(x_0)\| + \|f(x_0) - f(x_2)\| \end{aligned} \tag{28}$$

Because  $f$  is continuous in  $[x_1, x_0]$  and  $[x_0, x_2]$ , and derivable in  $(x_1, x_0)$  and  $(x_0, x_2)$ , we can obtain:

$$\|f(x_1) - f(x_0)\| \leq k_1 \|x_1 - x_0\| \tag{29}$$

$$\|f(x_0) - f(x_2)\| \leq k_2 \|x_0 - x_2\| \tag{30}$$

Then, we can have:

$$\|f(x_1) - f(x_2)\| \leq k(\|x_1 - x_0\| + \|x_0 - x_2\|) \tag{31}$$

where  $k = \max(k_1, k_2)$ . Considering the relationship that  $x_1 < x_0 < x_2$ , we can have:

$$\|f(x_1) - f(x_2)\| \leq k(\|x_1 - x_2\|) \tag{32}$$

As we can see, even though  $f$  is not derivable at  $x_0$ , for any  $x_1$  and  $x_2$ ,  $f$  still satisfies:  $\|f(x_1) - f(x_2)\| \leq k(\|x_1 - x_2\|)$ .

To sum up,  $f$  always satisfies Lipschitz constraint at the given conditions.

**Theorem 2** *Let  $P_r$  and  $P_g$  be the distributions of real images and generated images in  $X$ , a compact metric space. Let  $\Omega$  be the set of all  $f : X \rightarrow R$ , where  $f$  is a continuous and bounded function. And, the number of points where  $f$  is continuous but not differentiable is finite. Besides, if  $f$  is differentiable at point  $x$ , its derivative is finite. The set  $\Omega$  can be expressed as:*

$$\Omega : \left\{ f \mid \|f(x)\| \leq m, \text{ if } \frac{\partial f(x)}{\partial x} \text{ exists, } \left\| \frac{\partial f(x)}{\partial x} \right\| < \infty \right\} \tag{33}$$

where  $m$  represents the bound. Then, there must exist a  $k$ , and we have a computable  $k \cdot W(P_r, P_g)$ :

$$k \cdot W(P_r, P_g) = \sup_{f \in \Omega} E_{x \sim P_r} [f(x)] - E_{x \sim P_g} [f(x)] \tag{34}$$

where  $W(P_r, P_g)$  represents the Wasserstein distance [5, 23] between  $P_r$  and  $P_g$ .

**Proof** According to Theorem 1, for  $f$  in  $\Omega$ , there exists a  $k$  to satisfy Eq. (32). Then,  $\Omega$  is the set, which contains all the  $k$ -Lipschitz constrained functions  $f$ . Kantorovich–Rubinstein duality [5, 23] tell us that the supremum over all the functions in  $\Omega$  is  $k \cdot W(P_r, P_g)$ . As a result, we can obtain Eq. (34). To guarantee the boundedness and computability of  $k \cdot W(P_r, P_g)$ ,  $f$  is supposed to be a bounded function. Because, even though  $k$  in Theorem 1 is a finite number, it can be super large  $k \rightarrow \infty$ , leading to the incomputability of  $k \cdot W(P_r, P_g)$ . Enforcing  $f$  to be a bounded function can ensure the boundedness and computability of  $k \cdot W(P_r, P_g)$ :

$$k \cdot W(P_r, P_g) = \sup_{f \in \Omega} E_{x \sim P_r} [f(x)] - E_{x \sim P_g} [f(x)] \leq 2m \tag{35}$$

**Theorem 3** *Let  $P_r$  and  $P_g$  be the distributions of real and generated images,  $x$  is the image located in  $P_r$  and  $P_g$ , and  $f$  is the discriminator function, bounded by the BC conditions.  $I_{GD}$  in Eq. 15 is proportional to  $W(P_r, P_g)$ .*

**Proof**  $f$  is bounded by the BC conditions. Given enough iterations,  $f_{x \sim P_r}(x)$  would always converge to  $m$  and

