

On Positive-Unlabeled Classification in GAN

Tianyu Guo^{1,2}, Chang Xu², Jiajun Huang², Yunhe Wang³, Boxin Shi^{4,5}, Chao Xu¹, Dacheng Tao²

¹ Key Lab of Machine Perception (MOE), Dept. of Machine Intelligence, Peking University.

² UBTECH Sydney AI Centre, School of Computer Science, Faculty of Engineering, University of Sydney.

³ Noah's Ark Lab, Huawei Technologies. ⁴ NELVT, Dept. of CS, Peking University. ⁵ Peng Cheng Laboratory.

{tianyuguo, shiboxin}@pku.edu.cn, xuchao@cis.pku.edu.cn, yunhe.wang@huawei.com
{c.xu, dacheng.tao}@sydney.edu.au, jhua7177@uni.sydney.edu.au

Abstract

This paper defines a positive and unlabeled classification problem for standard GANs, which then leads to a novel technique to stabilize the training of the discriminator in GANs. Traditionally, real data are taken as positive while generated data are negative. This positive-negative classification criterion was kept fixed all through the learning process of the discriminator without considering the gradually improved quality of generated data, even if they could be more realistic than real data at times. In contrast, it is more reasonable to treat the generated data as unlabeled, which could be positive or negative according to their quality. The discriminator is thus a classifier for this positive and unlabeled classification problem, and we derive a new Positive-Unlabeled GAN (PUGAN). We theoretically discuss the global optimality the proposed model will achieve and the equivalent optimization goal. Empirically, we find that PUGAN can achieve comparable or even better performance than those sophisticated discriminator stabilization methods.

1. Introduction

Recently, deep generative models have received remarkable achievements in image generation tasks [14, 22, 25, 5]. As a representative generative model, GANs [5] approximated a target distribution via playing a min-max game. In the standard framework of GAN [5, 23], a generator takes noise vectors from a prior distribution (e.g. Gaussian distribution and normal distribution) as the input and tends to produce data that follows the distribution of the reference natural images, while the discriminator aims to distinguish the generated data from the real data. Various GAN methods have been developed in many interesting applications. For example, in the image-to-image translation task, generators in GANs map the input image to output image. Representative methods include Pix2pix [10] over paired training

images and cycleGAN [30] in an unsupervised way.

In vanilla GANs, the training progress usually lacks stability, and the quality of generated images is not always satisfactory (e.g. model collapse). For instance, DCGAN [23] carefully designed the neural architectures for the generator and the discriminator to alleviate these problems. Progressive GAN [12] generated high-resolution images by progressively deepening the network. BigGAN [3] produced high-quality images by improving training methods, e.g. enlarging batch size, and truncating the latent space. WGAN [1] and WGAN-GP [6] tried to fit and optimize the Wasserstein distance to stabilize the generation process. SNGAN [21] proved the necessity and benefits of introducing Lipschitz continuity to the discriminator.

These aforementioned methods to stabilize GAN could be roughly divided into two categories: designing stable network structures and training strategies and developing new effective optimization goals. However, neither of them has stepped away from the positive-negative classification problem initially established in standard GAN. Although WGAN and WGAN-GP no longer take the discriminator as a classifier for real data and generated data, the aim of the is still to separate the real and generated data as far as possible. To the best of our knowledge, existing GAN models attempt to strictly distinguish between generated data and real data and ignore the fact that the quality of the generated samples is not the same. It is unfair to treat high-quality samples equally with low-quality samples, especially when high-quality samples are sufficiently realistic. Although there are many theoretical results proposed to justify the final equilibrium, such as vanilla GAN [5] proving the existing of the equilibrium and WGAN [1] replacing the JS divergence with the Wasserstein distance, these analysis mainly focus on the final achievement rather than the intermediate status in the training process.

In this paper, we suggest that instead of an ordinary positive and negative classification (i.e. real v.s. fake) problem, GAN is actually in the face of a positive-unlabeled classi-

fication problem. With adequate training, generated data could look real and may appear to be even more realistic than real data at times. It would then be illogical to make a stereotype of generated data as fake data. To catch up with the continuously improved quality of generated data, we thus take them as unlabeled data, which consist of low-quality data and high-quality data. These high-quality data are considered to be close to or even better than some real data. Within the framework of positive-unlabeled classification, the classification objective of standard GAN can be re-defined, and different variants can be easily obtained by considering different scoring functions (*e.g.* those in LSGAN [20] and HingeGAN [21]). In addition, we get rid of the class balance constraint (*i.e.* half of the sample are fake), and observe impressive performance improvement by increasing the share of generated data in the mini-batch. Our theoretical analysis suggests that the proposed new algorithm has a guaranteed final equilibrium. Experimental results on benchmark datasets demonstrate that we actually enjoy more stable training progress and thus achieve better generated samples.

2. Proposed Approach

In this section, we first review preliminary works about the standard Generative Adversarial Network (SGAN). Then we analyze the problem existing in GAN and define a new role for the discriminator D . We also theoretically develop this idea into a new algorithm within the framework of SGAN, and then extended this algorithm to the general GAN, which shows the flexibility of our method.

