Deceive D: Adaptive Pseudo Augmentation for GAN Training with Limited Data

Anonymous Author(s)
Affiliation
Address
email

Abstract

Generative adversarial networks (GANs) typically require ample data for training in order to synthesize high-fidelity images. Recent studies have shown that training GANs with limited data remains formidable due to discriminator overfitting, the underlying cause that impedes the generator's convergence. This paper introduces a novel strategy called Adaptive Pseudo Augmentation (APA) to encourage healthy competition between the generator and the discriminator. As an alternative method to existing approaches that rely on standard data augmentations or model regularization, APA alleviates overfitting by employing the generator itself to augment the real data distribution with generated images, which deceives the discriminator adaptively. Extensive experiments demonstrate the effectiveness of APA in improving synthesis quality in the limited-data regime. We provide a theoretical analysis to examine the convergence and rationality of our new training strategy. APA is simple and effective. It can be added seamlessly to powerful contemporary GANs, such as StyleGAN2, with negligible computational cost. Codes will be released.

1 Introduction

While state-of-the-art GANs like StyleGAN2 [19] are constantly pushing forward the fidelity and resolution of synthesized images, they usually require a large amount of training data to fully unleash their power. Training GANs with insufficient data tend to generate poor-quality images, as shown in Figure 1. Collecting sufficient data samples for these GANs is sometimes infeasible, especially in domains where data are sparse and privacy-sensitive. To ease the practical deployment of powerful GANs, it is necessary to devise new strategies for training GANs with limited data while preserving the quality of synthesis.

Recent studies have shown that the overfitting of the discriminator is the critical reason that impedes effective GAN training on limited data [2, 17, 48, 41], rendering severe instability of training dynamics. Specifically, when the discriminator starts to overfit, the distributions of its outputs for real and generated samples gradually diverge from each other [17, 41], and its feedback to the generator becomes less informative. Consequently, the generator converges to an inferior point, compromising the quality of synthesized images. Recent solutions to this problem include the use of standard data augmentations, either conventional or differentiable, to real and generated images [17, 48, 50, 40] or applying an additional model regularization term [41]. Addressing the discriminator overfitting is still an open problem. We are interested in finding an alternative way to the aforementioned approaches.

In this paper, we present a simple yet effective way to regularize a discriminator without introducing any external augmentations or regularization terms. We call our method *Adaptive Pseudo Augmentation* (APA). In contrast to previous standard data augmentations [17, 48, 50, 40], we exploit the generator in a GAN itself to provide the augmentation, a more natural way to regularize the overfitting of the discriminator. Compared to the model regularization, our approach is more adaptive to fit

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

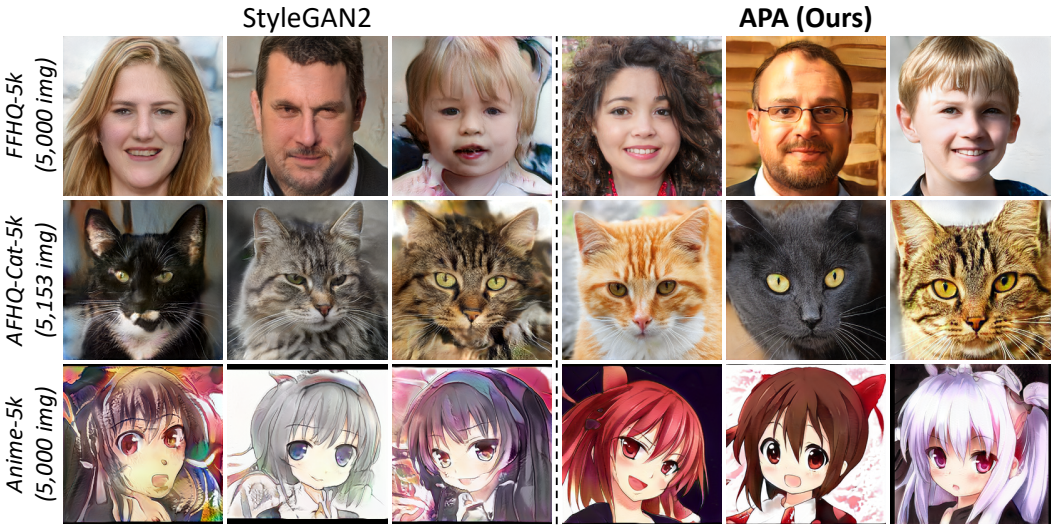

Figure 1: **StyleGAN2** [19] synthesized results (no truncation) deteriorate **given the limited amount of training data** (256×256), *i.e.*, FFHQ [18] (a subset of $5,000$ images, $\sim 7\%$ of full data), AFHQ-Cat [7] ($5,153$ images, which is small by itself), and Danbooru2019 Portraits (Anime) [4] (a subset of $5,000$ images, $\sim 2\%$ of full data). The proposed Adaptive Pseudo Augmentation (APA) effectively ameliorates the degraded performance of StyleGAN2 on limited data.

different settings and training status without manual tuning. Specifically, APA takes the fake/pseudo samples synthesized by the generator and moderately feeds them into the limited real data. Such pseudo data are adaptively presented to the discriminator as "real" instances. The goal of this pseudo augmentation for the real data is not to enlarge the real dataset but to suppress the discriminator's confidence in distinguishing real and fake distributions. The deceit is introduced adaptively, which is moderated by a deception probability according to the degree of overfitting. To quantify overfitting, we study a series of plausible heuristics derived from the discriminator raw output logits.

The main **contribution** of this work is a novel adaptive pseudo augmentation method for training GANs with limited data. This approach deceives the discriminator adaptively and mitigates the problem of discriminator overfitting. The proposed APA can be readily added to existing GAN training with negligible computational cost. We conduct extensive experiments to demonstrate the effectiveness of APA on training state-of-the-art StyleGAN2 [19] with limited data. The results are comparable or even better than other types of solutions [17, 41]. APA is also complementary to existing methods based on standard data augmentations for gaining a further performance boost. Besides, we theoretically connect APA with minimizing the JS divergence [22] between the smoothed data distribution and generated distribution, proving its convergence and rationality. We hope that our approach could extend the breadth and potential of solutions to GAN training with limited data.

