FakeCLR: Exploring Contrastive Learning for Solving Latent Discontinuity in Data-Efficient GANs

Ziqiang Li^{1[0000-0001-9484-2310]} *, Chaoyue Wang², Heliang Zheng², Jing Zhang³, and Bin Li¹

¹ University of Science and Technology of China, China ² JD Explore Academy, China ³ The University of Sydney, Australia iceli@mail.ustc.edu.cn, chaoyue.wang@outlook.com, zhenghl@mail.ustc.edu.cn, jing.zhang1@sydney.edu.au, binli@ustc.edu.cn

Abstract. Data-Efficient GANs (DE-GANs), which aim to learn generative models with a limited amount of training data, encounter several challenges for generating high-quality samples. Since data augmentation strategies have largely alleviated the training instability, how to further improve the generative performance of DE-GANs becomes a hotspot. Recently, contrastive learning has shown the great potential of increasing the synthesis quality of DE-GANs, yet related principles are not well explored. In this paper, we revisit and compare different contrastive learning strategies in DE-GANs, and identify (i) the current bottleneck of generative performance is the discontinuity of latent space; (ii) compared to other contrastive learning strategies, Instance-perturbation works towards latent space continuity, which brings the major improvement to DE-GANs. Based on these observations, we propose FakeCLR, which only applies contrastive learning on perturbed fake samples, and devises three related training techniques: Noise-related Latent Augmentation, Diversity-aware Queue, and Forgetting Factor of Queue. Our experimental results manifest the new state of the arts on both few-shot generation and limited-data generation. On multiple datasets, FakeCLR acquires more than 15% FID improvement compared to existing DE-GANs. Code is available at https://github.com/iceli1007/FakeCLR.

Keywords: Data-Efficient Training, GANs, Contrastive Learning

1 Introduction

Generative Adversarial Networks [1] are capable of generating realistic images [2] indistinguishable from real ones, which acquire remarkable achievements in image processing [3,4,5,6] and computer vision [7,8,9,10]. The success of GANs relies on large-scale training datasets, for example, CelebA consists of around

 $^{^{\}star}$ This work was performed when Ziqiang Li was visiting JD Explore Academy as a research intern.



Fig. 1: Training GANs with as little as 100 training samples (Obama dataset) typically results in severe discontinuity in latent space, *i.e.* under-diversity interpolation on the top four rows. StyleGAN2-ADA-Linear [13] is the baseline model with data augmentation. Generator trained with Instance-perturbation and ours show more accurate inversion, smoother latent space, more diverse interpolation, and better FID and PPL. The small grey images visualize the difference between the two face images. The red numbers are mean of pixel-wise difference, we highlighted the difference score > 0.2 with red border.

200K images, and LSUN has more than 1M images. However, many scenarios may have a limited amount of training data, *e.g.* there are just 10 examples per artist in the Artistic-Faces dataset [11]. Therefore, Data-Efficient Generative Adversarial Networks (DE-GANs [12]) are important to many real-world practices and attract more and more attention. More analyses can be found in the survey [12].

Previous studies [14,15] demonstrate that severe discriminator overfitting leads to the training failure of DE-GANs. The less training data there is, the earlier the discriminator overfitting happens. Specifically, the discriminator is overly confiding in a few real samples and shows abnormal decision boundaries. Then, the gradients passing to generator are inaccurate, and the training instability problems occur [16]. To mitigate it, many technologies [17,18,19,20,21] have been proposed. Among them, data augmentation strategies [14,15,22,23] largely alleviate the training instability issue.

Currently, how to improve the diversity and quality of DE-GANs generated images becomes the major objective of the community. Contrastive learning (CL) [24,25], as a representative self-supervised learning algorithm, has been applied to different components of GANs, *e.g.* real or fake samples. Although existing methods demonstrate promising performance on DE-GANs, the motivation and principle behind different CL strategies are not well explored and justified. In this paper, we revisit the principle of different contrastive learning strategies and explore a novel perspective of understanding the connections between contrastive learning and DE-GANs' training.

