

---

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

---

Liming Jiang<sup>1</sup>   Bo Dai<sup>1</sup>   Wayne Wu<sup>2</sup>   Chen Change Loy<sup>1\*</sup>

<sup>1</sup>S-Lab, Nanyang Technological University   <sup>2</sup>SenseTime Research  
{liming002, bo.dai, ccloy}@ntu.edu.sg   wuwenyan@sensetime.com

## 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 low-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. Code: <https://github.com/EndlessSora/DeceiveD>.

## 1 Introduction

While state-of-the-art GANs like StyleGAN2 [21] 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 [3, 19, 55, 46], 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 [19, 46], 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 [19, 55, 57, 45] or applying an additional model regularization term [46]. 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 [19, 55, 57, 45], we exploit the

---

\*Corresponding author.

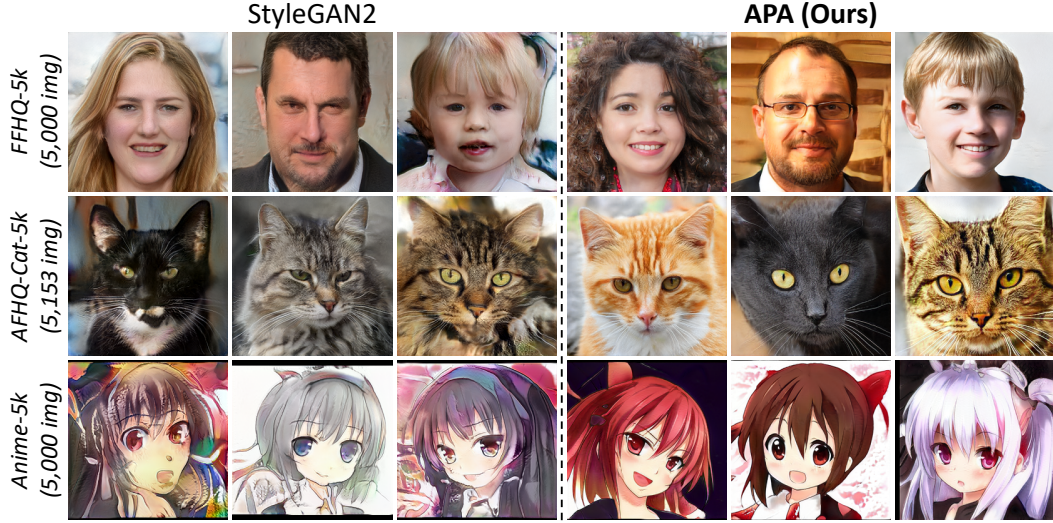


Figure 1: **StyleGAN2** [21] synthesized results (no truncation) deteriorate **given the limited amount of training data** ( $256 \times 256$ ), *i.e.*, FFHQ [20] (a subset of 5,000 images,  $\sim 7\%$  of full data), AFHQ-Cat [9] (5,153 images, which is small by itself), and Danbooru2019 Portraits (Anime) [5] (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.

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 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 for state-of-the-art GAN training with limited data. The results are comparable or even better than other types of solutions [19, 46]. 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 [27] 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) [10, 32, 39, 52] 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 [10, 42, 3, 31]. 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 [34]. The saturated form of vanilla GAN [10] is theoretically proven to minimize the JS divergence [27] between the two distributions. LSGAN [30] and EBGAN [54] correspond to the optimizations of  $\chi^2$ -divergence [38] and the total variation [4], respectively. On another note, WGAN [4] is designed for minimizing the Wasserstein distance.




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 [21] models on FFHQ [20] ( $256 \times 256$ ). “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 [12].

State-of-the-art methods, such as PGGAN [18], BigGAN [6], StyleGAN [20], and StyleGAN2 [21], employ large-scale training with contemporary techniques, achieving photorealistic results. These methods have been extended to various tasks, including face generation [18, 20, 21], image editing [1, 8, 37], semantic image synthesis [48, 36, 29], image-to-image translation [14, 58, 9, 17, 16, 40], style transfer [28, 13, 26], and GAN inversion [35, 40, 51]. 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 [49, 19, 55], which in turn deteriorates the stability of training dynamics in GANs, compromising the quality of generated images.

Many recent studies [53, 56, 45, 57, 55, 19] 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) [55] 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) [19] 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 [3, 44, 15], gradient penalty [31, 11, 41], one-sided label smoothing [42], spectral normalization [33], label noise [7], *etc.* These methods are designed for stabilizing training or preventing mode collapse [42]. The essence of their goals could be considered similar to our method since training GANs in the low-data regime exhibits similar behaviors as previously observed in early GANs with sufficient data. Additional discussions on these techniques are provided in the *Appendix*. Under the limited data setting, a very recent study proposes an LC-regularization term [46] 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 [21] models on the FFHQ dataset [20]. The settings of the two models differ only by the amount of data available to them for

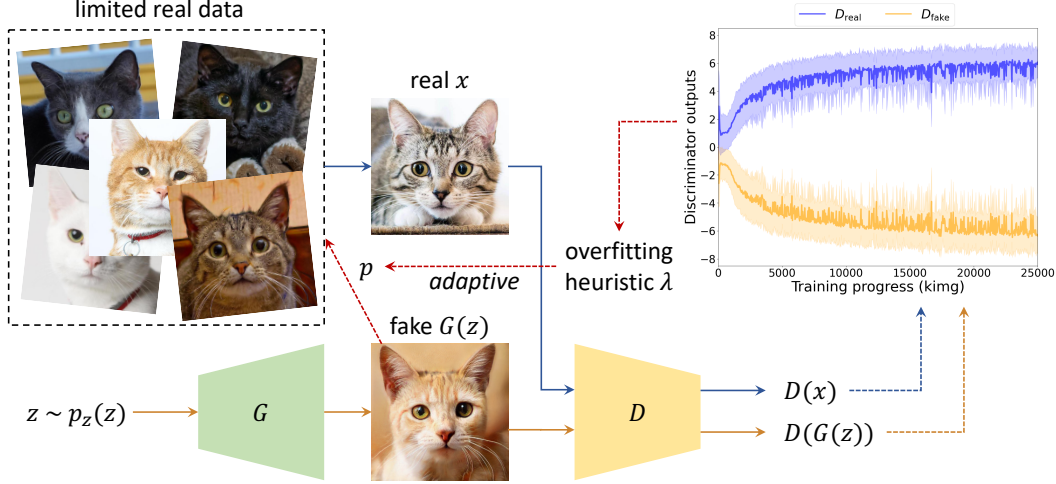


Figure 3: Adaptive pseudo augmentation (APA) for GAN training with limited data. We employ a GAN to augment itself using the generated images to deceive the discriminator adaptively. Specifically, APA feeds the images synthesized by the generator into the limited real data moderately, and these fakes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

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 fake samples. At the late stage of training,  $D$  can even judge all the input samples correctly with high confidence. Meanwhile, the evaluation FID [12] 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 [19], we apply an overfitting heuristic  $\lambda$  that quantifies the degree of  $D$ ’s overfitting. We extend the control scheme of ADA and provide three plausible variants:

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

where  $D_{\text{real}}$  and  $D_{\text{fake}}$  are defined as

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



where logit denotes the logit function. As shown in Figure 2, the  $\lambda_r$  in Eq. (1) estimates the portion of real images that obtain positive logit predictions by  $D$ , and that of generated images is captured by  $\lambda_f$ . Besides,  $\lambda_{rf}$  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  $\lambda_r$  in our main experiments and study other variants in the ablation study.