2 Related Work

Generative adversarial networks. Generative adversarial networks (GANs) [8, 27, 34] adopt an adversarial training scheme, where a generator keeps refining its capability in synthesizing images to compete with a discriminator (*i.e.*, a binary classifier) until the discriminator fails to classify the generated samples as fakes. GANs are known to suffer from training instability [8, 37, 2, 26]. Various approaches have been proposed to stabilize the training and improve the quality of synthesis by minimizing different f-divergences of the real and fake distributions [29]. The saturated form of vanilla GAN [8] is theoretically proven to minimize the JS divergence [22] between the two distributions. LSGAN [25] and EBGAN [47] correspond to the optimizations of χ^2-divergence [33] and the total variation [3], respectively. On another note, WGAN [3] is designed for minimizing the Wasserstein distance.

State-of-the-art methods, such as PGGAN [16], BigGAN [5], StyleGAN [18], and StyleGAN2 [19], employ large-scale training with contemporary techniques, achieving photorealistic results. These

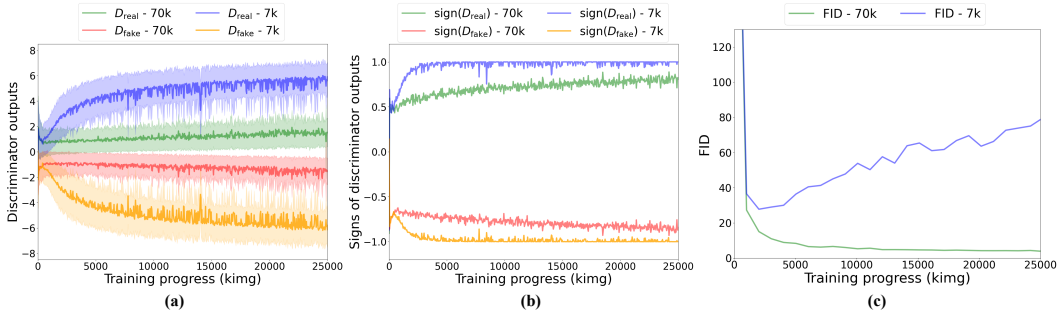

Figure 2: The overfitting of discriminator in GANs when limited training data are available. The three subplots report statistics of training snapshots of two StyleGAN2 [19] models on FFHQ [18] (256×256). "70k" indicates the full dataset, and "7k" means a subset of $7,000$ images (10% data). The "kimg" denotes thousands of real images shown to the discriminator. (a) Discriminator raw output logits. (b) Signs of discriminator outputs. (c) Training convergence measured by FID [10].

methods have been extended to various tasks, including face generation [16, 18, 19], image editing [1, 6, 32], semantic image synthesis [43, 31, 24], image-to-image translation [12, 51, 7, 15, 14, 35], style transfer [23, 11, 21], and GAN inversion [30, 35, 45]. Despite the remarkable success, the performance of GANs relies heavily on the amount of training data.

Training GANs with limited data. The significance and difficulty of training GANs with limited data have been attracting attention from many researchers recently. The issue of data insufficiency tends to cause overfitting in the discriminator [44, 17, 48], which in turn deteriorates the stability of training dynamics in GANs, compromising the quality of generated images.

Many recent studies [46, 49, 40, 50, 48, 17] propose to apply standard data augmentations for GAN training to enrich the diversity of the dataset to mitigate the overfitting of the discriminator. For instance, DiffAugment (DA) [48] adopts the same differentiable augmentation to both real and fake images for the generator and the discriminator without manipulating the target distribution. Adaptive discriminator augmentation (ADA) [17] shares a similar idea with DA, while it further devises an adaptive approach that controls the strength of data augmentations adaptively. In this work, we extend the study of such an adaptive approach.

Another type of solution is model regularization. Previous efforts on regularizing GANs include adding noise to the inputs of the discriminator [2, 39, 13], gradient penalty [26, 9, 36], one-sided label smoothing [37], spectral normalization [28], *etc*. These methods are designed for different goals, either for stabilizing training or preventing mode collapse [37]. Under the limited data setting, a very recent study proposes an LC-regularization term [41] to regulate the discriminator predictions using two exponential moving average variables that track the discriminator outputs throughout training.

Our work explores an alternative solution from a different perspective, which is also complementary to previous approaches based on standard data augmentations.

3 Methodology

In GAN's adversarial training, the goal of the generator G is to deceive the discriminator D and maximize the probability that D makes a wrong judgment. Therefore, G keeps refining its generated samples to better deceive D over time. When the training only accesses a limited amount of data, one would observe that D turns out to be overly confident and hardly makes any mistake, causing its feedback to G to become meaningless.

In Figure 2, we show the training "snapshots" of two StyleGAN2 [19] models on the FFHQ dataset [18]. The settings of the two models differ only by the amount of data available to them for training. As can be observed, both training processes start smoothly, and distributions of discriminator outputs for the real and generated images overlap at the early stage. As the training progresses, the discriminator, which only has access to limited data (7k images, 10% of full data), experiences diverged predictions much more rapidly, and the average sign boundary turns out to be more apparent. This divergence in prediction shows that D becomes increasingly confident in classifying real and

