

Improving the Fairness of the Min-Max Game in GANs Training

Zhaoyu Zhang, Yang Hua, Hui Wang, Seán McLoone
 Queen’s University Belfast

{zzhang55, Y.Hua, h.wang, s.mcloone}@qub.ac.uk

Abstract

Generative adversarial networks (GANs) have achieved great success and become more and more popular in recent years. However, understanding of the min-max game in GANs training is still limited. In this paper, we first utilize information game theory to analyze the min-max game in GANs and introduce a new viewpoint on the GANs training that the min-max game in existing GANs is unfair during training, leading to sub-optimal convergence. To tackle this, we propose a novel GAN called Information Gap GAN (IGGAN), which consists of one generator (G) and two discriminators (D_1 and D_2). Specifically, we apply different data augmentation methods to D_1 and D_2 , respectively. The information gap between different data augmentation methods can change the information received by each player in the min-max game and lead to all three players G , D_1 and D_2 in IGGAN obtaining incomplete information, which improves the fairness of the min-max game, yielding better convergence. We conduct extensive experiments for large-scale and limited data settings on several common datasets with two backbones, i.e., BigGAN and StyleGAN2. The results demonstrate that IGGAN can achieve a higher Inception Score (IS) and a lower Fréchet Inception Distance (FID) compared with other GANs. Codes are available at <https://github.com/zzhang05/IGGAN>

1. Introduction

Generative adversarial networks (GANs) [10] are a form of generative model consisting of a generator (G) and a discriminator (D). Specifically, G produces synthetic data with some given noise, while D distinguishes whether the data is from the generator’s output or real data.

GANs can produce visually appealing samples and have become more and more popular in image video synthesis tasks [4, 21, 23, 41, 45, 46]. Although GANs have achieved impressive results in recent years [5, 12, 14–17, 26, 27, 30, 44], understanding of the min-max game in GANs training is still limited [37, 39]. In this paper, we first apply information theory [34] to analyze the min-max game in GANs. We

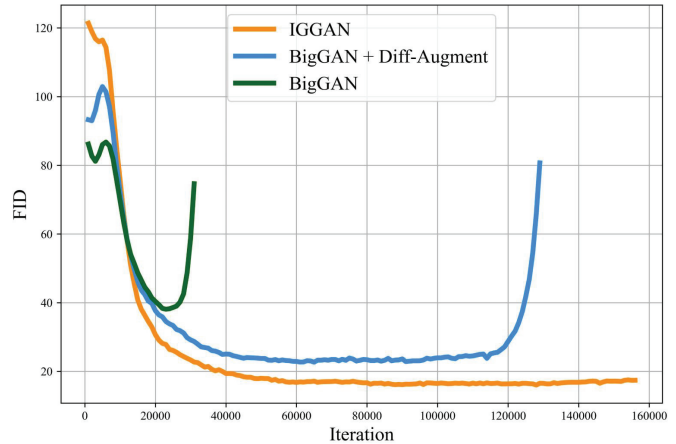


Figure 1. Training Fréchet Inception Distance (FID) [11] (lower is better) curves of BigGAN [3], BigGAN + Diff-Augment [50] and IGGAN on the 10% CIFAR-10 dataset. In this figure, IGGAN is based on the BigGAN backbone and achieves better convergence. *Best viewed in color.*

demonstrate that this game is unfair because the G obtains incomplete information while the D s always obtain complete information during the training of GANs. This can cause sub-optimal convergence in GANs [37, 39].

To assist with understanding this issue, we provide a brief illustration as follows. In the GANs model, the aim of the G is to make the D reward a high score for $D(G(z))$ [10], where z is the input noise of G . During the training of the unfair min-max GANs, there exists a type of generator output $G^*(z)$ which can cause D to reward a high score for $D(G^*(z))$ [1], indicating that G has deceived D successfully for this situation. We represent the support of this type of $G^*(z)$ under discriminator D as $P_D(G^*(z))$. In this case, producing samples based on $P_D(G^*(z))$ is the dominant strategy [6] for G in the min-max game. These generated similar distribution samples can be regarded as bad samples during training [36], which may finally cause the well-known mode collapse problem [22], leading to sub-optimal convergence [37].

To address this, we propose a novel GAN called Infor-

mation Gap GAN (IGGAN), a three-player min-max game consisting of one generator and two discriminators. We apply different data augmentation methods to D_1 and D_2 , respectively. The different information provided by different data augmentations causes the information gap between D_1 and D_2 . Because of this, all three players, i.e., G , D_1 and D_2 , obtain incomplete information in the min-max game. The incomplete information obtained by each player enables IGGAN to improve the fairness of the min-max game during training. In this case, there no longer exists a dominant strategy for G in the min-max game; therefore, G can produce more diverse images during training. To better understand this, suppose there exists a type of generator output $G^*(z)$ that can cause the output of one discriminator (D_1 or D_2) to reward a high score, this type of $G^*(z)$ typically cannot cause another discriminator to reward a high score due to the information gap between D_1 and D_2 , which can prevent G from producing similar distribution samples. In other words, IGGAN throws away the less diverse (bad) samples during training, similar to top-k GAN [36]. Hence, IGGAN can achieve a better optimization between the distribution of G (P_G) and the distribution of real data (P_{data}), leading to better convergence, as shown in Figure 1.

To sum up, the main contributions of this paper are as follows.