The strategy of using  $\lambda$  to adjust  $p$  is as follows. We set a threshold  $t$  (in most cases of our experiments,  $t = 0.6$ ) and initialize  $p$  to be zero. If  $\lambda$  signifies too much/little overfitting regarding  $t$  (i.e., larger/smaller than  $t$ ), 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 adjust  $p$  once every four iterations and clamp  $p$  from below to zero after each adjustment. 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 [10], 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  $p_{\text{data}}$  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  $\alpha$  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:

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

First, let us consider the optimal discriminator [10] for any given generator.

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

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

*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):

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

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  $\text{supp}(p_{\text{data}}) \cup \text{supp}(p_g)$ , where  $\text{supp}$  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  $p_{\text{data}}$  (i.e.,  $y = 1$ ) or from  $p_g$  (i.e.,  $y = 0$ ), we reformulate virtual training criterion [10] as:

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

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:

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

Observe that

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

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

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

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

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

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:

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

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  $0 \leq \alpha < p_{\text{max}} < 1$ , 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 [10], 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  $p_{\text{data}}$ . 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 four datasets in our main experiments: Flickr-Faces-HQ (FFHQ) [20] with 70,000 human face images, AFHQ-Cat [9] with 5,153 cat faces, Caltech-UCSD Birds-200-2011 (CUB) [47] with 11,788 images of birds, and Danbooru2019 Portraits (Anime) [5] 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  $256 \times 256$  using a high-quality Lanczos filter [25] 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 the *Appendix*.

**Evaluation metrics.** We follow the standard evaluation protocol [6, 46] for the quantitative evaluation. Specifically, we use the Fréchet Inception Distance (FID, lower is better) [12], which quantifies the distance between distributions for the real and generated images. FID evaluates the realness of synthesized images. Following [12, 19], we calculate FID for the models trained with limited

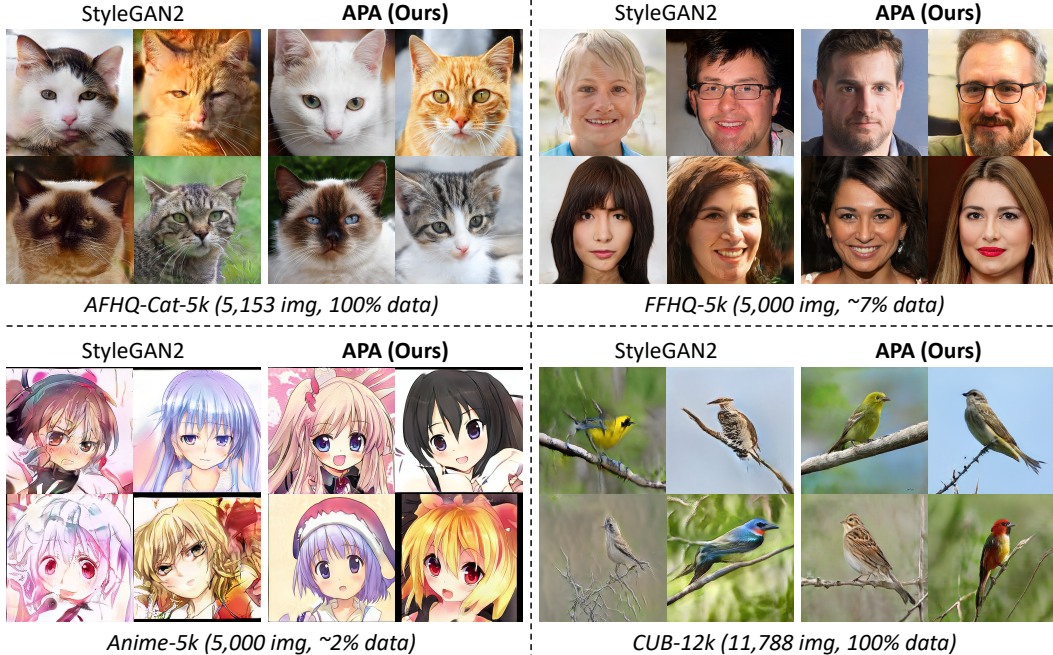


Figure 4: The proposed APA improves StyleGAN2 [21] synthesized results ( $256 \times 256$ , no truncation) on **various datasets** with limited data amounts. We randomly select subsets to confine the size of large datasets (*i.e.*, FFHQ-5k [20] and Anime-5k [5]) and directly use small datasets (*i.e.*, AFHQ-Cat-5k [9] and CUB-12k [47]) whose data amount is already insufficient for StyleGAN2.

Table 1: The FID (lower is better) and IS (higher is better) scores ( $256 \times 256$ ) 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 [21]	7.737	1.825	37.830	4.018	23.778	2.289	23.437	5.812
APA (Ours)	<b>4.876</b>	<b>2.156</b>	<b>13.249</b>	<b>4.487</b>	<b>13.089</b>	<b>2.330</b>	<b>12.889</b>	<b>5.869</b>

datasets using 50k generated images and all the real images in the original datasets. We also apply the Inception Score (IS, higher is better) [42]. IS considers the clarity and diversity of generated images.

**Implementation details.** We choose the state-of-the-art StyleGAN2 [21] 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 code. All the models are trained on 8 NVIDIA Tesla V100 GPUs. Please refer to the *Appendix* for more implementation details and additional benchmark results (*e.g.*, performance on BigGAN [6]).

#### 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 [20], with varying amounts of data, are shown in Figure 5. The corresponding quantitative test results are




Figure 5: The effectiveness of APA to improve StyleGAN2 [21] synthesized results ( $256 \times 256$ , no truncation) on the subsets of FFHQ [20] with **different data amounts**.

Table 2: The FID (lower is better) and IS (higher is better) scores ( $256 \times 256$ ) of our method on StyleGAN2 trained using the subsets of FFHQ [20] with **different data amounts**.

Method	1k ( $\sim 1.4\%$ )		5k ( $\sim 7\%$ )		7k (10%)		70k (100%)	
	FID $\downarrow$	IS $\uparrow$	FID $\downarrow$	IS $\uparrow$	FID $\downarrow$	IS $\uparrow$	FID $\downarrow$	IS $\uparrow$
StyleGAN2 [21]	86.407	2.806	37.830	4.018	27.738	4.264	3.862	5.243
APA (Ours)	<b>45.192</b>	<b>4.130</b>	<b>13.249</b>	<b>4.487</b>	<b>10.800</b>	<b>4.860</b>	<b>3.678</b>	<b>5.336</b>

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, 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. More overfitting and convergence analysis can be found in the *Appendix*.