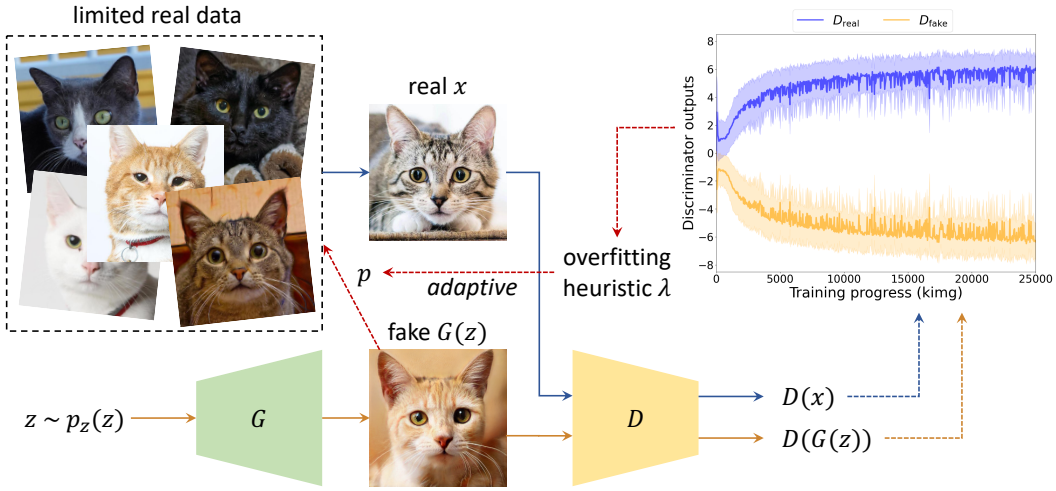

Figure 3: Adaptive pseudo augmentation (APA) for GAN training with limited data. We employ a GAN to augment itself using the generated images to deceive the discriminator adaptively. Specifically, APA feeds the images synthesized by the generator into the limited real data moderately, and these fakes are presented as "real" instances to the discriminator. Such deceits are introduced adaptively using an overfitting heuristic λ defined by the discriminator raw output logits. The augmentation/deception probability p can be adaptively controlled throughout training.

fake samples. At the late stage of training, D can even judge all the input samples correctly with high confidence. Meanwhile, the evaluation FID [10] scores (lower is better) deteriorate, consistent with the divergence of D's predictions. The phenomena above demonstrate how a discriminator gets overfitted quickly with limited data. As can be seen from the FID curves, the overfitting directly influences training dynamics and the convergence of G, leading to poor generation performance.

3.1 Adaptive Pseudo Augmentation

The generator itself naturally possesses the capability to counteract discriminator overfitting. To harness this capability, our method employs a GAN to augment itself using the generated samples (see Figure 3). Specifically, we feed the samples generated by G into the limited real data to form a pseudo-real set. The fake images in this set will be adaptively presented to D, pretending themselves as real data. The goal here is to deceive D with the pseudo-real set and consequently suppressing its confidence in distinguishing real and fake images.

Blindly presenting the fake images as reals to D may weaken the fundamental ability of D in adversarial training. In our approach, the deceit is introduced adaptively to avoid any potential adverse effects. To moderate the deception, we perform pseudo augmentation based on a probability, $p \in [0, 1)$, that quantifies the deception strength. Specifically, the pseudo augmentation will be applied with the probability p or be skipped with the probability $1 - p$.

We note that the overfitting state of D is dynamic throughout the training (see Figure 2). It is intuitive to let the deception probability p be adjusted adaptively based on the degree of overfitting. Ideally, p should be adjusted without manual tuning regardless of data scales and properties. To achieve this goal, inspired by ADA [17], we apply an overfitting heuristic λ that quantifies the degree of D's overfitting. We extend the control scheme of ADA and provide three plausible variants:

$$\lambda_r = \mathbb{E}\left(\operatorname{sign}\left(D_{\text{real}}\right)\right), \quad \lambda_f = -\mathbb{E}\left(\operatorname{sign}\left(D_{\text{fake}}\right)\right), \quad \lambda_{rf} = \frac{\mathbb{E}\left(\operatorname{sign}\left(D_{\text{real}}\right)\right) - \mathbb{E}\left(\operatorname{sign}\left(D_{\text{fake}}\right)\right)}{2}, \tag{1}$$

where D_{real} and D_{fake} are defined as

$$D_{\text{real}} = \operatorname{logit}\left(D\left(x\right)\right), \quad D_{\text{fake}} = \operatorname{logit}\left(D\left(G\left(z\right)\right)\right), \tag{2}$$

where logit denotes the logit function. As shown in Figure 2, the λ_r in Eq. (1) estimates the portion of real images that obtain positive logit predictions by D, and that of generated images is captured by λ_f. Besides, λ_{rf} indicates half of the distance between the signs of the real and fake logits. For all

these heuristics, $\lambda = 0$ represents no overfitting, and $\lambda = 1$ means complete overfitting. We use λ_r in our main experiments and study other variants in the ablation study.

The strategy of using λ to adjust p is as follows. We set a threshold t and initialize p to be zero. If λ signifies too much/little overfitting regarding t, the probability p will be increased/decreased by one fixed step. Using this step size, p can increase from zero to one in 500k images shown to D. We clamp p from below to zero after each step and use $t = 0.6$ in most cases. In this way, the strength of pseudo augmentation can be adaptively controlled based on the degree of overfitting (see Figure 3).

3.2 Theoretical Analysis

Let $p_z(z)$ be the prior on the input noise variable. The mapping from the latent space to the image space is denoted as $G(z)$. For sample x, $D(x)$ represents the estimated probability of x coming from the real data. To examine the rationality of APA, we analyze it in a non-parametric setting [8], where a model is represented with infinite capacity by exploring its convergence in the space of probability density functions. Ideally, the estimated probability distribution p_g defined by G should perfectly model the real data distribution p_{data} without bias if given enough capability and training time.

Since the deception strength p is adaptively adjusted, to facilitate this theoretical analysis, we make a mild assumption that α is the expected strength that approximates the effect of dynamic adjustment during the entire training procedure. Since $p \in [0, 1)$, we have $0 \leq \alpha < p_{\max} < 1$, where $p_{\max}$ is the maximum of deception strength throughout training. Hence, the value function $V(G, D)$ for the minimax two-player game of APA can be reformulated as:

$$\begin{aligned}\min_G \max_D V(G, D) = (1-\alpha)\, \mathbb{E}_{x \sim p_{\text{data}}(x)} \left[\log D(x)\right] + \alpha\, \mathbb{E}_{z \sim p_z(z)} \left[\log D(G(z))\right] \\ + \mathbb{E}_{z \sim p_z(z)} \left[\log\left(1 - D(G(z))\right)\right],\end{aligned} \tag{3}$$

First, let us consider the optimal discriminator [8] for any given generator.