1. By applying information min-max game theory [34] to analyze the existing GANs, we are the first to unveil that the min-max game in existing GANs is unfair. We believe it is one of the core issues in GANs training and, therefore, expect that this finding can inform future research on GANs.
2. To improve the min-max game in GANs, we propose a new GAN called IGGAN, which consists of one generator (G) and two discriminators (D_1 and D_2). We apply different data augmentation methods to D_1 and D_2 , resulting in the information gap between D_1 and D_2 , which can improve the fairness of the min-max game during training, thus yielding better convergence.
3. Experiments on several datasets, i.e., CIFAR-10/100 [18], STL10 [7], CelebA [24], FFHQ [17] and LSUN-CAT [43] with two backbones, i.e., BigGAN [3] and StyleGAN2 [17], demonstrate that IGGAN can obtain a higher Inception Score (IS) [31] and a lower Fréchet Inception Distance (FID) [11] compared with other state-of-the-art GANs.

2. Background

2.1. Generative Adversarial Networks (GANs)

Generative adversarial networks (GANs) [10] consist of a generator G and a discriminator D . Given the real image

data x drawn from the distribution p_{data} and a prior on the input noise $p_z(z)$, G attempts to generate an output image $G(z)$ that confuses D into believing that $G(z)$ comes from p_{data} . In contrast, D attempts to distinguish between samples from x and $G(z)$. Then, the two-player min-max game between D and G is formulated as:

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

The parameters of G and D are updated iteratively with gradient descent methods. The equilibrium can finally be reached when D cannot differentiate between the p_{data} and $p_z(z)$.

2.2. Training GANs with Augmentations

Recent successes in GANs have affirmed the importance of using data augmentation in GANs training [38]. We classify this type of GANs as GANs with augmentations. According to the data augmentation methods used, augmentation GANs can be categorized as either positive data augmentation (PDA) GANs [15, 38, 47, 50, 51] or negative data augmentation (NDA) GANs [35]. PDA-GANs guide the discriminator to avoid overfitting, while NDA-GANs lead the discriminator to learn the out-of-distribution samples in order to improve GANs training. These two types of augmentation methods can both benefit the training of GANs.

2.3. Training GANs with Multiple Discriminators

To improve the training of GANs, GANs with multiple discriminators [32, 33, 48, 49] have been proposed. These consist of one generator and several discriminators. The typical example is the D2GAN [28] which consists of one generator and two discriminators. D2GAN is different from the original GANs as it is a three-player min-max game. The generator G aims to produce realistic-looking samples to fool both of the discriminators. The first discriminator D_1 rewards high scores for samples from the data distribution, while the second one D_2 favors samples from the generator. All three players are parameterized by neural networks, wherein D_1 and D_2 do not share their parameters.

2.4. Complete and Incomplete Information Game

In game theory, a complete information game [34] is a game in which knowledge about other players is available to all participants. Complete information is the concept that each player in the game is aware of the sequence, strategies, and pay-offs throughout gameplay. On the contrary, a game with incomplete information [34] is a game where the players do not have common knowledge of the game being played.

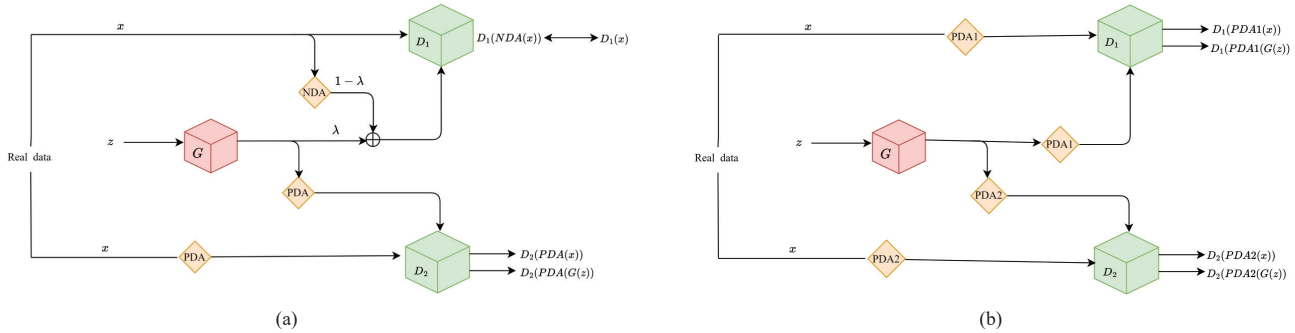


Figure 2. The overview of the IGGAN. (a) IGGAN with NDA in D_1 and PDA in D_2 : IGGAN (NDA + PDA). The information gap between PDA and NDA can help improve the fairness of the min-max game in the GAN. We use hyperparameter λ to balance the NDA data and fake data, and we directly apply the PDA to both real and fake samples as in Diff-Augment [50] and ADA [15] to avoid the leaking of augmentations [15]. Note that \leftrightarrow indicates a loss term pushing pairs to be apart. (b) IGGAN with different PDAs in D_1 and D_2 : IGGAN (PDA + PDA). The information gap between different PDAs can help improve the fairness of the min-max game in the GANs. We also directly apply the PDA to both real and fake samples to avoid leaking of augmentations. *Best viewed in color.*

3. Methodology

3.1. Analysis of the Unfair Min-Max Game in GANs

In the GANs min-max game, each player takes action in its game round. We define the game round of the generator and discriminator as R_G and R_D , and the action of the generator and discriminator as A_G and A_D , respectively. We then analyze the unfair min-max game problem in GANs as follows.

GANs can be regarded as a two-player min-max game where the two players are G and D . In R_G , G cannot know all the information of D for two reasons. First, according to Algorithm 1 of the original GANs [10], only the fake samples are used to optimize the parameters of G . On the contrary, both real and fake samples are used to optimize the parameters of D . Second, based on the observation as in [2, 20, 40], G only obtains part of the information of D in the min-max game. In this case, A_G is only based on incomplete information in the min-max game. In contrast, in R_D , D knows all the information of G (i.e., the input and output of G) and real data. Therefore, A_D is based on complete information in the min-max game. To sum up, in GANs, G obtains incomplete information, while D obtains complete information in the GANs min-max game, which means that the min-max game is unfair.