$f_{x \sim P_g}(x)$  would converge to  $-m$ . As a result,  $k \cdot W(P_r, P_g)$  will always converge to  $2m$ :

$$k \cdot W(P_r, P_g) = \sup_{f \in \Omega} E_{x \sim P_r} [f(x)] - E_{x \sim P_g} [f(x)] \approx 2m \tag{36}$$

It is clear that  $W(P_r, P_g)$  is proportional to  $E[\|x_r - x_g\|]$ , because both of them evaluate the difference between  $P_r$  and  $P_g$ . Then, we can use the following term  $GD$  to estimate  $W(P_r, P_g)$ :

$$GD = \frac{\|f_{x_r \sim P_r}(x_r) - f_{x_g \sim P_g}(x_g)\|}{\|x_r - x_g\|} \tag{37}$$

where  $x_r, x_g$  are the real image and generated image, respectively. As expressed above, the term  $\|f_{x_r \sim P_r}(x_r) - f_{x_g \sim P_g}(x_g)\|$  would always converge to  $2m$ , and  $W(P_r, P_g)$  is proportional to  $E[\|x_r - x_g\|]$ . Therefore,  $GD$  is inversely related to  $W(P_r, P_g)$ , and the reciprocal of  $GD$  can be used to roughly estimate  $W(P_r, P_g)$ .  $\square$

According to Lagrange’s mean value theorem,

$$GD = \frac{\|f_{x_r \sim P_r}(x_r) - f_{x_g \sim P_g}(x_g)\|}{\|x_r - x_g\|} = \|\nabla_x f(x)\|_2 \tag{38}$$

where  $x \in [x_g, x_r]$ . For the convenience of calculation,  $x$  is taken as  $x = \alpha \cdot x_r + (1 - \alpha) \cdot x_g$ , and  $\alpha \in [0, 1]$ . Then,  $\|\nabla_x f(x)\|_2$  is inversely related to  $W(P_r, P_g)$ . Finally,  $I_{GD}$  is proportional to  $W(P_r, P_g)$ .

## Appendix 2: Architecture

Discriminator in the toy model is listed in Table 2. Standard CNN architectures for CIFAR-10 and STL-10 are listed in Tables 3 and 4. ResNet-based CNN architectures for CIFAR10 and STL-10 are listed in Tables 5 and 6. Architectures for image generation on CELEBA dataset are listed in Tables 7 and 8.

**Table 2** Discriminator in the toy model

Input points : $x \in R^2$
Dense, Relu $\rightarrow 512 \times 2$
Dense, Relu $\rightarrow 512 \times 2$
Dense, Relu $\rightarrow 512 \times 2$
Dense $\rightarrow 1$

**Table 3** Generator of standard CNN architectures for CIFAR-10 and STL-10

---

Latent vector :  $z \in R^{128} \sim N(0, 1)$   
 Dense, BN, Relu  $\rightarrow 4 \times 4 \times 512$   
 $5 \times 5$ , stride = 2, Deconv, BN, Relu  $\rightarrow 8 \times 8 \times 256$   
 $5 \times 5$ , stride = 2, Deconv, BN, Relu  $\rightarrow 16 \times 16 \times 128$   
 $5 \times 5$ , stride = 2, Deconv, BN, Relu  $\rightarrow 32 \times 32 \times 64$   
 $3 \times 3$ , stride = 1, Conv, Tanh  $\rightarrow 32 \times 32 \times 3$

---

**Table 4** Discriminator of standard CNN architectures for CIFAR-10 and STL-10

---

Input RGB image :  $x \in R^{32 \times 32 \times 3}$   
 $3 \times 3$ , stride = 1, Conv, Leaky-Relu  $\rightarrow 32 \times 32 \times 64$   
 $5 \times 5$ , stride = 2, Conv, Leaky-Relu  $\rightarrow 16 \times 16 \times 128$   
 $5 \times 5$ , stride = 2, Conv, Leaky-Relu  $\rightarrow 8 \times 8 \times 256$   
 $5 \times 5$ , stride = 2, Conv, Tanh  $\rightarrow 4 \times 4 \times 512$   
 Dense  $\rightarrow 1$