Proposition 1. *If the generator G is fixed, the optimal discriminator D for APA is:*

$$D_G^*(x) = \frac{(1-\alpha)\, p_{\text{data}}(x) + \alpha\, p_g(x)}{(1-\alpha)\, p_{\text{data}}(x) + (1+\alpha)\, p_g(x)} \tag{4}$$

Proof. Applying APA, given any generator G, the training objective of the discriminator D is to maximize the value function $V(G, D)$ in Eq. (3):

$$\begin{aligned}V(G, D) &= (1-\alpha) \int_x p_{\text{data}}(x) \log D(x) dx + \alpha \int_z p_z(z) \log D(g(z)) dz + \int_z p_z(z) \log\left(1 - D(g(z))\right) dz \\ &= \int_x \left[(1-\alpha)\, p_{\text{data}}(x) \log D(x) + \alpha\, p_g(x) \log D(x) + p_g(x) \log\left(1 - D(x)\right)\right] dx \\ &= \int_x \left[\left((1-\alpha)\, p_{\text{data}}(x) + \alpha\, p_g(x)\right) \log D(x) + p_g(x) \log\left(1 - D(x)\right)\right] dx\end{aligned} \tag{5}$$

For any $(m, n) \in \mathbb{R}^2 \setminus \{0, 0\}$, the function $f(y) = m \log(y) + n \log(1 - y)$ achieves its maximum in the range $[0, 1]$ at $\frac{m}{m+n}$. Besides, the discriminator D is defined only inside of $\operatorname{supp}(p_{\text{data}}) \cup \operatorname{supp}(p_g)$, where supp is the set-theoretic support. Therefore, we conclude the proof for Proposition 1.

We have got the optimal discriminator $D_G^*(x)$ in Eq. (4) that maximizes the value function $V(G, D)$ given any fixed generator G. The goal of generator G in the adversarial training is to minimize the value function $V(G, D)$ in Eq. (3) when D achieves the optimum. Since the training objective of D can be interpreted as maximizing the log-likelihood for the conditional probability $P(Y = y|x)$, where Y estimates that x comes from p_{data} (*i.e.*, $y = 1$) or from p_g (*i.e.*, $y = 0$), we reformulate virtual training criterion [8] as:

$$\begin{aligned}C(G) &= (1-\alpha)\, \mathbb{E}_{x \sim p_{\text{data}}} \left[\log D_G^*(x)\right] + \alpha\, \mathbb{E}_{z \sim p_z} \left[\log D_G^*(G(z))\right] + \mathbb{E}_{z \sim p_z} \left[\log\left(1 - D_G^*(G(z))\right)\right] \\ &= (1-\alpha)\, \mathbb{E}_{x \sim p_{\text{data}}} \left[\log D_G^*(x)\right] + \alpha\, \mathbb{E}_{x \sim p_g} \left[\log D_G^*(x)\right] + \mathbb{E}_{x \sim p_g} \left[\log\left(1 - D_G^*(x)\right)\right]\end{aligned} \tag{6}$$

Then, let us consider the global minimum of $C(G)$ trained with the proposed APA.

Proposition 2. *Applying APA, the global minimum of the virtual training criterion $C(G)$ is still achieved if and only if $p_g = p_{\text{data}}$, where $C(G) = -\log 4$.*

Proof. 1) If $p_g = p_{\text{data}}$, we have $D_G^*(x) = \frac{1}{2}$ according to Eq. (4). By inspecting Eq. (6) at $D_G^*(x) = \frac{1}{2}$, we get $C^*(G) = (1-\alpha)\log\frac{1}{2} + \alpha\log\frac{1}{2} + \log\frac{1}{2} = \log\frac{1}{2} + \log\frac{1}{2} = -\log 4$.

2) To verify $C^*(G)$ is the global minimum of $C(G)$, and it can only be achieved when $p_g = p_{\text{data}}$, as in the derivation of Eq. (5), we obtain:

$$C(G) = \int_x \left((1-\alpha)\, p_{\text{data}}(x) + \alpha\, p_g(x)\right) \log D_G^*(x) dx + \int_x p_g(x) \log\left(1 - D_G^*(x)\right) dx \tag{7}$$

Observe that

$$\begin{aligned} -\log 4 &= (1-\alpha)\, \mathbb{E}_{x\sim p_{\text{data}}}[-\log 2] + \alpha\, \mathbb{E}_{x\sim p_g}[-\log 2] + \mathbb{E}_{x\sim p_g}[-\log 2] \\ &= -\int_x \left((1-\alpha)\, p_{\text{data}}(x) + \alpha\, p_g(x)\right) \log 2\, dx - \int_x p_g(x) \log 2\, dx \end{aligned} \tag{8}$$

Subtracting Eq. (8) from Eq. (7),

$$C(G) = -\log 4 + \int_x \left((1-\alpha)\, p_{\text{data}}(x) + \alpha\, p_g(x)\right) \log 2 \cdot D_G^*(x) dx + \int_x p_g(x) \log 2 \cdot \left(1 - D_G^*(x)\right) dx \tag{9}$$