## 4.2 Comparison with Other Solutions for GAN Training with Limited Data

**Performance and compatibility.** We compare the proposed APA with representative approaches designed for the low-data regime, including ADA [19] and LC-regularization (LC-Reg) [46], 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

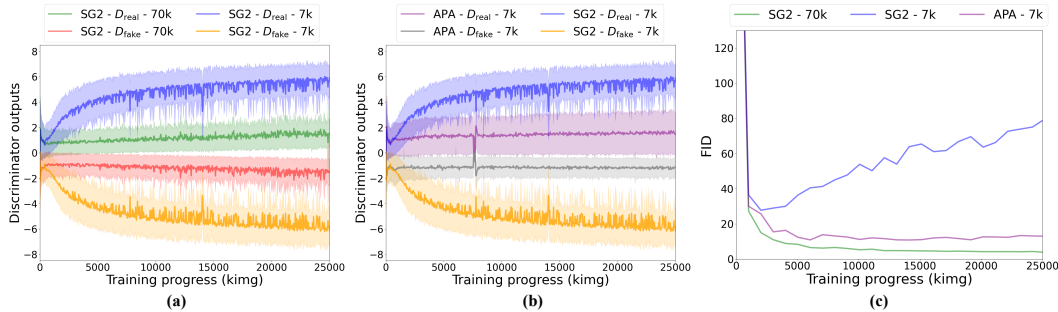


Figure 6: The **overfitting and convergence status** of APA compared to StyleGAN2 (SG2) on FFHQ [20] ( $256 \times 256$ ). (a) The discriminator raw output logits of StyleGAN2 on the full (70k) or limited (7k) datasets. (b) The discriminator raw output logits of StyleGAN2 and APA on the limited (7k) dataset. (c) The training convergence shown by FID.



Table 3: The FID (lower is better) and IS (higher is better) scores ( $256 \times 256$ ) 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 [21]	7.737	1.825	37.830	4.018	<u>3.862</u>	5.243
ADA [19]	6.053	2.119	<u>11.409</u>	<u>4.721</u>	4.018	<u>5.329</u>
LC-Reg [46]	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	<b>3.678</b>	<b>5.336</b>
ADA + APA (Ours)	<b>4.377</b>	<b>2.169</b>	<b>8.379</b>	<b>4.849</b>	3.811	5.321

Table 4: The FID (lower is better) and IS (higher is better) scores ( $256 \times 256$ ) of our method **compared to previous techniques for regularizing GANs** on StyleGAN2 trained with FFHQ-5k [20].

Metric	StyleGAN2 [21]	Instance noise [44]	One-sided LS [42]	APA (Ours)
FID ↓	37.830	40.981	33.978	<b>13.249</b>
IS ↑	4.018	4.231	4.029	<b>4.487</b>

improving the StyleGAN2 baseline on both the limited and full datasets. Although ADA [19] 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 [46] 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 [19]. As can be observed in Table 3, APA can further boost the performance of StyleGAN2 with ADA [19] given limited training data, suggesting the compatibility of our approach with standard data augmentations. Combining ADA [19] and APA on FFHQ full data outperforms StyleGAN2 but is slightly inferior to applying APA solely. The degraded performance is mainly affected by ADA [19], which we empirically found might slightly harm the performance when the training data is sufficient. Overall, these methods are not dedicated to improving performance under sufficient data. Nevertheless, as a beneficial side effect, the proposed APA may have this potential. More comparative results are included in the *Appendix*.

**Training cost.** We compare the computational cost of APA against ADA [19], using the same basic official codebase. There is no parameter or memory increment for both methods. As for the time consumption, we test the training cost on 8 NVIDIA Tesla V100 GPUs. On FFHQ-5k ( $256 \times 256$ ), the average training time of the StyleGAN2 [21] baseline is  $(4.740 \pm 0.100)$  sec/king (*i.e.*, seconds per thousand of images shown to the discriminator). The cost of our method is negligible, slightly increasing this value to  $(4.789 \pm 0.078)$  sec/king. As a reference, the value for ADA [19] is  $(5.327 \pm 0.116)$  sec/king, spending additional time for applying external augmentations.

### 4.3 Comparison with Previous Techniques for Regularizing GANs

As mentioned in Section 2, APA is closely related to previous techniques for regularizing GANs. The comparative results on APA and some representative conventional techniques, *i.e.*, instance noise [44] and one-sided label smoothing (LS) [42], are shown in Table 4. Applying instance noise [44] not boost the performance of StyleGAN2 much under limited data. One-sided label smoothing (LS) [42] (with the real label of 0.9) outperforms StyleGAN2 but still has a huge performance gap with our method. This further suggests the effectiveness and usefulness of APA.

### 4.4 Ablation Studies

**Ablation studies on variants of APA.** We study three key elements of APA, *i.e.*, the overfitting heuristic  $\lambda$ , the deception strength  $p$ , and the deception strategy. The version used in our main experiments is denoted as the “main” version, where  $\lambda = \lambda_r$ , and  $p$  is adjusted adaptively. Besides, the “main” version applies the deception strategy that is analogous to one-sided label flipping. As reported in Table 5, when using other variants of  $\lambda$  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

Table 5: **Ablation studies on variants of APA** on FFHQ-5k [20] ( $256 \times 256$ ). We study three key elements of APA, *i.e.*, the overfitting heuristic  $\lambda$ , the deception strength  $p$ , and the deception strategy. The “main” denotes the main version used in our previous experiments (*i.e.*,  $\lambda = \lambda_r$ ,  $p$  is adaptively adjusted, and the deception strategy is analogous to one-sided label flipping).

Metric	StyleGAN2	main	$\lambda = \lambda_f$	$\lambda = \lambda_{rf}$	$p = 0.5$ (fix)	two-sided
FID ↓	37.830	<u>13.249</u>	13.470	<b>12.679</b>	14.632	15.440
IS ↑	4.018	<b>4.487</b>	<u>4.420</u>	4.412	4.403	4.167

Table 6: **Ablation studies on the threshold  $t$**  on FFHQ-5k [20] ( $256 \times 256$ ). The “main” denotes the main version used in our previous experiments (*i.e.*,  $t = 0.6$ ).

Metric	StyleGAN2	$t = 0.4$	$t = 0.6$ (main)	$t = 0.8$
FID ↓	37.830	13.687	<b>13.249</b>	13.984
IS ↑	4.018	4.418	<b>4.487</b>	4.395

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. As for the deception strategy, we empirically observe that two-sided label flipping can still outperform StyleGAN2 but is inferior to the main version.

**Ablation studies on the threshold  $t$ .** We further provide the ablation studies on the threshold value  $t$  in Table 6. The version used in our main experiments is denoted as the “main” version, where  $t = 0.6$ . It can be seen that the models with different values of  $t$  achieve comparable results, outperforming the StyleGAN2 baseline. On FFHQ-5k ( $256 \times 256$ ),  $t = 0.6$  could be a more plausible choice. For a further explanation, we use  $t = 0.6$  as the default value since it works well in most cases. In practice, the value of  $t$  could be further adjusted to achieve even better results. Empirically, a smaller  $t$  can be chosen when one has fewer data. This means the deception strength  $p$  can be adjusted to increase more rapidly since the discriminator is more prone to overfitting when the data amount is fewer.