---

**Table 5** Generator of ResNet-based CNN architectures for CIFAR10 and STL-10

---

Latent vector :  $z \in R^{128} \sim N(0, 1)$   
 Dense  $\rightarrow 4 \times 4 \times 128$   
 ResBlock up  $\rightarrow 8 \times 8 \times 128$   
 ResBlock up  $\rightarrow 16 \times 16 \times 128$   
 ResBlock up  $\rightarrow 32 \times 32 \times 128$   
 BN, Relu  $\rightarrow 32 \times 32 \times 128$   
 $3 \times 3$ , stride = 1, Conv, Tanh  $\rightarrow 32 \times 32 \times 3$

---

**Table 6** Discriminator of ResNet-based CNN architectures for CIFAR10 and STL-10

---

Input RGB image :  $x \in R^{32 \times 32 \times 3}$   
 $3 \times 3$ , stride = 1, Conv  $\rightarrow 32 \times 32 \times 64$   
 ResBlock down  $\rightarrow 16 \times 16 \times 128$   
 ResBlock down  $\rightarrow 8 \times 8 \times 128$   
 ResBlock down  $\rightarrow 4 \times 4 \times 128$   
 Dense  $\rightarrow 1$

---

**Table 7** Generator architecture for image generation on CELEBA dataset

---

Latent vector :  $z \in R^{128} \sim N(0, 1)$   
 Dense, BN, Relu  $\rightarrow 4 \times 4 \times 512$   
 Upsample  $\rightarrow 8 \times 8 \times 512$   
 $3 \times 3$ , stride = 1, Conv, BN, Relu  $\rightarrow 8 \times 8 \times 256$   
 Upsample  $\rightarrow 16 \times 16 \times 256$   
 $3 \times 3$ , stride = 1, Conv, BN, Relu  $\rightarrow 16 \times 16 \times 128$   
 Upsample  $\rightarrow 32 \times 32 \times 128$   
 $3 \times 3$ , stride = 1, Conv, BN, Relu  $\rightarrow 32 \times 32 \times 64$   
 Upsample  $\rightarrow 64 \times 64 \times 64$   
 $3 \times 3$ , stride = 1, Conv, BN, Relu  $\rightarrow 64 \times 64 \times 32$   
 Upsample  $\rightarrow 128 \times 128 \times 32$   
 $3 \times 3$ , stride = 1, Conv, BN, Relu  $\rightarrow 128 \times 128 \times 32$   
 $3 \times 3$ , stride = 1, Conv, Tanh  $\rightarrow 128 \times 128 \times 3$

---

**Table 8** Discriminator architecture for image generation on CELEBA dataset

---

Input RGB image :  $x \in R^{128 \times 128 \times 3}$   
 $3 \times 3$ , stride = 1, Conv, Leaky-Relu  $\rightarrow 128 \times 128 \times 64$   
 Downsample  $\rightarrow 64 \times 64 \times 64$   
 $3 \times 3$ , stride = 1, Conv, Leaky-Relu  $\rightarrow 64 \times 64 \times 128$   
 Downsample  $\rightarrow 32 \times 32 \times 128$   
 $3 \times 3$ , stride = 1, Conv, Leaky-Relu  $\rightarrow 32 \times 32 \times 256$   
 Downsample  $\rightarrow 16 \times 16 \times 256$   
 $3 \times 3$ , stride = 1, Conv, Leaky-Relu  $\rightarrow 16 \times 16 \times 512$   
 Downsample  $\rightarrow 8 \times 8 \times 512$   
 $3 \times 3$ , stride = 1, Conv, Leaky-Relu  $\rightarrow 8 \times 8 \times 512$   
 Downsample  $\rightarrow 4 \times 4 \times 512$   
 Dense  $\rightarrow 1$