By substituting Eq. (4) into Eq. (9), we achieve:

$$\begin{aligned} C(G) = -\log 4 &+ \text{KLD}\left(\left((1-\alpha)\, p_{\text{data}} + \alpha\, p_g\right) \middle\| \frac{(1-\alpha)\, p_{\text{data}} + (1+\alpha)\, p_g}{2}\right) \\ &+ \text{KLD}\left(p_g \middle\| \frac{(1-\alpha)\, p_{\text{data}} + (1+\alpha)\, p_g}{2}\right), \end{aligned} \tag{10}$$

where KLD is the Kullback-Leibler (KL) divergence. Moreover, Eq. (10) further implies that the generation process of G by APA can be regarded as minimizing the Jensen-Shannon (JS) divergence between the smoothed data distribution and the generated distribution:

$$C(G) = -\log 4 + 2 \cdot \text{JSD}\left(\left((1-\alpha)\, p_{\text{data}} + \alpha\, p_g\right) \| p_g\right). \tag{11}$$

For the two distributions P and Q, their JS divergence $\text{JSD}(P\|Q) \geq 0$ and $\text{JSD}(P\|Q) = 0$ if and only if $P = Q$. Therefore, for $0 \leq \alpha < p_{\max} < 1$, we obtain that $C^*(G) = -\log 4$ is the global minimum of $C(G)$, and the only solution is $(1-\alpha)\, p_{\text{data}} + \alpha\, p_g = p_g$, *i.e.*, $p_g = p_{\text{data}}$. Q.E.D.

Given the proof in [8], if G and D have enough capacity to reach their optimum, Proposition 2 indicates that the generated distribution p_g can ideally converge to the real data distribution p_{data}. So far, we have proved the convergence of G trained with our proposed APA, which can perfectly model the real data distribution given sufficient capability and training time. Besides, the JS divergence term between the smoothed data distribution and the generated distribution in Eq. (11) implies that the judgment of D may be moderated to alleviate overfitting. These conclusions explain the rationality of the proposed APA for training GANs with limited data.

4 Experiments

Datasets. We use a total of four datasets: Flickr-Faces-HQ (FFHQ) [18] with $70,000$ human face images, AFHQ-Cat [7] with $5,153$ cat faces, Caltech-UCSD Birds-200-2011 (CUB) [42] with $11,788$ images of birds, and Danbooru2019 Portraits (Anime) [4] with $302,652$ anime portraits. We exploit some of their artificially limited subsets under different settings. All the images are resized to a moderate resolution of 256×256 using a high-quality Lanczos filter [20] to reduce the energy consumption for large-scale GAN training while preserving image quality. Additional dataset details, including the data source and license information, can be found in our *supplementary material*.

Evaluation metrics. We follow the standard evaluation protocol [5, 41] for the quantitative evaluation. Specifically, we use the Fréchet Inception Distance (FID, lower is better) [10], which quantifies the distance between distributions for the real and generated images. FID evaluates the realness of synthesized images. Following [10, 17], we calculate FID for the models trained with limited datasets using 50k generated images and all the real images in the original datasets. We also apply the Inception Score (IS, higher is better) [37]. IS considers the clarity and diversity of generated images.

Implementation details. We choose the state-of-the-art StyleGAN2 [19] as the backbone to verify the effectiveness of APA on limited data. We use the default setups of APA provided in Section 3.1 unless specified otherwise. For a fair and controllable comparison, we reimplement all baselines and run the experiments from scratch using official codes. All the models are trained on 8 NVIDIA Tesla V100 GPUs. Please refer to the *supplementary material* for more implementation details.

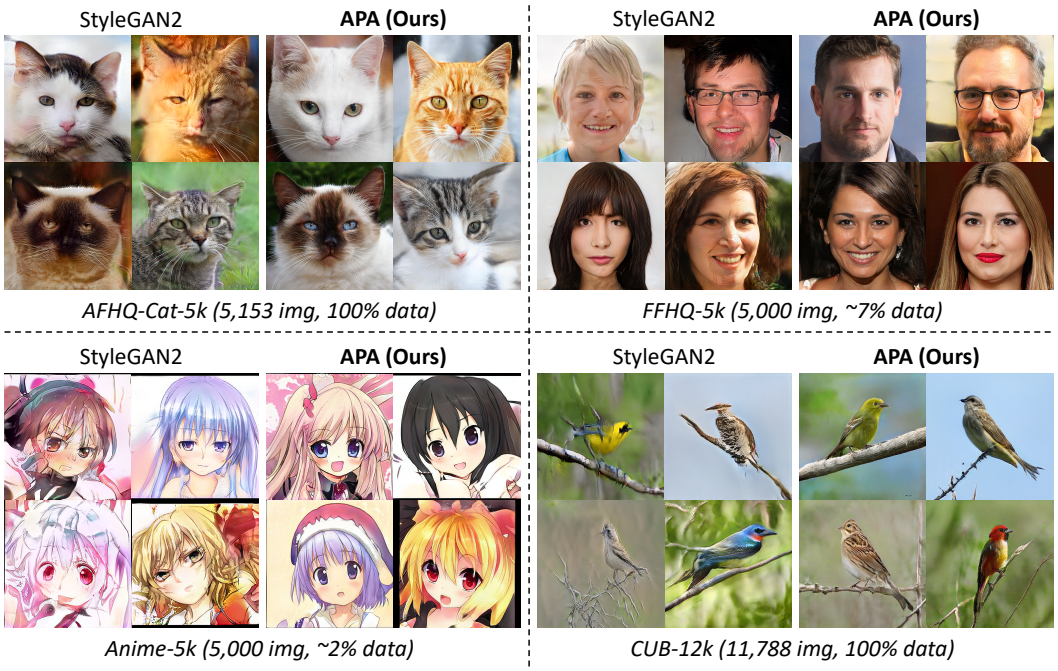

Figure 4: The proposed APA improves StyleGAN2 [19] synthesized results (256×256, no truncation) on **various datasets** with limited data amounts. We randomly select subsets to confine the size of large datasets (*i.e.*, FFHQ-5k [18] and Anime-5k [4]) and directly use small datasets (*i.e.*, AFHQ-Cat-5k [7] and CUB-12k [42]) whose data amount is already insufficient for StyleGAN2.

Table 1: The FID (lower is better) and IS (higher is better) scores (256×256) of our method compared to state-of-the-art StyleGAN2 on **various datasets** with limited data amounts.