2.1. Preliminary

GAN [5] was introduced by Goodfellow et al. (2014). It consists of two neural networks: discriminator network D and generator network G . The discriminator D aims to distinguish the provided real data and the fake data generated by the generator G . On the other hand, the generator G aims to generate fake data that can fool the discriminator D . Following this adversarial manner, we expect the G can generate high-quality data in the end. Formally, the objective function of GAN can be written as

$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))] \quad (1)$$

where p_{data} indicates the distribution of real data, \mathbf{z} is the random noise sampled from a prior distribution $p_{\mathbf{z}}$ (*i.e.* the Gaussian distribution), and $D(\mathbf{x})$ is the predicted probability of \mathbf{x} to be real by the discriminator. Since the minimax objective function might lead to gradient vanishing for G when D can perfectly distinguish two data set. More of GAN's variants (*e.g.* WGAN [1] and LSGAN [19]) transform this minimax game into a non-saturating game. In general, the

objective functions of these GANs can be concluded as follows:

$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [f_1(D(\mathbf{x}))] - \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [f_2(D(G(\mathbf{z})))] \quad (2)$$

where f_1 is the loss of classifying input as real and f_2 is the loss of classifying input as fake.

2.2. Problem Analysis

As shown above, existing GAN variants are trained to separate the real and generated data strictly. However, this does not match the actual situation in training. Some of the generated samples can achieve higher quality and are more realistic than others. This phenomenon usually lasts until the end of the training. As a result, the quality of samples generated by G is very different, and there are many high-quality samples and a considerable proportion of low-quality samples. For instance, there is a certain proportion of unsatisfactory samples in a well-trained generation space, and the gaps in the generated samples of different quality are relatively large. It is natural to model $p_g(x)$ as a mixture distribution [2, 18]. Therefore, the traditional method of strictly distinguishing the real sample from the generated sample does not conform to the actual situation of the training. In this paper, we propose an algorithm that is dedicated to picking out low-quality samples from the samples generated by G and promoting them, unlike traditional discriminators D that are dedicated to distinguishing real samples from generated samples. Our algorithm encourages the discriminator D to divide the generation space of G into high-quality samples and low-quality samples so that the generator G could improve the low-quality samples.

By doing so, the proposed method enjoys several desired properties: i) The discriminator pays more attention to poor quality samples, allowing the generator to focus on improving the quality of these bad samples. As a result, the quality of generated samples is more balanced, and the overall quality is expected to be enhanced, ii) The training strategy of our algorithm is more in line with the actual situation of samples generated by G in the training process, so the training process is expected to be more stable, iii) More importantly, the proposed algorithm is a flexible method which means that our algorithm could be easily integrated into existing frameworks of variety GAN, and we will show this desirable feature in the following subsection, and iv) Although the proposed method changes the function of the D network, we provide some theoretical results in Section 3 which demonstrate the proposed method also enjoys the same equilibrium condition with the other GAN and provide some guarantee for performance.

2.3. Positive-Unlabeled classification (PU) in SGAN

As mentioned above, we propose to allow the discriminator D to treat the high-quality generated samples as real data and focus on the bad generated samples. The discriminator D is required to learn how to distinguish high-quality samples with other low-quality samples. Identifying high-quality samples from generated samples under the guidance of real samples is very similar to Positive-Unlabeled classification problems [4, 15, 27], where only some positive samples were labeled, and the classifier tried to find positive samples from unlabeled samples consisting of positive and negative samples. According to the solution of the PU classification problem, we develop an algorithm learning a discriminator to recognize high-quality samples. Firstly, we denote the generated data as \mathbf{x}_g , consisted with high-quality samples \mathbf{x}_{gr} and the bad samples \mathbf{x}_{gf} . And we consider both \mathbf{x}_{gr} and real data \mathbf{x}_{data} to be real (*i.e.* $y_{gr} = y_{data} = 1$) while consider \mathbf{x}_{gf} to be fake (*i.e.* $y_{gf} = -1$). In addition, denote $p(x)$ as the marginal density of \mathbf{x}_g and $p_{gr}(\mathbf{x}) = p(\mathbf{x}|1)$ and $p_{gf}(\mathbf{x}) = p(\mathbf{x}|-1)$ are the class conditional densities of \mathbf{x}_{gr} and \mathbf{x}_{gf} respectively. Then the p_g which is the marginal density of \mathbf{x}_g , can be obtained with:

$$p_g(\mathbf{x}) = \pi p_{gr}(\mathbf{x}) + (1 - \pi)p_{gf}(\mathbf{x}), \quad (3)$$

where π is the unknown *class prior* (*i.e.*, the proportion of \mathbf{x}_{gr} in \mathbf{x}_g). Now we successfully separate the generated space p_g into two parts. To classify \mathbf{x}_{gr} and \mathbf{x}_{gf} from \mathbf{x}_g with D as a binary classifier learned the distribution of \mathbf{x}_g from p , we need to minimize its expected miss-classification rate $R(D)$. The loss function for minimizing $R(D)$ by a given π could be:

$$\begin{aligned} \min_D R(D) &= \pi \mathbb{E}_{\mathbf{x} \sim p_{gr}(\mathbf{x})} [\ell(D(\mathbf{x}), 1)] \\ &+ (1 - \pi) \mathbb{E}_{\mathbf{x} \sim p_{gf}(\mathbf{x})} [\ell(D(\mathbf{x}), -1)], \end{aligned} \quad (4)$$

where $\ell(D(\mathbf{x}), t)$ is the loss function measuring the loss of prediction $D(\mathbf{x})$ when the ground true label is t . However, we have less idea about which is \mathbf{x}_{gr} . In our definition, the good samples are similar to the real data, which means that p_{gr} can be replaced by p_{data} . Thus, the $p_g(\mathbf{x})$ can be calculated by:

$$p_g(\mathbf{x}) = \pi p_{data}(\mathbf{x}) + (1 - \pi)p_{gf}(\mathbf{x}). \quad (5)$$

Similarity, the bad generated samples \mathbf{x}_{gf} are also unknown, and we can only access the generated samples \mathbf{x}_g and the real data \mathbf{x}_{data} . $R(D)$ should be modified to avoid the term of p_{gf} . From Eq. (5), the low-quality part p_{gf} can be expressed as follows,

$$(1 - \pi)p_{gf}(\mathbf{x}) = p_g(\mathbf{x}) - \pi p_{data}(\mathbf{x}). \quad (6)$$

Then we can find out the follow equation:

$$\begin{aligned} (1 - \pi) \mathbb{E}_{p_{gf}} [\ell(D(\mathbf{x}), -1)] &= \mathbb{E}_{p_g} [\ell(D(\mathbf{x}), -1)] \\ &- \pi \mathbb{E}_{p_{data}} [\ell(D(\mathbf{x}), -1)]. \end{aligned} \quad (7)$$

Algorithm 1 Implementing PU learning in GAN

Input: The number of D iterations pre G iteration n_d ($n_d = 1$ in normal), the batch size m , the class prior knowledge π for the proportion of positive data in unlabeled data.

Input: Initialize the parameters θ_g of the generator G and the parameters θ_d of the discriminator D .

while θ_g has not converged **do**

for $t = 1, \dots, n_d$ **do**

Randomly sample $\{\mathbf{x}^i\}_{i=1}^m \sim p_{data}(\mathbf{x})$;

Randomly sample $\{\mathbf{z}^i\}_{i=1}^m \sim p_z(\mathbf{z})$;

Sample $\{\mathbf{z}^{(i)}\}_{i=1}^m \sim \mathbb{P}_z$

Calculate $\widehat{R}_p^+ \leftarrow \frac{1}{m} \sum_{i=1}^m [f_1(D(\mathbf{x}^i))]$

Calculate $\widehat{R}_p^- \leftarrow \frac{1}{m} \sum_{i=1}^m [f_2(D(\mathbf{x}^i))]$

Calculate $\widehat{R}_u^- \leftarrow \frac{1}{m} \sum_{i=1}^m [f_2(D(G(\mathbf{z}^i)))]$

Update $\theta_d \leftarrow \nabla_{\theta_d} \pi \widehat{R}_p^+ + \max(0, \widehat{R}_u^- - \pi \widehat{R}_p^-)$

end for

Randomly sample $\{\mathbf{z}^i\}_{i=1}^m \sim p_z(\mathbf{z})$;

Update $\theta_g \leftarrow -\nabla_{\theta_g} \frac{1}{m} \sum_{i=1}^m [f_2(D(G(\mathbf{z}^i)))]$

end while

Output: A generator network G .

By combining Eqs. (4) and (7), the new objective function will be:

$$\begin{aligned} \min_D R(D) &= \pi \mathbb{E}_{p_{data}} [\ell(D(\mathbf{x}), 1)] \\ &+ \mathbb{E}_{p_g} [\ell(D(\mathbf{x}), -1)] - \pi \mathbb{E}_{p_{data}} [\ell(D(\mathbf{x}), -1)]. \end{aligned} \quad (8)$$

By minimizing Eq. (8), the discriminator D can distinguish not only \mathbf{x}_{gr} but also \mathbf{x}_{data} from \mathbf{x}_{gf} , by only learning the distribution of \mathbf{x}_{data} based on p_{data} and distribution of \mathbf{x}_g (the generated samples) from p_g . We notice that the second and third terms of Eq. (8) are introduced from Eq. (6) and aim to calculate the loss over p_{gf} . The original loss $(1 - \pi) \mathbb{E}_{p_{gf}} [\ell(D(\mathbf{x}), -1)]$ is expected to be not less than zero, but the replacement loss function $\mathbb{E}_{p_g} [\ell(D(\mathbf{x}), -1)] - \pi \mathbb{E}_{p_{data}} [\ell(D(\mathbf{x}), -1)]$, may be negative. This abnormal value of loss may lead to over-fitting. It is important to avoid the it to be negative. Finally, the objective function of discriminator proposed in Eq. (8) will be:

$$\begin{aligned} \max_D V(D) &= \pi \mathbb{E}_{p_{data}} [\log(D(\mathbf{x}))] \\ &+ \max\{0, \mathbb{E}_{p_z} [\log(1 - D(G(\mathbf{z})))] \\ &- \pi \mathbb{E}_{p_{data}} [\log(1 - D(\mathbf{x}))]\}. \end{aligned} \quad (9)$$

Here we obtain the proposed objective function of the discriminator. Considering there is also an adversarial game between the discriminator D and the generator G , the generator G should still be trained to deceive the discriminator D . As a result, we can easily lead to the objective function of G as follows,

$$\min_G V(G) = \mathbb{E}_{p_z} [\log(1 - D(G(\mathbf{z})))] \quad (10)$$

Following Eqs. (9) and (10), we reached our objective to deal with the generated data in different ways, rather than treating it all as negative samples.