## 5 Discussion

We have shown the effectiveness of the proposed adaptive pseudo augmentation (APA) for state-of-the-art GAN training with limited data 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 [43] to enhance diversity. In addition, we only theoretically verified the convergence and rationality of APA. In future work, the theoretical analysis on the effectiveness of APA could be further explored.

**Broader impact.** On the one hand, the effectiveness of APA with negligible computational cost will benefit the practical deployment of GANs, especially in the low-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.

**Acknowledgments and funding disclosure.** This study is supported under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

## References

- [1] faceswap-GAN. <https://github.com/shaoanlu/faceswap-GAN/>. Accessed: 2019-08-16.
- [2] Anonymous, Danbooru community, and Gwern Branwen. Danbooru2019: A large-scale crowdsourced and tagged anime illustration dataset. <https://www.gwern.net/Danbooru2020#danbooru2019>. Accessed: 2021-04-10.
- [3] Martin Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks. In *ICLR*, 2017.
- [4] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *ICML*, 2017.
- [5] 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.
- [6] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *ICLR*, 2018.
- [7] Soumith Chintala. How to train a GAN? NeurIPS workshop on adversarial training. <https://sites.google.com/site/nips2016adversarial/>, 2016. Accessed: 2021-08-03.
- [8] 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.
- [9] Yunjey Choi, Youngjung Uh, Jaegul Yoo, and Jung-Woo Ha. StarGAN v2: Diverse image synthesis for multiple domains. In *CVPR*, 2020.
- [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.
- [11] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans. In *NeurIPS*, 2017.
- [12] 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.
- [13] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, 2018.
- [14] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017.
- [15] Simon Jenni and Paolo Favaro. On stabilizing generative adversarial training with noise. In *CVPR*, 2019.
- [16] Liming Jiang, Bo Dai, Wayne Wu, and Chen Change Loy. Focal frequency loss for image reconstruction and synthesis. In *ICCV*, 2021.
- [17] 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.
- [18] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *ICLR*, 2018.
- [19] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *NeurIPS*, 2020.
- [20] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019.
- [21] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020.
- [22] Vahid Kazemi and Josephine Sullivan. One millisecond face alignment with an ensemble of regression trees. In *CVPR*, 2014.
- [23] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*, arXiv:1412.6980, 2014.

- [24] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [25] 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.
- [26] 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.
- [27] Jianhua Lin. Divergence measures based on the Shannon entropy. *TIT*, 37:145–151, 1991.
- [28] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NeurIPS*, 2017.
- [29] 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.
- [30] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *ICCV*, 2017.
- [31] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *ICML*, 2018.
- [32] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint*, arXiv:1411.1784, 2014.
- [33] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018.
- [34] Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-GAN: Training generative neural samplers using variational divergence minimization. In *NeurIPS*, 2016.
- [35] 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.
- [36] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, 2019.
- [37] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. StyleCLIP: Text-driven manipulation of stylegan imagery. In *ICCV*, 2021.
- [38] 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.
- [39] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint*, arXiv:1511.06434, 2015.
- [40] 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.
- [41] Kevin Roth, Aurelien Lucchi, Sebastian Nowozin, and Thomas Hofmann. Stabilizing training of generative adversarial networks through regularization. In *NeurIPS*, 2017.
- [42] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training GANs. In *NeurIPS*, 2016.
- [43] Omry Sendik, Dani Lischinski, and Daniel Cohen-Or. Unsupervised multi-modal styled content generation. *arXiv preprint*, arXiv:2001.03640, 2020.
- [44] Casper Kaae Sønderby, Jose Caballero, Lucas Theis, Wenzhe Shi, and Ferenc Huszár. Amortised map inference for image super-resolution. In *ICLR*, 2017.
- [45] 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.
- [46] Hung-Yu Tseng, Lu Jiang, Ce Liu, Ming-Hsuan Yang, and Weilong Yang. Regularizing generative adversarial networks under limited data. In *CVPR*, 2021.



- [47] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The Caltech-UCSD Birds-200-2011 Dataset. 2011.
- [48] 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.
- [49] Ryan Webster, Julien Rabin, Loic Simon, and Frédéric Jurie. Detecting overfitting of deep generative networks via latent recovery. In *CVPR*, 2019.
- [50] Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. Caltech-UCSD Birds 200. 2010.
- [51] 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.
- [52] Yuanbo Xiangli, Yubin Deng, Bo Dai, Chen Change Loy, and Dahua Lin. Real or not real, that is the question. In *ICLR*, 2020.
- [53] Han Zhang, Zizhao Zhang, Augustus Odena, and Honglak Lee. Consistency regularization for generative adversarial networks. In *ICLR*, 2020.
- [54] Junbo Zhao, Michael Mathieu, and Yann LeCun. Energy-based generative adversarial networks. In *ICLR*, 2017.
- [55] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable augmentation for data-efficient gan training. In *NeurIPS*, 2020.
- [56] Zhengli Zhao, Sameer Singh, Honglak Lee, Zizhao Zhang, Augustus Odena, and Han Zhang. Improved consistency regularization for GANs. In *AAAI*, 2021.
- [57] Zhengli Zhao, Zizhao Zhang, Ting Chen, Sameer Singh, and Han Zhang. Image augmentations for GAN training. *arXiv preprint*, arXiv:2006.02595, 2020.
- [58] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.

## Appendix

This appendix provides supplementary information that is not elaborated in our main paper: Section A describes some preliminary concepts and definitions of our methodology. Section B provides additional discussions on previous techniques for regularizing GANs. Section C details our used datasets. Section D presents the implementation details of our experiments. Section E shows more results and analysis.

### A Preliminaries of Methodology

To further facilitate readers’ understanding of the theoretical analysis and the proposed Adaptive Pseudo Augmentation (APA), this section will provide some preliminary concepts and definitions for the methodology section in our main paper.

Generative adversarial networks (GANs) [10] aim at capturing the real data distribution to synthesize new data. Two networks are trained alternately via an adversarial process: a generator  $G$  learns to produce new samples, and a discriminator  $D$  (*i.e.*, a binary classifier) predicts the probability that a sample comes from the real data rather than from  $G$ . Following [10, 21], our main paper focuses on the fundamental problem of GANs, *i.e.*, unconditional image synthesis, which is generating random samples from a noise input in the latent space. The noise is sampled from a Gaussian distribution.

The goal of GANs is to learn an ideal generated distribution  $p_g$  from the real data distribution  $p_{\text{data}}$ . Let  $p_z(z)$  be the prior on the input noise variable. The mapping from the latent space to the image space can be denoted as  $G(z)$ . For sample  $x$ ,  $D(x)$  represents the estimated probability of  $x$  coming from the real data. Here, both  $G$  and  $D$  should be differentiable functions that are defined by the network parameters. To quantify the adversarial process,  $G$  and  $D$  play a minimax two-player game with the value function  $V(G, D)$ :

$$\min_G \max_D V(G, D) = \mathbb{E}_{x \sim p_{\text{data}}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(z)} [\log (1 - D(G(z)))] . \quad (12)$$

Let the virtual training criterion [10] for the generator  $G$  be  $C(G)$ . The global minimum of  $C(G)$  is achieved if and only if  $p_g = p_{\text{data}}$ , and the minimum value is  $-\log 4$ , as proved by [10]. This indicates that GANs can perfectly model the real data distribution if given sufficient capacity and time. In practice, we usually use a non-saturated form for  $G$  and train it to maximize  $\log D(G(z))$  instead of minimizing  $\log(1 - D(G(z)))$  to ensure a healthy gradient at the early training stage.