Method	AFHQ-Cat-5k		FFHQ-5k		Anime-5k		CUB-12k	
	FID ↓	IS ↑	FID ↓	IS ↑	FID ↓	IS ↑	FID ↓	IS ↑
StyleGAN2 [19]	7.737	1.825	37.830	4.018	23.778	2.289	23.437	5.812
APA (Ours)	**4.876**	**2.156**	**13.249**	**4.487**	**13.089**	**2.330**	**12.889**	**5.869**

4.1 The Effectiveness of APA

Effectiveness on various datasets. The comparative results of StyleGAN2 on various datasets with limited data amounts are shown in Figure 4. The quality of images synthesized by StyleGAN2 deteriorates under limited data. Ripple artifacts appear on the cat faces and human faces, and the facial features of the anime faces are misplaced. On the bird dataset with heavy background clutter, the generated images are completely distorted albeit trained with more data. The proposed APA significantly ameliorates image quality on all these datasets, producing much more photorealistic results. The quantitative evaluation results are reported in Table 1. Applying APA contributes to a performance boost of FID and IS in all cases, suggesting that the synthesized images are with higher quality and diversity on different datasets.

Effectiveness given different data amounts. The comparative results on subsets of FFHQ [18], with varying amounts of data, are shown in Figure 5. The corresponding quantitative test results are presented in Table 2. APA improves the image quality and metric performance in all cases. Notably, the quality of synthesized images by APA on 5k/7k data is visually close to StyleGAN2 results on the full dataset while with an order of magnitude fewer training samples. As for the quantitative results, the IS score of APA on 1k data is better than that of StyleGAN2 on 5k data, and both metrics of APA on 5k data outperform StyleGAN2 results on 7k data. APA can even improve StyleGAN2 performance on the full dataset, further indicating its potential.

Overfitting and convergence analysis. As shown in Figure 6, the divergence of StyleGAN2 discriminator predictions can be effectively restricted on FFHQ-7k (10% of full data) by applying the proposed APA. The curves of APA on FFHQ-7k become closer to that of StyleGAN2 on FFHQ-70k,

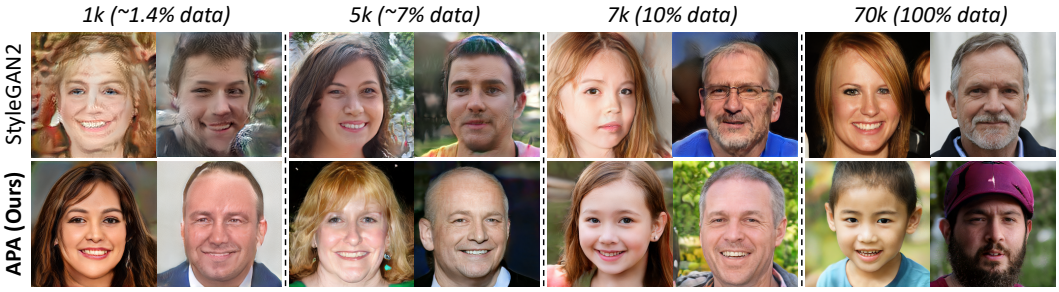

Figure 5: The effectiveness of APA to improve StyleGAN2 [19] synthesized results (256×256, no truncation) on the subsets of FFHQ [18] with **different data amounts**.

Table 2: The FID (lower is better) and IS (higher is better) scores (256×256) of our method on StyleGAN2 trained using the subsets of FFHQ [18] with **different data amounts**.

	1k ($\sim 1.4\%$)		5k ($\sim 7\%$)		7k (10%)		70k (100%)	
Method	FID ↓	IS ↑	FID ↓	IS ↑	FID ↓	IS ↑	FID ↓	IS ↑
StyleGAN2 [19]	86.407	2.806	37.830	4.018	27.738	4.264	3.862	5.243
APA (Ours)	**45.192**	**4.130**	**13.249**	**4.487**	**10.800**	**4.860**	**3.678**	**5.336**

suggesting the effectiveness of APA in curbing the overfitting of the discriminator. Besides, APA improves the training convergence of StyleGAN2 on limited data, shown by the FID curves.

4.2 Comparison with Other Solutions

Performance and compatibility. We compare the proposed APA with representative approaches designed for the limited-data regime, including ADA [17] and LC-regularization (LC-Reg) [41], which perform standard data augmentations and model regularization, respectively. The results are reported in Table 3. As a single method, APA outperforms previous solutions in most cases, effectively improving the StyleGAN2 baseline on both the limited and full datasets. Although ADA [17] achieves slightly better results than our method on FFHQ-5k, it yields a worse FID score with StyleGAN2 on the full dataset. Applying LC-Reg [41] needs careful manual tuning, and its own effect on limited data is not apparent compared to other methods.

It is noteworthy that APA is also complementary to existing methods based on standard data augmentations, *e.g.*, ADA [17]. As can be observed in Table 3, APA can further boost the performance of StyleGAN2 with ADA [17] given limited training data, suggesting the compatibility of our approach with standard data augmentations.

Training cost. We compare the computational cost of APA against ADA [17], using the same basic official codebase. There is no parameter or memory increment for both methods. As for the time

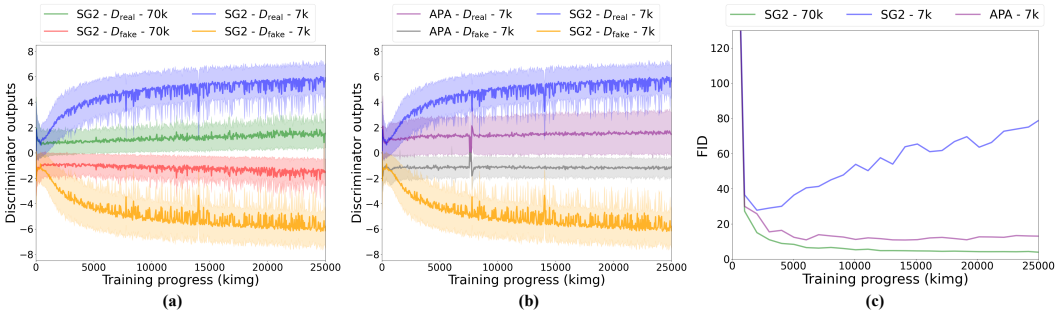

Figure 6: The **overfitting and convergence status** of APA compared to StyleGAN2 (SG2) on FFHQ [18] (256×256). (a) The discriminator raw output logits of StyleGAN2 on the full (70k) or limited (7k) datasets. (b) The discriminator raw output logits of StyleGAN2 and APA on the limited (7k) dataset. (c) The training convergence shown by FID.