Recently, many techniques have been applied to GANs to improve GANs training, including employing augmentations or multiple discriminators. In the next section, we will show that these modifications do not prevent the D s from obtaining complete information in the min-max game.

GANs with multiple discriminators. We select the most simple GANs with multi discriminators, i.e., D2GAN [28], for our analysis. D2GAN can be regarded as a three-player min-max game where the three players are G , D_1 , and D_2 .

In R_G , G cannot know all the information of D_1 and D_2 . Thus, A_G is based on incomplete information in the min-max game. In contrast, in R_{D_1} , D_1 knows all the information of G and D_2 , because the real data and fake data to D_1 and D_2 are the same. Therefore, A_{D_1} is based on complete information in the min-max game. Similarly, in R_{D_2} , D_2 knows all the information of G and D_1 , hence, A_{D_2} is also based on complete information in the min-max game. In summary, in the GANs with multi-discriminator, G obtains incomplete information, while D_1 and D_2 obtain complete information, which means their min-max game is unfair.

GANs with augmentations. Compared to GANs, GANs with augmentations apply data augmentations to improve the training. The data augmentations do not influence information obtained by the players in the min-max game. Therefore, for the player G , A_G is based on incomplete information. On the contrary, for the player D , A_D is based on complete information in the min-max game. To conclude, in augmentation GANs, G obtains incomplete information, while D obtains complete information, which means the min-max game in the GANs with augmentations is unfair.

Based on the theory developed in [37], this unfair min-max game can harm the training of GANs and finally cause sub-optimal convergence in GANs.

3.2. IGGAN

The overview of IGGAN is shown in Fig.2. IGGAN consists of one generator and two discriminators with different data augmentations. According to the type of data augmentation methods, IGGAN consists of two cases. In case 1, as shown in Figure 2(a), NDA is applied to the first discriminator D_1 to produce the out-of-distribution samples. At the same time, PDA is applied to the second discrimina-

tor D_2 to improve the training. We utilize the information gap between NDA and PDA to improve the fairness of the min-max game. In case 2, as shown in Figure 2(b), different PDAs are applied to D_1 and D_2 to produce the information gap, respectively. We utilize the information gap between different PDAs to improve the fairness of the min-max game. According to NDA-GAN [35], all of the NDA methods aim to produce similar out-of-distribution samples. Thus, we do not apply different NDAs in IGGAN because it cannot produce the desired information gap (referring to the experimental results in Table 11). Next, we show that the information gap in the IGGAN can indeed help improve the fairness of the min-max game as follows.

Although the architecture of the IGGAN is similar to the D2GAN [28], the min-max game it employs is different. In R_G , G cannot know all the information of D_1 and D_2 . Thus, A_G is only based on incomplete information in the min-max game. At the same time, in R_{D_1} , D_1 knows all the information of G but does not know all of the information of D_2 , since D_1 cannot know the data augmentation information of D_2 . Therefore, A_{D_1} is based on incomplete information in the min-max game. Similarly, in R_{D_2} , D_2 knows all the information of the G but does not know all the information of D_1 , because D_2 cannot know the data augmentation information of D_1 . In this case, A_{D_2} is also based on incomplete information in the min-max game. Therefore, in IGGAN, all players obtain incomplete information which improves the fairness of the min-max game. To assist with understanding how improving the fairness of the min-max game, as achieved in IGGAN, can benefit the training of GANs, we conduct a brief experiment on the 10% CIFAR-10 dataset using the BigGAN backbone, the results of which are shown in Fig 1. IGGAN can achieve better convergence.

We further demonstrate that the changes of information for each player do not influence IGGAN to reach the Nash equilibrium in both cases, as shown in Theorem 1.

Theorem 1. Theoretical analysis of IGGAN.

Case 1: IGGAN with NDA in D_1 and PDA in D_2

Let $\bar{P} \in \mathbb{P}(\chi)$ be any distribution over χ with disjoint support than p_{data} , such that $supp(p_{data}) \cap supp(\bar{P}) = \emptyset$. Let $P^T \in p_{data}$ be any distribution over real data. Let $D_1 : \chi \rightarrow \mathbb{R}$ and $D_2 : \chi \rightarrow \mathbb{R}$ be the set of discriminators over χ , $f : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ be a convex, semi-continuous function such that $f(1) = 0$, f^* be the convex conjugate of f , f' be the derivative of f , G_θ be a distribution with sample space χ , and $G_\theta^T \in G_\theta$ be any distribution over sample space χ . T is one kind of PDA method. Then $\forall \lambda \in (0, 1]$, we have

$$\begin{aligned} & \underset{G_\theta \in P(\chi)}{\operatorname{argmin}} \max_{D_1, D_2 : \chi \rightarrow \mathbb{R}} L_f(G_\theta, D_1, D_2) \\ &= \underset{G_\theta \in P(\chi)}{\operatorname{argmin}} \max_{D_1, D_2 : \chi \rightarrow \mathbb{R}} L_f(\lambda G_\theta + (1 - \lambda)\bar{P}, D_1, D_2) \quad (2) \\ &= p_{data}, \end{aligned}$$

where $L_f(G_\theta, D_1, D_2) = E_{x \sim p_{data}}[D_1(x)] - E_{x \sim G_\theta}[f^*(D_1(x))] + E_{x \sim p_{data}}[D_2(T(x))] - E_{x \sim G_\theta}[f^*(D_2(T(x)))]$ is the objective function for IGGAN following NDA-GAN [35] and f-GAN [29]. The optimal discriminators for D_1 and D_2 are different, shown as follows:

$$\begin{aligned} & \underset{D_1 : \chi \rightarrow \mathbb{R}}{\operatorname{argmax}} L_f(\lambda G_\theta + (1 - \lambda)\bar{P}, D_1) \\ &= f'(p_{data}/(\lambda G_\theta + (1 - \lambda)\bar{P})). \end{aligned} \quad (3)$$

$$\underset{D_2 : \chi \rightarrow \mathbb{R}}{\operatorname{argmax}} L_f(G_\theta, D_2) = f'(P^T/G_\theta^T). \quad (4)$$

Proof. See supplementary materials.