In all the figures of this work, we show raw output logits of  $D$  before the last Sigmoid activation to better visualize its prediction confidence. Let  $\text{logit}$  denotes the logit function, we define:

$$D_{\text{real}} = \text{logit}(D(x)), \quad D_{\text{fake}} = \text{logit}(D(G(z))). \quad (13)$$

## B Additional Discussions on Previous Techniques for Regularizing GANs

As mentioned in our main paper, previous techniques for regularizing GANs include adding noise to the inputs of the discriminator [3, 44, 15], gradient penalty [31, 11, 41], one-sided label smoothing [42], spectral normalization [33], label noise [7], *etc.* These approaches are designed for stabilizing training or preventing mode collapse [42]. The essence of their objectives could be considered similar to our method since training GANs in the low-data regime exhibits similar behaviors as previously observed in early GANs with sufficient data. However, several differences are worth highlighting.

First of all, the actual goals of previous strategies and the proposed APA may not be completely the same. Specifically, APA is specialized for training GANs in the low-data regime, which was not carefully considered by prior studies. This difference in the data setting is quite important since it is the core problem we wish to address. Even the performance of state-of-the-art StyleGAN2 [21] deteriorates when trained with a limited amount of data, although it has exploited many advanced techniques for stabilizing training or preventing mode collapse, such as R1 regularization [31]. The main challenge that lies within the low-data regime is the overfitting of the discriminator. Although this issue might also appear on early GANs, it becomes more severe when data is limited.

Besides, we have presented the comparative studies on APA and conventional techniques in the main paper. Empirically, we showed that previous methods could not boost the performance of StyleGAN2 much under limited data. Some of them showed a huge performance gap in comparison to the proposed APA. The experiments indicate that previous relevant strategies cannot handle the low-data regime well, further suggesting the effectiveness and usefulness of APA.

Last but not least, APA and these techniques themselves are not exactly the same. The proposed APA is an effective practice and improvement of these ideas on modern GANs, whose implementations are very different from the early ones. Compared to previous techniques, APA is more adaptive to fit different settings and the overfitting status in training. Although using adaptive heuristics was also explored in the past, it had been found unpractical at the time [7]. APA makes the adaptive control scheme possible in practice.

We believe that the proposed APA could contribute to the community for its effectiveness, simplicity, and adaptability for training state-of-the-art GANs in the low-data regime. Hopefully, our approach could extend the breadth and potential of solutions to GAN training with limited data.

## C Dataset Details

This section will detail our explored datasets in the main paper. We randomly select subsets to confine the size of large datasets and directly use small datasets for GAN training with limited data under different settings. Our main paper focuses on the fundamental unconditional image synthesis task with powerful contemporary GANs. Thus, there is no need to split a separate test set. We exploit a high-quality Lanczos filter [25] for image resizing and save images in the uncompressed PNG format. We will detail each dataset separately as follows.

- **AFHQ-Cat.** We use the AFHQ-Cat dataset released by [9], which is the cat category of the high-quality Animal Faces-HQ (AFHQ) dataset [9]. The original authors mentioned that they collected images with permissive licenses from the Flickr and Pixabay websites. All images were vertically and horizontally aligned at the center. The dataset was released under Creative Commons BY-NC 4.0 license by NAVER Corporation. As a small dataset,

we exploit 5,153 training images of cat faces with various breeds. The original resolution is  $512 \times 512$ , which is scaled to  $256 \times 256$  for training.

- **FFHQ.** The full Flickr-Faces-HQ (FFHQ) dataset [20] consists of 70,000 high-quality ( $1024 \times 1024$ ) human face images. The faces contain considerable variation in terms of age, ethnicity, background, and accessories. The original authors mentioned that the images were crawled from the Flickr website, and only the ones under permissive licenses were collected. The images were automatically aligned [22] and cropped. The individual images were published in Flickr by their respective authors under either Creative Commons BY 2.0, Creative Commons BY-NC 2.0, Public Domain Mark 1.0, Public Domain CC0 1.0, or U.S. Government Works license. The dataset itself was made available under Creative Commons BY-NC-SA 4.0 license by NVIDIA Corporation. We randomly select different subsets of FFHQ, *i.e.*, FFHQ-1k (1,000 images,  $\sim 1.4\%$  data), FFHQ-5k (5,000 images,  $\sim 7\%$  data), FFHQ-7k (7,000 images,  $10\%$  data), and FFHQ-70k (70,000 images,  $100\%$  data) to perform GAN training given different data amounts. The original images are resized to  $256 \times 256$  for training in the experiments of our main paper. We will also show some synthesized examples with the original resolution in this Appendix.
- **Danbooru2019 Portraits (Anime).** The Danbooru2019 Portraits (Anime) [5] is a dataset consisting of  $512 \times 512$  anime faces cropped from solo SFW Danbooru2019 images [2]. The full dataset contains a total of 302,652 images in a broad portrait style encompassing ears, necklines, hats, *etc.*, rather than the tightly cropped faces. The dataset was released under the Creative Commons public domain (CC-0) license. We artificially confine the dataset into a subset for GAN training, *i.e.*, Anime-5k (5,000 images,  $\sim 2\%$  data). The training image size is  $256 \times 256$ .
- **Caltech-UCSD Birds-200-2011 (CUB).** The Caltech-UCSD Birds-200-2011 (CUB) [47] is an extended version of CUB-200 [50], a challenging dataset containing 200 bird species. The dataset consists of 11,788 bird images at diverse locations with heavy background clutter. The images were harvested using the Flickr image search. We were unable to find the license information about this dataset. We employ all the available images for GAN training since the dataset is already small and difficult. The resolutions of original images vary, and we uniformly resize them to  $256 \times 256$ .

## D Implementation Details

In the main paper, we choose the state-of-the-art StyleGAN2 [21] as the backbone to verify the effectiveness of the proposed APA on limited data. Besides, we compare our method with representative approaches designed for the low-data regime, including the adaptive discriminator augmentation (ADA) [19] and LC-regularization (LC-Reg) [46], which perform standard data augmentations and model regularization, respectively. In addition, we compare APA with representative conventional techniques for regularizing GANs, *i.e.*, instance noise [44] and one-sided label smoothing [42]. For a fair and controllable comparison, we reimplement all baselines and run the experiments from scratch using official code. The qualitative and quantitative results of each method are reported using the best model throughout training.