---

## References

1. Anil C, Lucas J, Grosse R (2018) Sorting out lipschitz function approximation. arXiv preprint [arXiv:1811.05381](https://arxiv.org/abs/1811.05381)
2. Arjovsky M, Chintala S, Bottou L (2017) Wasserstein gan. arXiv preprint [arXiv:1701.07875](https://arxiv.org/abs/1701.07875)
3. Barratt S, Sharma R (2018) A note on the inception score. arXiv preprint [arXiv:1801.01973](https://arxiv.org/abs/1801.01973)
4. Berthelot D, Schumm T, Metz L (2017) Began: boundary equilibrium generative adversarial networks. arXiv preprint [arXiv:1703.10717](https://arxiv.org/abs/1703.10717)

5. Goodfellow I, Pouget-Abadie J, Mirza M (2014) Generative adversarial nets. *Advances in neural information processing systems*
6. Gulrajani I, Ahmed F, Arjovsky M (2017) Improved training of wasserstein gans. In: *Advances in neural information processing systems*, pp 5769–5779
7. Heinonen J (2005) *Lectures on lipschitz analysis*. University of Jyväskylä
8. Heusel M, Ramsauer H, Unterthiner T (2017) Gans trained by a two time-scale update rule converge to a nash equilibrium. In: *Advances in neural information processing systems*, pp 6626–6637
9. Kingma D, Ba J (2015) Adam: a method for stochastic optimization. In: *International conference on learning representations (ICLR)*
10. Li Q, Haque S, Anil C, Lucas J, Grosse RB, Jacobsen J (2019) Preventing gradient attenuation in lipschitz constrained convolutional networks. In: *Advances in neural information processing systems*, pp 15364–15376
11. Mao X, Li Q, Xie H (2017) Least squares generative adversarial networks. In: *2017 IEEE international conference on computer vision (ICCV)*, pp 2813–2821
12. Mescheders L, Geiger A, Nowozin S (2018) Which training methods for gans do actually converge? *arXiv preprint arXiv:1801.04406*
13. Miyato T, Kataoka T, Koyama M (2018) Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*
14. Olsson C, Bhupatiraju S, Brown T (2018) Skill rating for generative models. *arXiv preprint arXiv:1808.04888*
15. Qi GJ (2017) Loss-sensitive generative adversarial networks on lipschitz densities. *arXiv preprint arXiv:1701.06264*
16. Radford A, Metz L, Chintala S (2015) Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*
17. Salimans T, Goodfellow I, Zaremba W (2016) Improved techniques for training gans. In: *Advances in neural information processing systems*, pp 2234–2242
18. Salimans T, Kingma DP (2016) Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Advances in Neural Information Processing Systems* pp. 901–909
19. Sokolić J, Giryès R, Sapiro G, Rodrigues M (2017) Robust large margin deep neural networks. *IEEE Trans Signal Process* 65(16):4265–4280
20. Torralba A, Fergus R, Freeman WT (2005) 80 million tiny images: a large data set for non-parametric object and scene recognition. *IEEE Trans Pattern Anal Mach Intel* 30(11):901–909
21. Tsuzuku Y, Sato I, Sugiyama M (2018) Lipschitz-margin training: scalable certification of perturbation invariance for deep neural networks. In: *Advances in neural information processing systems*, pp 6541–6550
22. Warde-Farley D, Bengio Y (2016) Improving generative adversarial networks with denoising feature matching
23. Wu J, Huang Z, Thoma J (2017) Energy-relaxed Wasserstein gans (energywgan): towards more stable and high resolution image generation. *arXiv preprint arXiv:1712.01026*
24. Yang J, Kannan A, Batra D (2017) Lr-gan: layered recursive generative adversarial networks for image generation. *arXiv preprint arXiv:1703.01560*
25. Yang S, Luo P, Loy CC (2015) From facial parts responses to face detection: a deep learning approach. In: *IEEE international conference on computer vision (ICCV)*, pp 3676–3684

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.