Case 2: IGGAN with different PDAs in D_1 and D_2

Let $P^{T_1}, P^{T_2} \in p_{data}$ be any distribution over real data. Let $D_1 : \chi \rightarrow \mathbb{R}$ and $D_2 : \chi \rightarrow \mathbb{R}$ be the set of discriminators over χ , $f : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ be a convex, semi-continuous function such that $f(1) = 0$, f^* be the convex conjugate of f , f' be the derivative of f , G_θ be a distribution with sample space χ , and $G_\theta^{T_1}, G_\theta^{T_2} \in G_\theta$ be any distribution over sample space χ . T_1 and T_2 are different PDA methods. Then $\forall \lambda \in (0, 1]$, we have

$$\underset{G_\theta \in P(\chi)}{\operatorname{argmin}} \max_{D_1, D_2 : \chi \rightarrow \mathbb{R}} L_f(G_\theta, D_1, D_2) = p_{data}, \quad (5)$$

where $L_f(G_\theta, D_1, D_2) = E_{x \sim p_{data}}[D_1(T_1(x))] - E_{x \sim G_\theta}[f^*(D_1(T_1(x)))] + E_{x \sim p_{data}}[D_2(T_2(x))] - E_{x \sim G_\theta}[f^*(D_2(T_2(x)))]$ is the objective function for IGGAN following NDA-GAN [35] and f-GAN [29]. The optimal discriminators for D_1 and D_2 are different, shown as follows:

$$\underset{D_1 : \chi \rightarrow \mathbb{R}}{\operatorname{argmax}} L_f(G_\theta, D_1) = f'(P^{T_1}/G_\theta^{T_1}). \quad (6)$$

$$\underset{D_2 : \chi \rightarrow \mathbb{R}}{\operatorname{argmax}} L_f(G_\theta, D_2) = f'(P^{T_2}/G_\theta^{T_2}). \quad (7)$$

Proof. See supplementary materials.

4. Experiment

We demonstrate the superiority of IGGAN by comparing it with other state-of-the-art GANs on several datasets, i.e., CIFAR-10 [18], CIFAR-100 [18], STL-10 [7], CelebA [24], FFHQ [17] and LSUN-CAT [43]. More details of experiments can be found in the supplementary materials.

4.1. Datasets Preparation

We select vanilla BigGAN [3] and vanilla StyleGAN2 [17] as our backbones. In IGGAN, the two discriminators have the same architecture. Moreover, if we apply the NDA in IGGAN, following NDA-GAN [35], we set hyperparameter λ as 0.25 for the CIFAR-10, STL-10 and CelebA datasets, and 0.5 for the CIFAR-100 dataset.

For the BigGAN [3] backbone, we follow NDA-GAN [35] to prepare datasets: (a) CIFAR-10 [18] contains 60K

Method	IS [31]	FID [11]
BigGAN [3]	-	18.64
BigGAN + Diff-Augment [50]	-	15.23
CR-BigGAN [47]	-	14.56
GN-BigGAN [42]	8.72	13.71
ICR-SNGAN [51]	-	13.36
NDA-GAN (Jigsaw) [35]	-	12.61
Vision-aided-BigGAN [19]	-	11.17
DAG-GAN [38]	-	10.89
IGGAN (PDA + PDA)	8.96	11.04
IGGAN (NDA + PDA)	8.99	10.68

Table 1. IS (higher is better) and FID (lower is better) on the CIFAR-10 dataset with unconditional generation. For a fair comparison, **IS and FID are measured using 10k samples; the test set is the reference distribution.** Here, BigGAN is selected as the backbone for all the methods. The Best FID of all the methods is reported in the Table.

32×32 images with 10 labels, out of which 50K are used for training, and 10K are used for testing; (b) CIFAR-100 [18] contains 60K 32×32 images with 100 labels, out of which 50K are used for training, and 10K are used for testing; (c) CelebA [24] contains 162,770 training images and 19,962 test images, which are resized to 64×64 ; (d) STL-10 [7] contains 100K (unlabeled) training images and 8K (labeled) test images, which are resized to 32×32 . In our experiments, following the number of images in the test set of each dataset, we use 10K generated images for CIFAR-10, 10K for CIFAR-100, 19,962 for CelebA, and 8K for STL-10. **The test set for each dataset is used as the reference distribution for FID calculation, following prior work [47].** For the SyleGAN2 [17] backbone, we follow StyleGAN2 + ADA [15] and Diff-Augment [50] to prepare datasets (FFHQ [17] and LSUN-CAT [43]). FFHQ contains 70K images and LSUN-CAT contains 200K images. The images in FFHQ and LSUN-CAT are resized to 256×256 . Following Diff-Augment [50], we perform experiments on 30K, 10K, 5K and 1K training samples. **The full dataset is used as the reference distribution for FID calculation, following prior work [50].**

4.2. Results using BigGAN Backbone

4.2.1 Results on several standard datasets

The results on the commonly used CIFAR-10 dataset are shown in Table 1 (unconditional results) and Table 2 (conditional results). We compare two cases of IGGAN with other state-of-the-art GANs on BigGAN backbone. For IGGAN (NDA + PDA), we apply the widely used Diff-Augment [50] as PDA to avoid leaking of augmentations [15] during training, and one of the methods shown in Table III as NDA. The results of applying different NDA methods,