Hou *et al.* [8] aims to use GAN to improve the PU classification, where two generators were introduced to fit the positive and the unlabeled distributions, respectively. By contrast, we aim to stabilize the training of GANs by treating the generated samples as unlabeled samples consisting of both high-quality and low-quality samples and formulating GAN itself as a PU classification. The main idea and purpose of the two papers are definitely different. To our best knowledge, we are the first paper to treat the generated samples as unlabeled samples and focus on improving the low-quality sample to learn a better generator.

2.4. PU classification for general GANs

Above we conclude our objective function within the standard GAN framework. The proposed method can also be integrated into other general GAN frameworks flexibly. In this part, we combine the proposed method with other loss functions of discriminator in GAN.

In general, the objective function Eq. (2) contains two loss functions f_1 and f_2 . Those concrete loss functions can be changed for a different variance of GANs, but all these loss functions are following the same concepts that $f_1(\mathbf{x})$ and $f_2(\mathbf{x})$ are trying to separate the real data from the generated data as far as possible. Similar to SGAN, we propose that the discriminator in GAN is better to focus on generated samples with low quality and recognize the high-quality samples from the generated samples. Following this concept, we implement the proposed method for the general framework of GAN with the following equation:

$$\begin{aligned} \max_G \min_D V(D, G) = & \pi \mathbb{E}_{p_{data}} [f_1(D(\mathbf{x}))] \\ & + \max \{0, \mathbb{E}_{p_z} [f_2(D(G(\mathbf{z})))]\} \\ & - \pi \mathbb{E}_{p_{data}} [f_2(D(\mathbf{x}))]. \end{aligned} \quad (11)$$

With the help of Eq. (11), we can now integrate the proposed method into various frameworks of GAN, such as WGAN-GP [24], LSGAN [19], and SpectralGAN [21]. This flexibility that combining with other models provides the proposed method a chance to get further improvement on existing excellent models. Loss functions corresponding to the specific model can be found in the supplementary material.

Although we finally got a new loss function, in the next section, we theoretically prove that our proposed algorithm is also designed to minimize the distance between the generated distribution and the real distribution, which provides a theory for the effectiveness of our algorithm.

3. Theoretical Analysis

In the proposed method, the discriminator D is encouraged to not only distinguish the real samples from the gener-

ated samples but also allow a certain proportion of generated samples of high quality to be recognized as real data, which reduces the instability problem during the training progress. Following this principle, we have obtained a novel loss function Eq. (9) and Eq. (10) in the framework of the standard GAN. Although Eq. (9) and Eq. (10) are designed to achieve the above-mentioned desirable characteristics, it is unclear whether the final convergence of the proposed method satisfies the requirements of the generation task. In this section, we provide a formal technical analysis of the convergence of the proposed objective function and prove that the proposed algorithm will perfectly lead the generated distribution to the real one. See the supplementary material for proof.

Now we consider the standard GAN based framework and analyze the optimal discriminator and generator. The discriminator D is optimized by Eq. (9). Following the analysis proposed in GAN [5], the optimal distribution D will balance between the true distribution p_{data} and the learned distribution p_{gf} .

Theorem 1. *For the generator G fixed, the optimal discriminator D is*

$$D^*(\mathbf{x}) = \frac{p_{data}(\mathbf{x})}{p_{data}(\mathbf{x}) + \frac{1-\pi}{\pi} p_{gf}(\mathbf{x})},$$

where $p_{gf}(\mathbf{x})$ is the distribution of low-quality generated samples of G .

With the optimal discriminator D fixed, we can reformulate the objective function by replacing $D(\mathbf{x})$ in Eq. (10) according to Theorem 1. By doing so, we can summarize the behavior of G in the following theorem.

Theorem 2. *With the optimal discriminator D fixed, the optimization of generator G is equivalent to minimize $\pi \log \pi + (1 - \pi) \log(1 - \pi) + \pi KL(p_{data} || p_g) + (1 - \pi) KL(p_{gf} || p_g)$.*

Theorem 2 suggests that the optimal generator G will pay attention to reducing the divergence within the generated space by minimizing the distance between p_g and p_{gf} . Moreover, the generator will also guide the generated distribution p_g as close as possible to the real distribution p_{data} , which ensures the quality of the generated sample. Combining Theorem 1 and Theorem 2, we can summarize the following corollary.