Table 3: The FID (lower is better) and IS (higher is better) scores (256×256) of our method **compared to other state-of-the-art solutions** designed for GAN training with limited data on StyleGAN2. The bold number indicates the best value, and the underline marks the second best.

Method	AFHQ-Cat-5k		FFHQ-5k		FFHQ-70k (full)	
	FID ↓	IS ↑	FID ↓	IS ↑	FID ↓	IS ↑
StyleGAN2 [19]	7.737	1.825	37.830	4.018	<u>3.862</u>	5.243
ADA [17]	6.053	2.119	<u>11.409</u>	<u>4.721</u>	4.018	<u>5.329</u>
LC-Reg [41]	6.699	1.943	35.148	3.926	3.933	5.312
APA (Ours)	<u>4.876</u>	<u>2.156</u>	13.249	4.487	**3.678**	**5.336**
ADA + APA (Ours)	**4.377**	**2.169**	**8.379**	**4.849**	-	-

Table 4: **Ablation studies** of different variants of APA on FFHQ-5k [18] (256×256). We study the two key elements of APA, *i.e.*, the overfitting heuristic λ and deception probability p. The "main" denotes the main version used in our previous experiments (*i.e.*, $\lambda = \lambda_r$, and p is adaptively adjusted).

Metric	StyleGAN2	main	$\lambda = \lambda_f$	$\lambda = \lambda_{rf}$	$p = 0.5$ (fixed)
FID ↓	37.830	<u>13.249</u>	13.470	**12.679**	14.632
IS ↑	4.018	**4.487**	<u>4.420</u>	4.412	4.403

consumption, we test the training cost on 8 NVIDIA Tesla V100 GPUs. On FFHQ-5k (256×256), the average training time of the StyleGAN2 [19] baseline is (4.740 ± 0.100) sec/kimg (*i.e.*, seconds per thousand of images shown to the discriminator). The cost of our method is negligible, slightly increasing this value to (4.789 ± 0.078) sec/kimg. As a reference, the value for ADA [17] is (5.327 ± 0.116) sec/kimg, spending additional time for applying external augmentations.

4.3 Ablation Studies

We further explore variants of APA through ablation studies. We study the two key elements of APA, *i.e.*, the overfitting heuristic λ and the deception strength p. The version used in our main experiments is denoted as the "main" version, where $\lambda = \lambda_r$, and p is adjusted adaptively. As reported in Table 4, when using other variants of λ we suggested in Eq. (1) (*i.e.*, $\lambda = \lambda_f$ and $\lambda = \lambda_{rf}$), the models achieve comparable performance as the main version. The FID score becomes even better for $\lambda = \lambda_{rf}$, indicating the flexibility of our provided heuristics to extend and modify APA. More interestingly, even a fixed moderate deception probability (*e.g.*, $p = 0.5$) can still work much better than original StyleGAN2 on limited data, albeit slightly inferior to the adaptively adjusted p. This implies the importance of the pseudo augmentation, and the adaptive control scheme can further boost performance without manual tuning.

5 Discussion

We have shown the effectiveness of the proposed adaptive pseudo augmentation (APA) for state-of-the-art GAN training with limited data both theoretically and empirically. With negligible computational cost, APA achieves comparable or even better performance than other types of solutions on various datasets. APA is also complementary to existing methods based on standard data augmentations.

Limitations. Despite promising results, the quality of synthesized images by APA on the datasets with extremely limited data amount (*e.g.*, hundreds of images) can still be improved. Besides, on certain datasets such as FFHQ-5k, applying APA solely may be slightly inferior to approaches based on standard data augmentations. Since we do not apply any external augmentations, these two limitations are both due to the insufficiency of the dataset's intrinsic diversity. These limitations may be approached in the future in two ways: 1) Incorporating better standard data augmentations to APA. 2) Exploring the issue of data insufficiency from the generator aspect, *e.g.*, using a multi-modal generator [38] to enhance diversity.

Broader impact. On the one hand, the effectiveness of APA with negligible computational cost will benefit the practical deployment of GANs, especially in the limited-data regime. APA may also extend the breadth and potential of solutions to training GANs with limited data and benefit downstream tasks, such as conditional synthesis. On the other hand, APA may also bring potential concerns on its capability to ease the higher-quality fake media synthesis using only limited data, as technology is usually a double-edged sword. However, we believe that these concerns can be resolved by developing better media forensics methods and datasets as countermeasures.

References

[1] faceswap-GAN. https://github.com/shaoanlu/faceswap-GAN/. Accessed: 2019-08-16.

[2] Martin Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks. In *ICLR*, 2017.

[3] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *ICML*, 2017.

[4] Gwern Branwen, Anonymous, and Danbooru Community. Danbooru2019 Portraits: A large-scale anime head illustration dataset. https://www.gwern.net/Crops#danbooru2019-portraits. Accessed: 2021-04-10.

[5] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *ICLR*, 2018.

[6] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation. In *CVPR*, 2018.

[7] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. StarGAN v2: Diverse image synthesis for multiple domains. In *CVPR*, 2020.

[8] 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.

[9] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans. In *NeurIPS*, 2017.

[10] 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.

[11] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, 2018.

[12] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017.