Method	IS [31]	FID [11]
BigGAN [3]	9.06	11.51
GN-BigGAN [42]	9.22	10.05
CR-BigGAN [47]	-	11.48
NDA-BigGAN (Jigsaw) [35]	-	9.42
BCR-BigGAN [51]	9.29	9.21
BigGAN + Diff-Augment [50]	9.22	8.47
Cntr+BCR-BigGAN [52]	9.41	8.30
Vision-aided-BigGAN [19]	-	8.27
IGGAN (PDA + PDA)	9.32	8.26
IGGAN (NDA + PDA)	9.43	8.15

Table 2. IS (higher is better) and FID (lower is better) on the CIFAR-10 dataset with conditional generation. For a fair comparison, **IS and FID are measured using 10k samples; the test set is the reference distribution.** Here, BigGAN is selected as the backbone for all the methods. The Best FID of all the methods is reported in the Table.

i.e., Jigsaw, Stitching, Mixup, and Cutmix in the IGGAN on several standard datasets are shown in Table 4, we report the best NDA and PDA combination results in IGGAN in Tables 1 and 2 for the comparison. For IGGAN (PDA + PDA), based on the conclusion in [52], we apply Translation and Cutout as PDA, respectively, to avoid leaking of augmentations [15]. As shown in Tables 1 and 2, IGGAN (NDA + PDA) obtains the highest Inception Score (IS) and lowest FID compared with existing state-of-the-art GANs on both unconditional and conditional CIFAR-10 datasets.

Furthermore, to further show the effectiveness of the information gap between D_1 and D_2 , we build several GAN experimental settings based on the PDA and NDA for the ablation study, as shown in Table 5. In Table 5, we do not apply different NDAs in one D as our experimental settings because NDA-GAN [35] has already shown that applying different NDAs in one D can achieve lower performance. At the same time, we also do not apply different PDAs in one D because DAG-GAN [38] has demonstrated that different PDAs in one D can result in a worse performance. We compared the two types of IGGAN with other GAN settings on CIFAR10/100, STL10 and CelebA datasets, and the results are shown in Table 3. IGGAN achieves considerable improvement compared with other GAN settings on these datasets. In addition, different PDAs can unavoidably produce some similar distribution samples. Thus, IGGAN (NDA + PDA) can provide a greater information gap than IGGAN (PDA + PDA), hence resulting in better results. More generated results of IGGAN can be found in the supplementary materials.

	NDA-GAN	PDA-GAN	Mix-GAN	IGGAN (PDA + PDA)	IGGAN (NDA + PDA)
C10 (U)	12.61	11.87	12.08	11.04	10.68
C10	9.42	8.47	8.78	8.26	8.15
C100 (U)	19.72	16.94	17.28	16.78	16.00
C100	13.90	11.93	12.31	11.73	11.30
STL10 (U)	23.94	22.93	23.34	21.97	21.39
CelebA (U)	22.62	21.36	21.89	21.01	20.63

Table 3. FID Score (lower is better) of IGGAN and other GAN settings on CIFAR-10 (C10), CIFAR-100 (C100), STL10 and CelebA datasets. Here, we select BigGAN as the backbone for all settings and report the best NDA and PDA combination results in IGGAN; (U) means unconditional dataset settings. **The test set for each dataset is used as the reference distribution for FID calculation, as in prior work [47].** For a fair comparison, results are the best run result.

	Jigsaw	Stitching	Mixup	Cutmix
C10U	10.68	10.90	11.01	11.19
C10	8.32	8.22	8.31	8.15
C100U	16.00	16.57	16.70	16.22
C100	11.30	11.75	11.92	11.53
STL10	21.39	22.24	22.74	21.70
CelebA	20.67	21.03	20.85	20.63

Table 4. FID Score (lower is better) on applying different NDA methods to IGGAN on several datasets. For NDA, we apply the NDA methods shown in [35]. The results in red and green represent the best result and second-best results, respectively.

Method	Number of D	NDA	PDA
NDA-GAN	1	Yes	No
PDA-GAN	1	No	Yes
Mix-GAN	1	Yes	Yes
IGGAN (PDA + PDA)	2	No	Yes
IGGAN (NDA + PDA)	2	Yes	Yes

Table 5. Experimental settings for NDA-GAN, PDA-GAN, Mix-GAN, and two types of IGGAN. NDA-GAN consists of NDA in one D , and we apply the jigsaw, which has been shown to obtain the best performance in this setting [35]. PDA-GAN consists of PDA in one D , and the commonly used Diff-Augment [50] is applied in this setting. For Mix-GAN, we apply NDA (jigsaw) and PDA (Diff-Augment) in one D . For IGGAN with NDA in D_1 and PDA in D_2 , we apply Diff-Augment as PDA, and one of the NDA methods as in [35] as NDA. For IGGAN with different PDAs in D_1 and D_2 , based on the conclusion in [52], we apply Translation and Cutout as PDA, respectively.

4.2.2 Results on limited data

We also apply the IGGAN to the limited CIFAR-10 and CIFAR-100 datasets. We use the same limited data settings as in Diff-Augment [50], i.e., only use 10% and 20% of the CIFAR-10 and CIFAR-100 training set to train the model. From the results shown in Table 6, we can find that IG-

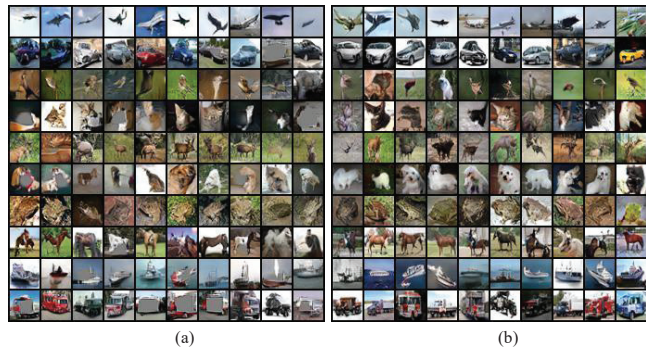


Figure 3. A comparison of the generated images on the 10% CIFAR-10 dataset (each line representing one class): (a) Images generated by BigGAN + Diff-Augment; (b) Images generated by IGGAN. *Best viewed in color.*

GAN achieves state-of-the-art performance without Massive Augmentation (MA) and obtains comparable results with Massive Augmentation (MA) on the 10% and 20% CIFAR-10 and CIFAR-100 datasets. Furthermore, to further show the effectiveness of the information gap between D_1 and D_2 , we also compared IGGAN with other GAN settings, and the results are shown in Table 7. IGGAN achieves considerable improvement compared with these GANs on the limited CIFAR-10/100 datasets.