Corollary 1. *The global minimum of the proposed objective function $V(G, D)$ is achieved if and only if $p_{gf} = p_g = p_{data}$. At that point, $C(G)$ achieves the value of $\pi \log \pi + (1 - \pi) \log(1 - \pi)$, and $D(\mathbf{x})$ achieves the value of π .*

The above theoretical results prove that the proposed method will achieve equilibrium if and only if $p_g = p_{data}$, which points out that our method enjoys the same global equilibrium point as other GAN frameworks. These results justify our approach. The next experiment section further illustrates the effectiveness of our approach.

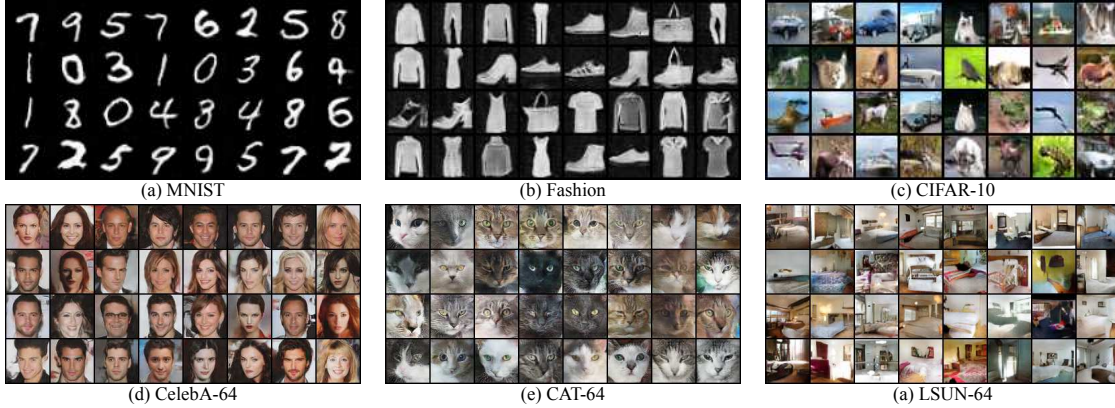


Figure 1: Generated samples obtained by the proposed method on image datasets.

Table 1: Performance comparison on several benchmark datasets.

Loss	SGAN	PUSGAN	HingeGAN	PUHingeGAN	LSGAN	PULSGAN	WGAN-GP	PUWGAN-GP
MNIST	18.65	16.53	21.48	16.91	13.47	13.96	17.42	14.77
Fashion	25.72	24.33	28.39	25.31	32.05	26.72	26.50	23.12
CIFAR-10	43.39	31.02	43.85	36.37	27.64	22.32	36.86	31.85
CelebA-64	48.44	43.89	46.13	43.94	51.67	47.80	36.09	35.72
CAT-64	46.13	16.90	29.52	25.72	57.22	26.79	21.86	17.35

4. Experiment

In this section, we evaluate the proposed method on a range of datasets including MNIST [17], FMNIST [26], CIFAR-10 [16], CAT [29], and LSUN-bedroom [28]. We resize images in the MNIST and FMNIST datasets to 32×32 for convenience. For these datasets with more than one kind of resolution, we mark them with the resolution, such as CAT-64. Experiments on CAT-128, CAT-256, CelebA-128, and LSUN-128 datasets are also conducted to evaluate the high-resolution generation ability of the proposed approach. Moreover, due to the limitation of computational resource, for LUSN-128 we randomly sample 100,000 images from the dataset as training set, instead of using all of them. The experiment is implemented in pytorch, and we use FID (Fréchet Inception Distance) [7] as the quantitative indicator to evaluate the performance based the quality of the generated results (lower value of FID indicating higher generated quality). FID scores are calculated with 10,000 generated samples and 10,000 real images randomly sampled from the dataset in advance.

As mentioned in Section 2, our approach enjoys a high degree of flexibility can be integrated into most kind of GAN frameworks. We chose some variants of GAN as basic frameworks, such as standard GAN (SGAN) [5], LSGAN [19], WGAN-GP [6], and HingeGAN [21], and then integrate our method into these frameworks for comparison. For a fair comparison, we follow the same settings and architectures of GANs when we integrate our method and make sure the loss functions are the only changed part. We also compare our method with Relativistic GAN [11], which is another flexible GAN framework, and we use the average version

(RaGAN) in the experiment. All objective function of the proposed frameworks could be found in the supplementary material.

All models used in experiments will be trained with Adam optimizer [13], and the random selecting seed is set to 1. In addition, the discriminator follows the CNN structure described by Miyato, *et al.* [21] while the generator will follow the structure of standard DCGAN [23] for all models generating images whose resolution less than 128 except WGAN-GP whose structure we leave in the supplementary material. We use the stable setting for DCGAN [23] as the *basic setting* for training, which the learning rate lr is set to 0.0002, the $\beta_1 = 0.5$ and $\beta_2 = 0.999$ for Adam optimizer, and the number of training time for discriminator and generator will both equal to 1. In addition, batch normalization [9] is also implemented. Moreover, we set a general case for the hyper-parameter π growth pattern called the *basic pattern* in this experiment. The basic pattern will be initialized π with 0.1 and will increase smoothly at each iteration until it reaches 0.7. Detailed network structures used on other datasets could be found in the supplementary material.