We ran our experiments on an internal computing cluster with Slurm Workload Manager. All the models are trained on 8 NVIDIA Tesla V100 GPUs with 32 GB memory capacity. We follow [19] and employ its mixed-precision FP16/FP32 training scheme in all our experiments. The actual memory consumption for each model is around 11 GB per GPU. We perform 25,000 kimg (*i.e.*, thousands of images shown to the discriminator, measuring the training progress [20, 21, 19]) of training for each model. For the average training time cost of different models, please refer to the Training cost of Section 4.2 in the main paper.

We implemented the proposed APA on top of the official implementation of StyleGAN2 [21]. The network architecture is kept unchanged. The mapping network contains 8 fully connected layers, and the dimensionality of the input and intermediate latent space is 512. We use the combination of a generator with output skips and a residual discriminator. The detailed structures of the generator and the discriminator are the same as [21], *e.g.*, using weight demodulation [21] in the generator. The activation function is Leaky ReLU with a negative slope of 0.2. We apply several other standard techniques in [18, 20, 21], including the mini-batch standard deviation layer at the end of the

Table 7: The FID (lower is better) and IS (higher is better) scores ( $256 \times 256$ ) of **transfer learning** on MetFaces [19] with limited data amounts from the pre-trained StyleGAN2 model on FFHQ-70k [20].

Method	MetFaces-1336 (full data)		MetFaces-500 ( $\sim 37\%$ data)	
	FID $\downarrow$	IS $\uparrow$	FID $\downarrow$	IS $\uparrow$
StyleGAN2 [21]	30.988	3.719	54.691	3.218
APA (Ours)	<b>21.050</b>	<b>4.103</b>	<b>29.508</b>	<b>3.986</b>

Table 8: The FID (lower is better) and IS (higher is better) scores ( $256 \times 256$ ) **under a lower amount of training data** on FFHQ-500 [20] (a subset of 500 images,  $\sim 0.7\%$  of full data).

Method	FFHQ-500 ( $\sim 0.7\%$ data)	
	FID $\downarrow$	IS $\uparrow$
StyleGAN2 [21]	119.815	2.446
APA (Ours)	<b>50.989</b>	<b>4.099</b>

discriminator [18], equalized learning rate for all the trainable parameters [18], pixel-wise feature vector normalization [18], the exponential moving average of generator weights [18], style mixing regularization [20], path length regularization [21], and lazy regularization [21]. The training loss is the non-saturating logistic loss [10, 21] with  $R_1$  regularization [31]. The batch size is 64 for our experiments trained with the resolution of  $256 \times 256$ . The Adam [23] optimizer is applied with  $\beta_1 = 0$ ,  $\beta_2 = 0.99$ . For other network and training details, we follow the original paper and official code of StyleGAN2 [21].

For APA, we set the overfitting heuristic  $\lambda = \lambda_r$  in our main experiments and study other variants (*i.e.*,  $\lambda = \lambda_f$  and  $\lambda = \lambda_{rf}$ ) through the ablation study. Aside from the ablation study of a fixed deception probability  $p = 0.5$ ,  $p$  is adaptively adjusted according to  $\lambda$ . The adaptive adjustment of  $p$  is as follows. We first initialize  $p$  to be zero and set a threshold value  $t$  ( $t = 0.6$  in most cases unless specified otherwise). If  $\lambda$  signifies too much/little overfitting regarding  $t$  (*i.e.*, larger/smaller than  $t$ ), the probability  $p$  will be increased/decreased by one fixed step. Using this step size,  $p$  can increase quickly from zero to one, *i.e.*, in 500k images shown to  $D$ . We adjust  $p$  once every four iterations. We clamp  $p$  from below to zero after each adjustment so that  $p$  can always be larger than zero. We do not set an upper bound for  $p$  while it can be naturally restricted under a safe limit. Then, the pseudo augmentation of each instance will be applied with the probability  $p$  or be skipped with the probability  $1 - p$ . In this way, the strength of pseudo augmentation can be adaptively controlled based on the degree of overfitting throughout training.

As for other methods used for comparison, we strictly follow all the details in their papers and released official code using recommended setups.

## E More Results and Analysis

**Effectiveness of APA for transfer learning.** Table 7 reports the additional transfer learning results on MetFaces [19] from the pre-trained StyleGAN2 model on FFHQ-70k [20]. All the metrics can be boosted by APA, further verifying its effectiveness for transfer learning besides training from scratch.

**The performance of APA under a lower amount of data (*i.e.*, fewer than 1k).** We have reported transfer learning results on MetFaces-500 [19] (a subset of 500 images,  $\sim 37\%$  of full data) in Table 7. Table 8 shows more results on FFHQ-500 [20] (a subset of 500 images,  $\sim 0.7\%$  of full data). Under fewer data, APA can still boost StyleGAN2 performance by a large margin. However, the quality itself of synthesized images by APA can still be improved, consistent with our discussion on limitations in the main paper.

**More examples of the effectiveness of APA on various datasets.** We show more comparative results of StyleGAN2 [21] and the proposed APA to demonstrate the effectiveness of our method to improve the state-of-the-art baseline on various datasets with limited data amounts, *i.e.*, AFHQ-Cat-5k [9] (Figure 7), FFHQ-5k [20] (Figure 8), Anime-5k [5] (Figure 9), and CUB-12k [47] (Figure 10). Regardless of applying the truncation trick [6, 20, 21] or not, the proposed APA can significantly ameliorate the degraded synthesis quality of StyleGAN2 with limited training data in all cases. The generated images by our method are highly photorealistic using only limited training data, being closer to the real data distributions.

**Qualitative results of comparison with other solutions for GAN training with limited data.** In the main paper, we have quantitatively shown the effectiveness of the proposed APA over other state-





Figure 7: More examples of the effectiveness of our method to improve state-of-the-art StyleGAN2 [21] synthesized results ( $256 \times 256$ ) on **AFHQ-Cat-5k** [9] (5, 153 images, which is small by itself).



Figure 8: More examples of the effectiveness of our method to improve state-of-the-art StyleGAN2 [21] synthesized results ( $256 \times 256$ ) on **FFHQ-5k** [20] (a subset of 5, 000 images,  $\sim 7\%$  of full data).





Figure 9: More examples of the effectiveness of our method to improve state-of-the-art StyleGAN2 [21] synthesized results ( $256 \times 256$ ) on **Anime-5k** [5] (a subset of 5,000 images,  $\sim 2\%$  of full data).



Figure 10: More examples of the effectiveness of our method to improve state-of-the-art StyleGAN2 [21] synthesized results ( $256 \times 256$ ) on **CUB-12k** [47] (11,788 images, which is small by itself).