To further show the superiority of IGGAN, the comparison of generated images on the 10% CIFAR-10 dataset is shown in Figure 3.

4.3. Results using StyleGAN2 Backbone

According to OMASGAN [9], NDA is not useful for all of the GAN backbones. Thus, we first apply the NDA on StyleGAN2 and find that NDA results in worse performance on StyleGAN2, as shown in Table 8. This means that NDA fails to produce the out-of-distribution samples. These generated in-distribution samples not only harm GANs training but also prevent us from utilizing PDA and NDA to build the information gap in StyleGAN2. Therefore, we only apply different PDAs (case 2) in IGGAN to improve the min-max

Method	MA	FID (10% C10)	FID (20% C10)	FID (10% C100)	FID (20% C100)
Non-saturated GAN [10]	No	41.99	18.59	70.50	32.64
LS-GAN [25]	No	41.68	21.60	54.69	27.09
RAHinge GAN [13]	No	48.13	23.90	52.72	28.79
BigGAN [3]	No	48.08	21.86	66.71	32.99
StyleGAN + ADA [15]	No	36.02	23.08	45.87	32.30
BigGAN + GenCo [8]	No	28.08	16.57	40.98	26.15
IGGAN (PDA + PDA)	No	25.17	15.48	38.61	22.57
IGGAN (NDA + PDA)	No	23.56	13.91	37.48	21.64
BigGAN + Diff-Augment [50]	Yes	23.34	14.53	35.39	22.55
DAG-GAN [38]	Yes	21.30	13.10	51.14	26.51
BigGAN + GenCo [8]	Yes	18.10	12.61	25.22	18.44
IGGAN (PDA + PDA)	Yes	19.61	13.27	29.98	20.04
IGGAN (NDA + PDA)	Yes	17.91	12.44	27.02	19.36

Table 6. FID (lower is better) on the limited CIFAR-10 (C10) and CIFAR-100 (C100) datasets (10% and 20%). MA means Massive Augmentation, which has the same meaning as in Genco [8]. Here, we select BigGAN as the backbone for IGGAN. **The test set for each dataset is used as the reference distribution for FID calculation, as in prior work [47].** For a fair comparison, results are averaged over three evaluation runs; all standard deviations are less than 1% relatively.

	NDA-GAN	PDA-GAN	Mix-GAN	IGGAN (PDA + PDA)	IGGAN (NDA + PDA)
10% C10	33.26	23.34	26.93	19.61	17.91
20% C10	16.83	14.53	14.85	13.27	12.44
10% C100	42.29	35.39	33.86	29.98	27.02
20% C100	26.95	22.55	25.99	20.04	19.36

Table 7. FID Score (lower is better) of IGGAN and other GAN settings on limited CIFAR-10 (C10) and CIFAR-100 (C100) settings. For a fair comparison, BigGAN is applied as the backbone for all settings. MA is applied for all settings. The FIDs are averaged over three runs; all standard deviations are less than 1% relatively.

Method	FFHQ (1K)	LSUN-CAT (1K)
StyleGAN2	62.16	182.85
+NDA (Jigsaw)	65.31	190.77
+NDA (Stitch)	67.68	196.47
+NDA (Mixup)	64.22	185.54
+NDA (Cutmix)	68.93	199.06

Table 8. FID score (lower is better) on applying different NDAs in StyleGAN2. Massive Augmentation (MA) is applied in both settings. The FIDs are averaged over five runs; all standard deviations are less than 1% relatively.

game. The commonly used data augmentation methods, i.e., Diff-Augment [50] and ADA [15], are selected as different PDAs in the IGGAN to avoid the leaking of augmentations. The results on 256×256 FFHQ and LSUN-CAT datasets are shown in Table 9. IGGAN achieves state-of-the-art performance compared with other methods. More generated images by IGGAN on FFHQ and LSUN-CAT datasets can be found in the supplementary materials.

Furthermore, Genco [8] also applies multiple discriminators and data augmentations in their method. For a fair comparison with Genco on FFHQ and LSUN datasets, the results of IGGAN without Massive Augmentation (MA) are shown in Table 10. IGGAN achieves a considerable improvement compared with Genco on the FFHQ-1K and LSUN-CAT-1K datasets. To better show the superiority of IGGAN, the generated images on FFHQ-1K are shown in Figure 4. IGGAN produces higher-quality images compared with Genco.

4.4. Ablation Study

To show that the information gap is helpful for GANs training, we first apply IGGAN (NDA + NDA) with two different NDAs (Jigsaw and Stitching) to D_1 and D_2 , respectively. Then, we apply the same PDA (Diff-Augment) to D_1 and D_2 , respectively. These two experimental settings cannot produce the information gap between D_1 and D_2 in the IGGAN. The results of the ablation study are shown in Table 11. Without the information gap between D_1 and D_2 , IGGAN cannot obtain a better FID than the baseline.