4.1. Quantitative Image Generation Results

In this section, we evaluate the generation ability of the proposed method on multi-category image sets MNIST [17], FMNIST [26], and CIFAR-10 [16] and single-category image sets CAT [29] and LSUN-bedroom [28]. In this part, the resolution of images in the MNIST, FMNIST, and CIFAR-10 is 32 and is 64 in the rest datasets. We choose four representative adversarial models SGAN, HingeGAN, LSGAN, and WGAN-GP as basic frameworks and compare them with our

Table 2: Stability experiment results (reported by FID) on the CAT dataset with three resolution.

Loss	64×64 images (N=9304)				128×128 images (N=6645)				256×256 images (N=2011)			
	Min	Max	Mean	SD	Min	Max	Mean	SD	Min	Max	Mean	SD
SGAN	19.32	95.31	48.36	33.92	-	-	-	-	-	-	-	-
RaSGAN	19.11	42.20	32.39	7.53	23.18	38.74	26.12	5.66	33.67	133.94	60.42	23.66
PUSGAN	12.71	18.99	16.38	1.99	17.76	28.91	23.61	3.43	35.75	63.68	50.76	9.9
LSGAN	21.84	69.03	39.22	13.64	22.75	62.39	43.48	12.68	-	-	-	-
RaLSGAN	15.73	25.20	19.85	4.47	16.52	43.87	23.74	7.33	42.02	282.15	72.32	86.11
PULSGAN	12.86	26.79	21.28	4.36	18.27	38.52	26.24	6.06	33.23	269.27	73.68	89.04
WGAN-GP	17.73	31.62	24.61	5.19	21.91	43.06	34.48	8.24	58.31	262.24	109.21	64.38
RaSGAN-GP	16.83	24.32	20.59	2.84	18.78	44.53	31.63	10.04	39.42	115.07	69.65	24.19
PUWGAN-GP	13.60	20.67	17.13	1.73	17.32	36.17	27.32	5.52	35.29	93.48	59.01	17.71

Table 3: Performance comparison based on FID on the CIFAR-10 dataset with different settings.

Loss	basic	$lr = .001$	No BN	Tanh
SGAN	43.39	74.75	47.83	57.42
RaSGAN	33.63	44.67	42.21	55.38
PUSGAN	31.02	40.93	37.85	54.70
LSGAN	27.64	56.36	40.81	68.40
RaLSGAN	23.48	35.75	37.28	55.92
PULSGAN	22.32	37.65	37.34	51.45
HingeGAN	43.85	41.66	39.57	58.43
RaHingeGAN	38.03	44.25	41.54	51.70
PUHingeGAN	36.37	35.60	38.09	50.84

algorithms, which are denoted as PUSGAN, PUHingeGAN, PULSGAN, and PUWGAN-GP, respectively.

Table 1 reports a comparison of the FID score obtained by the proposed method and basic models. Our models enjoy the ability to combine with most variants of GAN, which allows us to achieve the best performance on most data sets. Table 1 shows that most of our methods exceed their corresponding basic frameworks, which demonstrates the effectiveness and the flexibility of our approach. In Figure 1, we show a few images generated by the PUSGAN models. We observe that the proposed method generates high-quality images on various datasets, which is consistent with the quantitative results in Table 1.

4.2. Evaluating Stability

We evaluate the stability of the proposed method on three resolution of the CAT dataset, such as 64×64 , 128×128 , and 256×256 pixels. As there are only 6654 and 2011 samples in the CAT-128 dataset and CAT-256 dataset respectively, some variants of GAN are unable to converge on these datasets. We choose SGAN, LSGAN, and WGAN-GP as basic models. We compare the proposed method with both these basic models as well as the corresponding Relativistic GAN (RaGAN). For each model, We calculate the FID score of the current model every 10,000 iterations.

The results will be presented with the minimum, maximum, mean, and standard deviation (SD) of these obtained FID values. Table 2 shows the FID results for different networks in different resolutions of data sets.

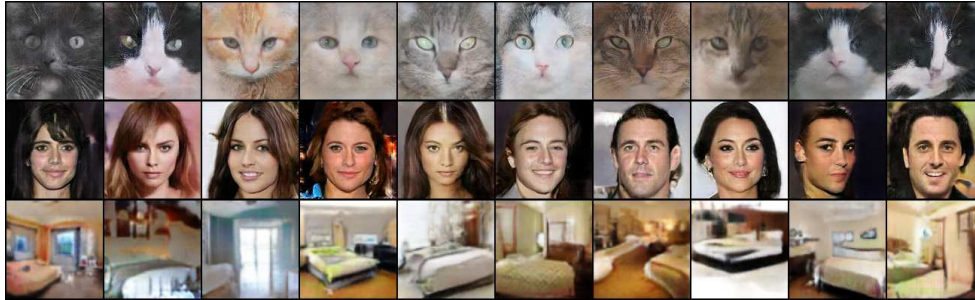
For 64×64 resolution dataset, all models trained by the proposed method except PULSGAN, can achieve much lower FID in minimum, maximum and mean compared with its original version and even can further achieve lower FID values than their relativistic versions, which indicates that our algorithm can effectively improve the training stability and improve the quality of the generation.