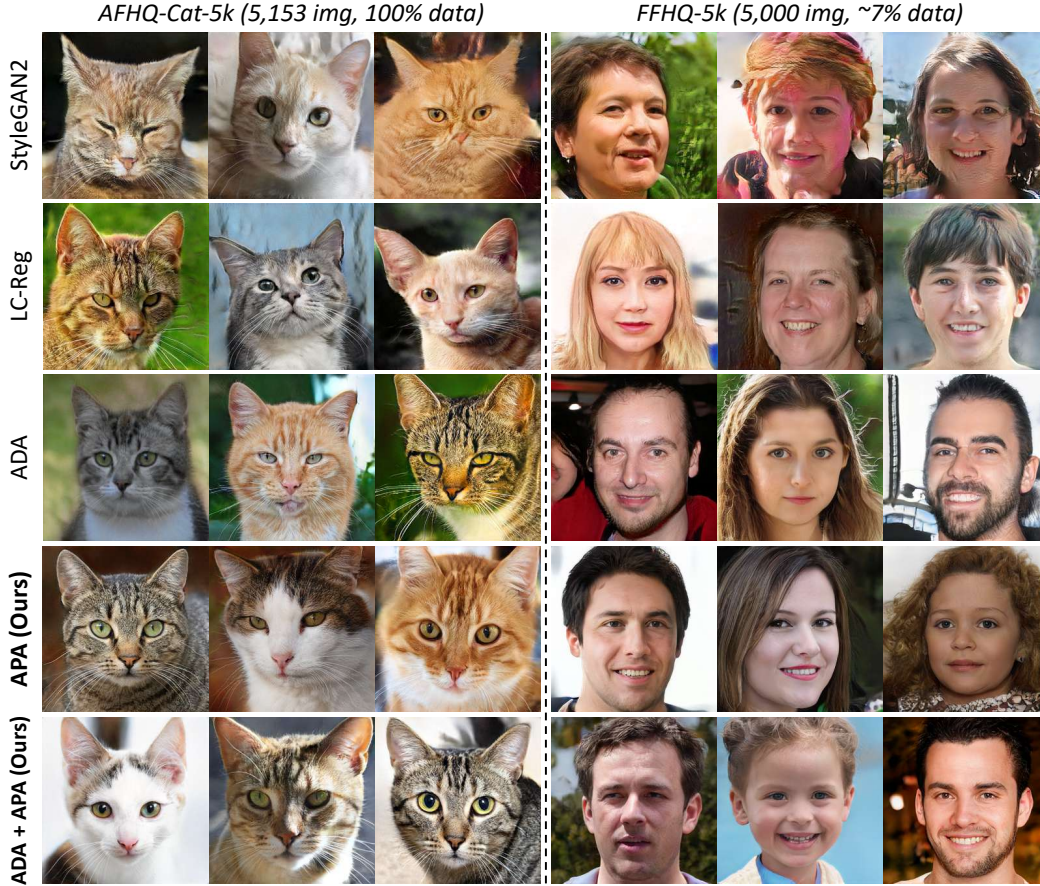


Figure 11: The synthesized results ( $256 \times 256$ , no truncation) of our method **compared to other state-of-the-art solutions designed for GAN training with limited data** on StyleGAN2 [21]. We confine the data amount of FFHQ [20] and directly use AFHQ-Cat [9] that is small by itself.

Table 9: The FID (lower is better) and IS (higher is better) scores ( $256 \times 256$ ) of **additional comparison with the state-of-the-art ADA [19]** tailored for GAN training with limited data on StyleGAN2 trained with FFHQ [20] and MetFaces [19]. The bold number indicates the best value, and the underline marks the second best.

Method	FFHQ-7k		FFHQ-1k		MetFaces-1336		MetFaces-500	
	FID ↓	IS ↑	FID ↓	IS ↑	FID ↓	IS ↑	FID ↓	IS ↑
StyleGAN2 [21]	27.738	4.264	86.407	2.806	30.988	3.719	54.691	3.218
ADA [19]	10.275	4.813	22.590	4.239	20.834	4.005	30.368	3.974
APA (Ours)	10.800	<u>4.860</u>	45.192	4.130	21.050	<u>4.103</u>	<u>29.508</u>	<u>3.986</u>
ADA + APA (Ours)	<b>7.333</b>	<b>4.994</b>	<b>18.892</b>	<b>4.316</b>	<b>18.865</b>	<b>4.207</b>	<b>28.408</b>	<b>4.044</b>

of-the-art approaches designed for the low-data regime, including ADA [19] and LC-regularization (LC-Reg) [46], which perform standard data augmentations and model regularization, respectively. In Figure 11, we show the qualitative results for further illustration. The synthesis quality of the StyleGAN2 [21] baseline deteriorates on the limited amount of training data. Ripple artifacts and substantial distortions appear on the generated images by StyleGAN2. LC-Reg [46] slightly improves the quality of synthesis by reducing the distortions on the images. The amelioration of visual quality is more evident by applying ADA [19], where the ripple artifacts are clearly subsided while some minor artifacts exist on the hair and beard. The proposed APA achieves comparable or even better visual quality than LC-Reg [46] and ADA [19], effectively improving the StyleGAN2 synthesized results on limited data. Notably, APA is also complementary to ADA [19] for gaining a further performance boost, suggesting the compatibility of our method with standard data augmentations. The visual results we present here are in line with the quantitative performance in our main paper.

**More comparative results with ADA [19].** Table 9 reports additional comparative results with the state-of-the-art ADA [19] tailored for GAN training with limited data to further highlight the

StyleGAN2



APA (Ours)



Figure 12: The effectiveness of our method to improve state-of-the-art **StyleGAN2** [21] higher-resolution synthesized results ( $1024 \times 1024$ , no truncation) **trained with the limited data**, *i.e.*, FFHQ-5k [20] (a subset of 5,000 images,  $\sim 7\%$  of full data). Our method achieves an FID score of **9.545**, outperforming the original StyleGAN2 of 18.296.






Figure 13: The **evolution of the deception probability  $p$**  during training on FFHQ [20] ( $256 \times 256$ ) with different data amounts. The “kimg” denotes thousands of real images shown to the discriminator.

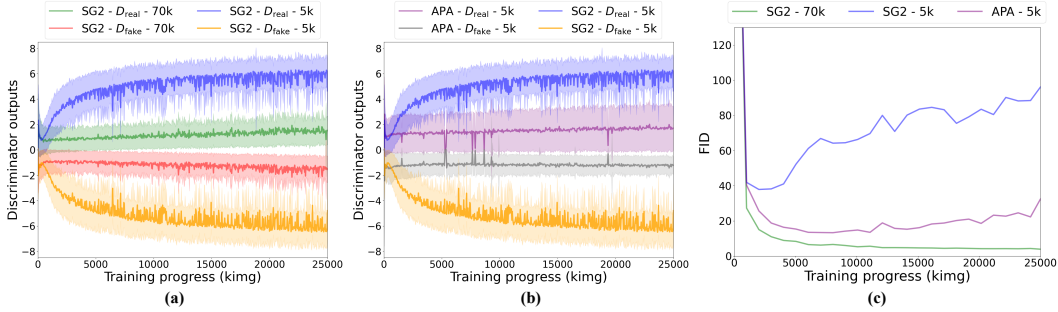


Figure 14: The **overfitting and convergence status** of APA compared to StyleGAN2 (SG2) on FFHQ [20] ( $256 \times 256$ ). (a) The discriminator raw output logits of StyleGAN2 on the full (70k) or limited (5k) datasets. (b) The discriminator raw output logits of StyleGAN2 and APA on the limited (5k) dataset. (c) The training convergence shown by FID.

advantages of the proposed APA. We include additional comparisons on FFHQ-7k [20] and FFHQ-1k (see our main paper for 5k and 70k). To make the comparison settings more comprehensive, we also provide the transfer learning results on MetFaces [19] from the pre-trained StyleGAN2 model on FFHQ-70k. We also make a comparison with a fewer data amount (*i.e.*, 500 images). Combined with the results we presented in the main paper, the proposed APA achieves comparable or even better performance than ADA [19] while with less computational cost. Both methods outperform the StyleGAN2 baseline under limited data. Although applying APA solely may be inferior to ADA [19] on FFHQ-1k (in line with our discussion in the main paper), it is worth mentioning that APA is also complementary to ADA [19], which is very important to boost the performance further.