Method	FFHQ				LSUN-CAT			
	30K	10K	5K	1K	30K	10K	5K	1K
StyleGAN2 [17]	6.16	14.75	26.60	62.16	10.12	17.93	34.69	182.85
StyleGAN2 + Diff-Augment [50]	5.05	7.86	10.45	25.66	9.68	12.07	16.11	42.26
StyleGAN2 + ADA [15]	5.46	8.13	10.96	21.29	10.50	13.13	16.95	43.25
IGGAN (PDA + PDA)	4.89	7.14	9.47	20.16	9.14	11.20	15.85	30.80

Table 9. FID score (lower is better) on 256×256 FFHQ and LSUN-CAT datasets. Following Diff-Augment [50], we perform experiments on 30K, 10K, 5K and 1K training samples. Massive Augmentation (MA) is applied in all of the methods. For a fair comparison, StyleGAN2 is selected as the backbone for IGGAN. The FIDs are averaged over five runs; all standard deviations are less than 1% relatively.




Figure 4. A comparison of the generated images on the FFHQ-1K dataset. (a) Images generated by Genco; (b) Images generated by IGGAN (PDA + PDA). Massive Augmentation (MA) is not applied in both methods for a fair comparison. *Best viewed in color.*

Method	FFHQ (1K)	LSUN-CAT (1K)
StyleGAN2 [17]	100.13	186.88
Genco [8]	65.31	140.08
IGGAN (PDA + PDA)	24.10	34.47

Table 10. FID score (lower is better) on FFHQ-1K and LSUN-1K datasets. For a fair comparison, StyleGAN2 is selected as the backbone for all settings and Massive Augmentation (MA) is not applied in all methods. We report FID over three runs; all standard deviations are less than 1% relatively.

In contrast, with the information gap between D_1 and D_2 , IGGAN achieves a considerable improvement. This shows the information gap between different data augmentations is helpful for GAN training and can improve performance.

5. Conclusion

In this paper, by analyzing the min-max game of GANs, we unveil a novel insight into an issue that inhibits GAN training, namely the min-max game is unfair in existing GANs during training, leading to sub-optimal convergence. To address this, we propose a new GAN called IGGAN, consisting of one generator and two discriminators, where we apply different data augmentations to each discriminator and utilize the information gap between different data

Method	FID (C10U)	FID (C10)
NDA-BigGAN	12.61	9.42
IGGAN (NDA + NDA)	12.82	9.61
BigGAN + Diff-Augment	11.87	8.47
IGGAN (Diff-Augment in D_1 and D_2)	11.98	8.62
IGGAN (NDA + PDA)	10.68	8.15

Table 11. Experiment results for an ablation study on the effectiveness of the information gap in IGGAN. C10U means the CIFAR-10 dataset with unconditional generation, and C10 means the CIFAR-10 dataset with conditional generation. Here, we select BigGAN as the backbone for all settings and report the results of **IGGAN (NDA + PDA)** with the best NDA and PDA combination.

augmentations to make all three players obtain incomplete information and thus improve the fairness of the min-max game, yielding better convergence. Experiments on CIFAR-10/100, STL10, CelebA with BigGAN and experiments on FFHQ and LSUN-CAT datasets with StyleGAN2 demonstrate the superiority of IGGAN.

Acknowledgement

We are grateful for use of the computing resources from the Northern Ireland High Performance Computing (NI-HPC) service funded by EPSRC (EP/T022175).

References

- [1] Martin Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks. In *ICLR*, 2017. 1
- [2] Qingyan Bai, Ceyuan Yang, Yinghao Xu, Xihui Liu, Yujin Yang, and Yujun Shen. Glead: Improving gans with a generator-leading task. In *CVPR*, 2023. 3
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. In *ICLR*, 2019. 1, 2, 4, 5, 7
- [4] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *CVPR*, 2021. 1
- [5] Tianlong Chen, Yu Cheng, Zhe Gan, Jingjing Liu, and Zhangyang Wang. Data-efficient gan training beyond (just) augmentations: A lottery ticket perspective. In *NeurIPS*, 2021. 1
- [6] Kim-Sau Chung and Jeffrey C Ely. Foundations of dominant-strategy mechanisms. *The Review of Economic Studies*, 74(2):447–476, 2007. 1
- [7] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 2011. 2, 4, 5
- [8] Kaiwen Cui, Jiaying Huang, Zhipeng Luo, Gongjie Zhang, Fangneng Zhan, and Shijian Lu. Genco: Generative co-training for generative adversarial networks with limited data. In *AAAI*, 2022. 7, 8
- [9] Nikolaos Dionelis. Omasgan: Out-of-distribution minimum anomaly score gan for sample generation on the boundary. *arXiv preprint arXiv:2110.15273*, 2021. 6
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014. 1, 2, 3, 7
- [11] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017. 1, 2, 5
- [12] Yifan Jiang, Shiyu Chang, and Zhangyang Wang. Transgan: Two transformers can make one strong gan. In *NeurIPS*, 2021. 1
- [13] Alexia Jolicoeur-Martineau. The relativistic discriminator: a key element missing from standard gan. In *ICLR*, 2019. 7
- [14] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *ICLR*, 2018. 1
- [15] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *NeurIPS*, 2020. 1, 2, 3, 5, 7, 8
- [16] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 1
- [17] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, 2020. 1, 2, 4, 5, 8
- [18] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. The cifar-10 dataset. *online: http://www.cs.toronto.edu/kriz/cifar.html*, 2014. 2, 4, 5
- [19] Nupur Kumari, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Ensembling off-the-shelf models for gan training. In *CVPR*, 2022. 5
- [20] Gayoung Lee, Hyunsu Kim, Junho Kim, Seonghyeon Kim, Jung-Woo Ha, and Yunje Choi. Generator knows what discriminator should learn in unconditional gans. In *ECCV*, 2022. 3
- [21] Ru Li, Chuan Wang, Jue Wang, Guanghui Liu, Heng-Yu Zhang, Bing Zeng, and Shuaicheng Liu. Uphdr-gan: Generative adversarial network for high dynamic range imaging with unpaired data. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(11):7532–7546, 2022. 1
- [22] Kanglin Liu, Wenming Tang, Fei Zhou, and Guoping Qiu. Spectral regularization for combating mode collapse in gans. In *ICCV*, 2019. 1
- [23] Shiguang Liu, Sijia Li, and Haonan Cheng. Towards an end-to-end visual-to-raw-audio generation with gan. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(3):1299–1312, 2021. 1
- [24] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Large-scale celebfaces attributes (celeba) dataset. *Retrieved August*, 15(2018):11, 2018. 2, 4, 5
- [25] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *ICCV*, 2017. 7
- [26] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018. 1
- [27] Takeru Miyato and Masanori Koyama. cgans with projection discriminator. In *ICLR*, 2018. 1
- [28] Tu Nguyen, Trung Le, Hung Vu, and Dinh Phung. Dual discriminator generative adversarial nets. In *NeurIPS*, 2017. 2, 3, 4
- [29] Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. In *NeurIPS*, 2016. 4
- [30] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*, 2016. 1
- [31] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *NeurIPS*, 2016. 2, 5
- [32] Axel Sauer, Kashyap Chitta, Jens Müller, and Andreas Geiger. Projected gans converge faster. In *NeurIPS*, 2021. 2
- [33] Axel Sauer, Katja Schwarz, and Andreas Geiger. Stylegan-xl: Scaling stylegan to large diverse datasets. In *SIGGRAPH 2022*, pages 1–10, 2022. 2
- [34] Hongxia Shen and Tamer Başar. Incentive-based pricing for network games with complete and incomplete information. In *Advances in Dynamic Game Theory*, 2007. 1, 2