For higher resolution data sets, SGAN failed to converge in 128×128 and 256×256 resolution datasets while LSGAN will be stuck at the early stage in 256×256 resolution dataset [11]. The standard version of our model (PUSGAN) shows further stability with lower values of maximum, mean, and SD in all three resolution datasets. On the most challenging 256 resolution dataset, the proposed method achieves both the satisfactory quality and stability. In experiments, we found that although the PULSGAN can converge in the CAT-256 dataset, the convergence is much slower than other GANs. Nevertheless, PULSGAN still can achieve competitive results compared with other states of arts GANs such as WGAN-GP and RaGAN.

Overall, our algorithm presents desirable stability for all three data sets, and it can achieve similar or even better results compare to relativistic versions. It is impressive that all these GANs trained by the proposed method have improved. As a result, we conclude that the above stability experiment demonstrates that the proposed method provides could provide stability for the training progress for a variety of GANs and thus improve the quality of generated images.

4.3. Evaluating on Hard Training Setting

As we have claimed, our approach focuses on improving low-quality samples and lead to more stable training progress. Thus our approach enjoys the ability to generalize to many training settings. To demonstrate this, we evaluate the proposed method on several hard training settings and compare it with the other GAN frameworks. In this



(a) Generation samples on 128 resolution datasets.



(b) Generation samples on the CAT-256 dataset.

Figure 2: High-resolution generated samples.




Figure 3: High-resolution interpolation results.

part, we implement SGAN (RaSGAN, PUSGAN), LSGAN (RaLSGAN, PULSGAN), and HingeGAN (RaHingeGAN, PUHingeGAN) on the CIFAR-10 dataset. The experiment is conducted on one basic setting and three hard settings. The basic setting is same as above, and three hard settings are i) changing learning rate to 0.001 ($lr=0.001$), ii) removing Batch normalization layers in G and D (No BN), and iii) replacing all activation functions with Tanh in G and D (Tanh).

The results are showed in Table 3. In the stable setting, we can find that the PUSGAN has better performance than SGAN and its other variants. The PUHingeGAN has a huge improvement compared with the original HingeGAN with a gap of 14, and it also performs better than RaHingeGAN. On the other hands, the performance of PULSGAN is slightly worse than other LSGAN versions.

When the learning rate is increased to 0.001, all three PU versions of GANs perform well, compared with the original one. While PUSGAN and PUHingeGAN can perform better than their relativistic versions. However, by changing optimization settings such as removing batch normalization or replace ReLU activation function with Tanh activation function (No BN and Tanh in columns respectively), the performances of PUGANs will be worse. It might indicate that the PUGANs rely on optimization terms for stable training.

4.4. High-resolution Results

Generating high-resolution images is a complicated task. To demonstrate the generation ability of our algorithm, we evaluate the proposed method on CAT-128, CelebA-128, and

LSUN-128 datasets with 128×128 pixels and the CAT-256 dataset with 256×256 pixels. There are 202,599 images in the CelebA-128 dataset, and 3,033,042 in the LSUN-128 dataset (only 100,000 samples are used for training). As mentioned above, the CAT-128 and CAT-256 datasets are considered as more challenging high-resolution datasets because there are only 6,645 and 2,011 samples, respectively. Images shown in Figure 2 are generated by the proposed method within the architecture of SGAN (PUSGAN), while SGAN failed to generate such high-resolution images, especially on the CAT-256 dataset. Moreover, interpolation is also an impressive feature of the generative models, which indicates that the generative model successfully learns to fit the distribution of natural images instead of overfitting to the training samples. We show high-resolution interpolation results obtained by the proposed method in Figure 3. It shows that our model generates smooth interpolation images. Figures 2 and 3 demonstrate that the proposed method could provide improvement on the quality of generated samples.

4.5. Evaluating the impact of class prior π

In Eq. (3), we introduce a class prior π into our algorithm. The π indicates the proportion of high-quality fake data in fake data, and we treat it as a hyper-parameter. In this section, we further evaluate the impact of class prior π with PUSGAN framework on the CAT-64 dataset. The structure and training settings are the same as the previous sections. We set four different increasing patterns for π during training. The first pattern is the basic pattern we used in previous sections. The




Figure 4: The trend of Fréchet Inception Distance (FID) of PUSGAN using different π growth pattern.

second version will set π to be 0.3 at the beginning, and it will be increased with 0.1 at every 10k iterations until 0.7 is reached. The second version is used to evaluate the impact of the fast growth of π . For the third and fourth patterns, π will be fixed at 0.3 and 0.5 during the training process.

In Figure 4, we found that the fast-growing pattern achieves the worst average performance, and its FID scores remain relatively high. As a comparison, the basic pattern can reach lower FID values than the fast pattern, and it is also relatively stable in the later training stage. The fixed π value of 0.3 could present more stability and generate a competitive result, compared with the previous two patterns. In addition, the fourth version has the worst performance at the beginning, but it was keep going better and achieved the best performance in the end, within all four patterns. The shortage of this version mainly lies in the large fluctuations and slow convergence. It is interesting that all first three patterns have similar FID values at the early stage, while the performance with a higher π is much worse at the same stage. This may be because the proportion of high-quality samples in the G network at the beginning of training is far from 0.5, and setting π to 0.5 is against the real situation, leading to an unstable training. On the other hands, the fourth one achieves a lower FID value in the end, while others have similar values at the same time. It might show a too large (e.g. 0.7), or a too small (e.g. 0.3) values of π reduce the performances. As a result, the proposed method enjoys considerable tolerance for the selection of hyperparameters.