**Higher-resolution examples on StyleGAN2.** In Figure 12, we show some higher-resolution ( $1024 \times 1024$ ) synthesized images on FFHQ-5k [20] (a subset of 5,000 images,  $\sim 7\%$  of full data) to further illustrate the effectiveness of our approach in improving StyleGAN2 with limited training data. The truncation trick [6, 20, 21] is not applied. The proposed APA evidently improves StyleGAN2 in both perceptual quality and the FID score, indicating its effectiveness for GAN training with limited data.

**Evolution of the deception probability.** In Figure 13, we visualize the evolution of the deception probability  $p$  in training on FFHQ [20] ( $256 \times 256$ ) with different data amounts. The evolution of  $p$  may be relatively more unstable when the data amount is very limited (*i.e.*, 1k). Notably, the proposed APA possesses a desired property that the deception probability  $p$  can be naturally restricted under a safe limit ( $\sim 0.8$ ) regardless of the training data amounts. Hence, the fundamental capability of discriminator in adversarial training may be better preserved thanks to the adaptive control scheme.

**More overfitting and convergence analysis.** Aside from the overfitting and convergence analysis on FFHQ-7k (10% of full data) in our main paper, we present the additional analysis on FFHQ-5k ( $\sim 7\%$  of full data) and FFHQ-1k ( $\sim 1.4\%$  of full data) in Figure 14 and Figure 15, respectively. The trend of model overfitting and convergence on FFHQ-5k and FFHQ-1k is similar to FFHQ-7k.

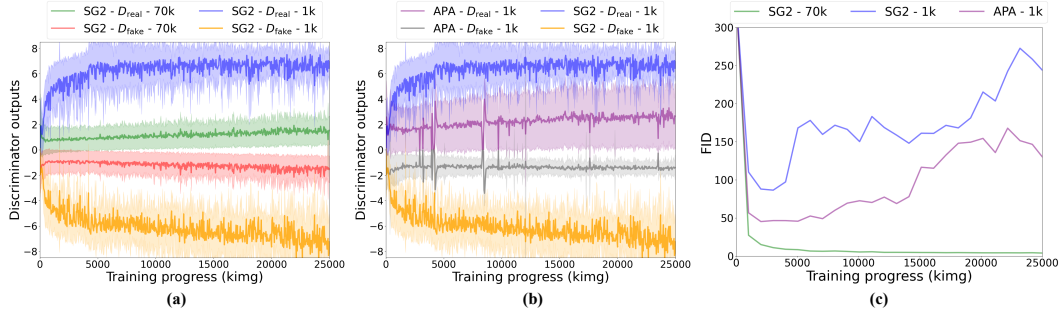


Figure 15: The **overfitting and convergence status** of APA compared to StyleGAN2 (SG2) on FFHQ [20] ( $256 \times 256$ ). (a) The discriminator raw output logits of StyleGAN2 on the full (70k) or limited (1k) datasets. (b) The discriminator raw output logits of StyleGAN2 and APA on the limited (1k) dataset. (c) The training convergence shown by FID.




Figure 16: The **overfitting and convergence status** of APA compared to previous techniques for regularizing GANs on FFHQ-5k [20] ( $256 \times 256$ ,  $\sim 7\%$  data). (a) The discriminator raw output logits of StyleGAN2 (SG2) and APA. (b) The discriminator raw output logits of instance noise (ISN) [44] and APA. (c) The discriminator raw output logits of one-sided label smoothing (OLS) [42] and APA.

Table 10: The FID (lower is better) and IS (higher is better) scores ( $32 \times 32$ ) of **class-conditional image synthesis on BigGAN** trained with the subsets of **CIFAR-10** [24] with limited data amounts.

Method	full data		20% data		10% data	
	FID ↓	IS ↑	FID ↓	IS ↑	FID ↓	IS ↑
BigGAN [6]	9.531	9.078	22.024	8.343	44.061	7.589
APA (Ours)	<b>8.283</b>	<b>9.362</b>	<b>15.316</b>	<b>8.822</b>	<b>25.987</b>	<b>8.410</b>

Table 11: The FID (lower is better) and IS (higher is better) scores ( $32 \times 32$ ) of **class-conditional image synthesis on BigGAN** trained with the subsets of **CIFAR-100** [24] with limited data amounts.

Method	full data		20% data		10% data	
	FID ↓	IS ↑	FID ↓	IS ↑	FID ↓	IS ↑
BigGAN [6]	13.281	10.525	35.590	8.706	64.828	6.635
APA (Ours)	<b>11.429</b>	<b>11.243</b>	<b>23.506</b>	<b>9.811</b>	<b>45.794</b>	<b>8.114</b>

The divergence of StyleGAN2 discriminator predictions can be effectively restricted by applying the proposed APA, indicating its effectiveness in mitigating the discriminator overfitting. Meanwhile, APA improves the training convergence measured by FID.

Furthermore, we show the comparative overfitting analysis on APA and previous techniques for regularizing GANs in Figure 16. We observe that the effectiveness to counteract overfitting is in line with the generation performance of these methods in the main paper. Specifically, StyleGAN2 experiences diverged predictions most rapidly, and APA obtains the most effective restriction on the divergence of discriminator outputs. Applying instance noise (ISN) [44] produces curves that are very close to StyleGAN2, and one-sided label smoothing (OLS) [42] only restricts the real predictions. This further suggests: 1) the importance of addressing the discriminator overfitting for training GANs with limited data; 2) the effectiveness of APA in alleviating the discriminator overfitting, outperforming previous strategies.

**Additional training convergence visualizations.** Please refer to our [supplementary video](#) for additional training convergence visualizations on FFHQ-7k [20] (7,000 images, 10% of full data). The truncation trick [6, 20, 21] with  $\psi = 0.7$  is applied to the synthesized images. The proposed APA effectively improves the training convergence of StyleGAN2 on limited data.

**Class-conditional image synthesis with BigGAN [6].** To further enrich our benchmark and demonstrate the effectiveness of the proposed APA under diverse settings, we show additional class-conditional image synthesis results on the state-of-the-art BigGAN [6] trained with CIFAR-10 [24] and CIFAR-100 [24] in Table 10 and Table 11, respectively. It can be observed that APA outperforms BigGAN under limited training data in all cases, further suggesting its adaptability and effectiveness for class-conditional image synthesis with other powerful contemporary GANs, such as BigGAN.