- [35] Abhishek Sinha, Kumar Ayush, Jiaming Song, Burak Uzkent, Hongxia Jin, and Stefano Ermon. Negative data augmentation. In *ICLR*, 2021. [2](#), [4](#), [5](#), [6](#)
- [36] Samarth Sinha, Zhengli Zhao, Anirudh Goyal ALIAS PARTH GOYAL, Colin A Raffel, and Augustus Odena. Top-k training of gans: Improving gan performance by throwing away bad samples. In *NeurIPS*, 2020. [1](#), [2](#)
- [37] Ruoyu Sun, Tiantian Fang, and Alexander Schwing. Towards a better global loss landscape of gans. In *NeurIPS*, 2020. [1](#), [3](#)
- [38] Ngoc-Trung Tran, Viet-Hung Tran, Ngoc-Bao Nguyen, Trung-Kien Nguyen, and Ngai-Man Cheung. On data augmentation for gan training. *IEEE Transactions on Image Processing*, 30:1882–1897, 2021. [2](#), [5](#), [7](#)
- [39] Ananya Uppal, Shashank Singh, and Barnabás Póczos. Non-parametric density estimation & convergence rates for gans under besov ipm losses. In *NeurIPS*, 2019. [1](#)
- [40] Jianyuan Wang, Ceyuan Yang, Yinghao Xu, Yujun Shen, Hongdong Li, and Bolei Zhou. Improving gan equilibrium by raising spatial awareness. In *CVPR*, 2022. [3](#)
- [41] Pengyu Wang, Hongqing Zhu, Hui Huang, Han Zhang, and Nan Wang. Tms-gan: A twofold multi-scale generative adversarial network for single image dehazing. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(5):2760–2772, 2021. [1](#)
- [42] Yi-Lun Wu, Hong-Han Shuai, Zhi-Rui Tam, and Hong-Yu Chiu. Gradient normalization for generative adversarial networks. In *ICCV*, 2021. [5](#)
- [43] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, 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. [2](#), [4](#), [5](#)
- [44] Ning Yu, Guilin Liu, Aysegul Dundar, Andrew Tao, Bryan Catanzaro, Larry S Davis, and Mario Fritz. Dual contrastive loss and attention for gans. In *ICCV*, 2021. [1](#)
- [45] Mingkuan Yuan and Yuxin Peng. Bridge-gan: Interpretable representation learning for text-to-image synthesis. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(11):4258–4268, 2019. [1](#)
- [46] Gongjie Zhang, Kaiwen Cui, Tzu-Yi Hung, and Shijian Lu. Defect-gan: High-fidelity defect synthesis for automated defect inspection. In *WACV*, 2021. [1](#)
- [47] Han Zhang, Zizhao Zhang, Augustus Odena, and Honglak Lee. Consistency regularization for generative adversarial networks. In *ICLR*, 2020. [2](#), [5](#), [6](#), [7](#)
- [48] Zhaoyu Zhang, Mengyan Li, Haonian Xie, Jun Yu, Tongliang Liu, and Chang Wen Chen. Twgan: Twin discriminator generative adversarial networks. *IEEE Transactions on Multimedia*, 24:677–688, 2021. [2](#)
- [49] Zhaoyu Zhang, Mengyan Li, and Jun Yu. On the convergence and mode collapse of gan. In *SIGGRAPH Asia 2018 Technical Briefs*, 2018. [2](#)
- [50] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable augmentation for data-efficient gan training. In *NeurIPS*, 2020. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [51] Zhengli Zhao, Sameer Singh, Honglak Lee, Zizhao Zhang, Augustus Odena, and Han Zhang. Improved consistency regularization for gans. In *AAAI*, 2021. [2](#), [5](#)
- [52] Zhengli Zhao, Zizhao Zhang, Ting Chen, Sameer Singh, and Han Zhang. Image augmentations for gan training. *arXiv preprint arXiv:2006.02595*, 2020. [5](#), [6](#)