4.6. Evaluating the impact of increasing number of fake data for training

Normally, the size of real data is much smaller than the size of the generated data in the adversarial generative task. It is an interesting problem about how to make the most of this large amount of generated data. In general training progress of GAN, the batch size of real samples and generated samples are the same. Here, we try to increase the batch size of the generated data and maintain that of the actual data to take advantage of this large number of generated samples. We investigate the impact of increasing the batch size of the generated data for training. The evaluation is based on




Figure 5: The trend of Fréchet Inception Distance (FID) of PUSGAN trained by different batch sizes of fake data.

three versions of PUSGAN with different batch sizes of fake data. The first version is the basic version that the batch size of real and fake data are the same. The second and third versions will use twice and three times more fake data than real data, respectively. The structure and training settings are the same as the one we used in previous sections. We report these interesting results Figure 5.

From results, we find that the second and third versions of PUSGAN can reach their best performance at the very beginning, which proves that PUSGAN can converge faster by increasing the batch size of fake data. Although it shows that an increase in the number of false samples can provide a small performance boost and faster convergence, it seems to be detrimental to stability. Considering the stability and for a fair comparison, we insist on using the same batch size for both the real and generated data in the above experiments.

5. Conclusion

In this paper, we present a positive-unlabeled generative adversarial network (PUGAN), where the discriminator is trained to recognize the high-quality samples from the generated data. The proposed method addresses problems in traditional methods that neglecting the gradual increase in sample quality and the imbalance of generated sample quality, which provides more stable training progress and higher generation quality. We further demonstrate that our approach has the flexibility to combine with most existing GAN frameworks without additional computational cost. Experiments conducted on real-world image datasets suggest that the proposed method successfully improve both the stability and the quality of generated samples. We also provide theoretical results to illustrate the justification of our approach.

Acknowledgment

We thank anonymous area chairs and reviewers for their helpful comments. This research was supported in part by National Natural Science Foundation of China under Grant No. 61876007 and 61872012, Beijing Academy of Artificial Intelligence (BAAI), National Key R&D Program of China (2019YFF0302902), and Australian Research Council Grant DE-180101438.

References

- [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *Proc. of International Conference on Machine Learning*, pages 214–223, 2017. 1, 2
- [2] Javad Behboodian. On the modes of a mixture of two normal distributions. *Technometrics*, 12(1):131–139, 1970. 2
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 1
- [4] Marthinus C Du Plessis, Gang Niu, and Masashi Sugiyama. Analysis of learning from positive and unlabeled data. In *Advances in neural information processing systems*, pages 703–711, 2014. 3
- [5] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 1, 2, 4, 5
- [6] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in neural information processing systems*, pages 5767–5777, 2017. 1, 5
- [7] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, Günter Klambauer, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a nash equilibrium. *arXiv preprint arXiv:1706.08500*, 2017. 5
- [8] Ming Hou, Brahim Chaib-Draa, Chao Li, and Qibin Zhao. Generative adversarial positive-unlabelled learning. *arXiv preprint arXiv:1711.08054*, 2017. 4
- [9] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. 5
- [10] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *arXiv preprint*, 2017. 1
- [11] Alexia Jolicoeur-Martineau. The relativistic discriminator: a key element missing from standard GAN. In *Proc. of International Conference on Learning Representations*, 2019. 5, 6
- [12] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 1
- [13] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [14] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 1
- [15] Ryuichi Kiryo, Gang Niu, Marthinus C du Plessis, and Masashi Sugiyama. Positive-unlabeled learning with non-negative risk estimator. In *Advances in neural information processing systems*, 2017. 3
- [16] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. 2009. 5
- [17] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998. 5
- [18] Bruce G Lindsay. Mixture models: theory, geometry and applications. In *NSF-CBMS regional conference series in probability and statistics*, pages i–163. JSTOR, 1995. 2
- [19] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proc. of International Conference on Computer Vision*, 2017. 2, 4, 5
- [20] Xudong Mao, Qing Li, Haoran Xie, Raymond Y. K. Lau, and Zhen Wang. Multi-class generative adversarial networks with the L2 loss function. *CoRR*, abs/1611.04076, 2016. 2
- [21] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018. 1, 2, 4, 5
- [22] Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*, 2016. 1
- [23] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. 1, 5
- [24] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016. 4
- [25] Aaron van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixelcnn decoders. In *Advances in neural information processing systems*, 2016. 1
- [26] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017. 5
- [27] Yixing Xu, Chang Xu, Chao Xu, and Dacheng Tao. Multi-positive and unlabeled learning. In *IJCAI*, 2017. 3
- [28] Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 5
- [29] Weiwei Zhang, Jian Sun, and Xiaoou Tang. Cat head detection-how to effectively exploit shape and texture features. In *Proc. of European Conference on Computer Vision*, pages 802–816. Springer, 2008. 5
- [30] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *arXiv preprint*, 2017. 1