[13] Simon Jenni and Paolo Favaro. On stabilizing generative adversarial training with noise. In *CVPR*, 2019.

[14] Liming Jiang, Bo Dai, Wayne Wu, and Chen Change Loy. Focal frequency loss for image reconstruction and synthesis. *arXiv preprint*, arXiv:2012.12821, 2020.

[15] Liming Jiang, Changxu Zhang, Mingyang Huang, Chunxiao Liu, Jianping Shi, and Chen Change Loy. TSIT: A simple and versatile framework for image-to-image translation. In *ECCV*, 2020.

[16] 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.

[17] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *NeurIPS*, 2020.

[18] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019.

[19] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020.

[20] Cornelius Lanczos. *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*. United States Governm. Press Office Los Angeles, CA, 1950.

[21] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *ECCV*, 2018.

[22] Jianhua Lin. Divergence measures based on the Shannon entropy. *TIT*, 37:145–151, 1991.

[23] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NeurIPS*, 2017.

[24] Xihui Liu, Guojun Yin, Jing Shao, Xiaogang Wang, and Hongsheng Li. Learning to predict layout-to-image conditional convolutions for semantic image synthesis. In *NeurIPS*, 2019.

[25] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *ICCV*, 2017.

[26] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *ICML*, 2018.

[27] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint*, arXiv:1411.1784, 2014.

[28] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018.

[29] Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-GAN: Training generative neural samplers using variational divergence minimization. *arXiv preprint arXiv:1606.00709*, 2016.

[30] Xingang Pan, Xiaohang Zhan, Bo Dai, Dahua Lin, Chen Change Loy, and Ping Luo. Exploiting deep generative prior for versatile image restoration and manipulation. In *ECCV*, 2020.

[31] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, 2019.

[32] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. StyleCLIP: Text-driven manipulation of stylegan imagery. *arXiv preprint*, arXiv:2103.17249, 2021.

[33] Karl Pearson. X. on the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 50:157–175, 1900.

[34] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint*, arXiv:1511.06434, 2015.

[35] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *CVPR*, 2021.

[36] Kevin Roth, Aurelien Lucchi, Sebastian Nowozin, and Thomas Hofmann. Stabilizing training of generative adversarial networks through regularization. In *NeurIPS*, 2017.

[37] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training GANs. In *NeurIPS*, 2016.

[38] Omry Sendik, Dani Lischinski, and Daniel Cohen-Or. Unsupervised multi-modal styled content generation. *arXiv preprint*, arXiv:2001.03640, 2020.

[39] Casper Kaae Sønderby, Jose Caballero, Lucas Theis, Wenzhe Shi, and Ferenc Huszár. Amortised map inference for image super-resolution. *arXiv preprint*, arXiv:1610.04490, 2016.

[40] Ngoc-Trung Tran, Viet-Hung Tran, Ngoc-Bao Nguyen, Trung-Kien Nguyen, and Ngai-Man Cheung. On data augmentation for gan training. *TIP*, 30:1882–1897, 2021.

[41] Hung-Yu Tseng, Lu Jiang, Ce Liu, Ming-Hsuan Yang, and Weilong Yang. Regularizing generative adversarial networks under limited data. In *CVPR*, 2021.

[42] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The Caltech-UCSD Birds-200-2011 Dataset. 2011.

[43] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional GANs. In *CVPR*, 2018.

[44] Ryan Webster, Julien Rabin, Loic Simon, and Frédéric Jurie. Detecting overfitting of deep generative networks via latent recovery. In *CVPR*, 2019.

[45] Weihao Xia, Yulun Zhang, Yujiu Yang, Jing-Hao Xue, Bolei Zhou, and Ming-Hsuan Yang. Gan inversion: A survey. *arXiv preprint*, arXiv: 2101.05278, 2021.

[46] Han Zhang, Zizhao Zhang, Augustus Odena, and Honglak Lee. Consistency regularization for generative adversarial networks. In *ICLR*, 2020.

[47] Junbo Zhao, Michael Mathieu, and Yann LeCun. Energy-based generative adversarial network. *arXiv preprint*, arXiv:1609.03126, 2016.

[48] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable augmentation for data-efficient gan training. In *NeurIPS*, 2020.

[49] Zhengli Zhao, Sameer Singh, Honglak Lee, Zizhao Zhang, Augustus Odena, and Han Zhang. Improved consistency regularization for gans. *arXiv preprint*, arXiv:2002.04724, 2020.

[50] Zhengli Zhao, Zizhao Zhang, Ting Chen, Sameer Singh, and Han Zhang. Image augmentations for gan training. *arXiv preprint*, arXiv:2006.02595, 2020.

[51] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] Please refer to Section 4 and Section 3.2.
   (b) Did you describe the limitations of your work? [Yes] Please refer to Section 5.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] Please refer to Section 5.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] Please refer to Line 139-145.
   (b) Did you include complete proofs of all theoretical results? [Yes] Please refer to Section 3.2.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Please refer to **Datasets**, **Evaluation metrics**, and **Implementation details** in Section 4. Additional details can be found in the supplementary material. We will release the code upon the publication of this paper.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Please refer to the dataset and implementation details in Section 4 and the supplementary material.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] Error bars were reported for the computational cost test (see Section 4.2 **Training cost**). For the performance metrics, error bars were not reported because it would be too computationally expensive for the large-scale GAN training. Our quantitative results were based on over 50k generated images and all the real images in the original dataset to reduce the error.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Please refer to the implementation details in Section 4 and the supplementary material.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] Please refer to the dataset and implementation details in Section 4 and the supplementary material.
   (b) Did you mention the license of the assets? [Yes] Please refer to the dataset details in the supplementary material.
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A] We will release the code upon the publication of this paper.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No] We use public datasets provided by the original authors. The consent information is not applicable for the animal and anime datasets. We did not find the consent information of the human face dataset from the original authors.

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No] This information was not provided by the original authors. We carefully checked the images presented in this paper, which contain no personally identifiable information or offensive content.

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]