In particular, we experimentally analyze the results of thee different contrastive learning strategies (Instance-real: Eq (6), Instance-fake: Eq (7), and Instance-perturbation: Eq (8)) on DE-GANs. Among them, Instance-perturbation contributes most and achieves significant gains in terms of FID, while others bring slight or even negative improvement (Table 1). To dig out how Instanceperturbation works, we analyze the latent space by GAN inversion and interpolation. As shown in Fig. 1 (top four rows), the training data of DE-GANs is sparse and discrete, and the generator tends to remember these samples, which leads to the discontinuity of the latent space. It can be observed that Instanceperturbation (the fifth row) can alleviate this problem, which achieves better inversion results and more continuous latent interpolation. Such results indicate that current DE-GANs, e.g. the most representative method StyleGAN2-ADA, are troubled by the discontinuities on latent space, and reasonable contrastive learning methods may help to solve this problem. Based on these findings, we propose a novel contrastive learning method termed FakeCLR for DE-GANs, which only applies Instance-perturbation and additional three exquisite strategies inspired by continuously changing fake images. These three strategies are Noiserelated Latent Augmentation, Diversity-aware Queue, and Forgetting Factor of Queue. Finally, the proposed FakeCLR achieves state-of-the-art performance and contributes to more continuous latent space. Overall, our contributions include:

- Illustrating the major bottleneck of DE-GANs. We explore and emphasize the bottleneck, discontinuous of the latent space, of existing DE-GAN methods.
- Exploring the connections between contrastive learning and DE-GANs. Through comprehensive experiments, we revisit three popular contrastive learning strategies that can be applied to DE-GANs, and identify that only Instanceperturbation brings the major improvement on generative performance.
- A new contrastive learning method for DE-GANs. We propose new contrastive learning with noise perturbation on fake images for DE-GANs called FakeCLR. Meanwhile, three technical innovations have been proposed for improving generative performance.
- Comprehensive experimental validations on various training settings. Besides data augmentation, the proposed FakeCLR achieves significant improvements on both limited-data generation and few-shot generation tasks.

2 Background

2.1 Data-Efficient Generative Adversarial Networks

Generative Adversarial Networks (GANs) [1] act a two-player adversarial game, where the generator G(z) is a distribution mapping function that transforms low-dimensional latent distribution p_z to target distribution p_g . And the discriminator D(x) evaluates the divergence between the generated distribution p_q and real distribution p_r . The generator and discriminator minimize and maximize the adversarial loss function. This min-max game can be expressed as:

$$\min_{\phi} \max_{\theta} f(\phi, \theta) = \mathbb{E}_{\mathbf{x} \sim p_r} \left[g_1 \left(D_{\theta}(\mathbf{x}) \right) \right] + \mathbb{E}_{\mathbf{z} \sim p_z} \left[g_2 \left(D_{\theta} \left(G_{\phi}(\mathbf{z}) \right) \right) \right], \tag{1}$$

where ϕ and θ are parameters of the generator G and discriminator D, respectively. g_1 and g_2 are different functions for various GANs. Data-Efficient GANs is a special case of GANs, which targets at obtaining a generator that can fit the real data distribution with limited, few-shot training samples. Formally, given the limited data distribution p_l (Generally, samples of p_l is the subset of p_r), we expect the generated distribution p_g obtained by training with p_l to be as close as possible to p_r . Concretely, loss functions of each network are formalized as:

$$\mathcal{L}_{D} = -\mathbb{E}_{\mathbf{x} \sim p_{l}} \left[g_{1} \left(D_{\theta}(\mathbf{x}) \right) \right] - \mathbb{E}_{\mathbf{z} \sim p_{z}} \left[g_{2} \left(D_{\theta} \left(G_{\phi}(\mathbf{z}) \right) \right) \right],$$

$$\mathcal{L}_{G} = \mathbb{E}_{\mathbf{z} \sim p_{z}} \left[g_{2} \left(D_{\theta} \left(G_{\phi}(\mathbf{z}) \right) \right) \right].$$
(2)

Previous studies attribute the degradation of DE-GANs to the overfitting of the discriminator. Many data augmentation, regularization, architectures, and pre-training techniques [12] have been proposed to mitigate this issue. (i) Data Augmentation [14,26,15,27,28] is a striking method for mitigating overfitting of the discriminator and orthogonal to other ongoing researches on training, architecture, and regularization. Popular augmentation strategies, such as Adaptive Data Augmentation (ADA), are employed as the essential complements in most DE-GANs. (ii) Regularization [29,30,13] is also a kind of popular technology for enhancing generalization in discriminator by introducing priors or extra supervision tasks. For instance, Tseng et al. [21] proposed an anchors-based regularization term to slow down the training of discriminator under limited data. (iii) Architecture is also the key point to improve the performance and stability of GANs. Some light-weight networks [31] have been employed in DE-GANs. Additionally, some model compression methods [32] have also been introduced in DE-GANs [20]. (iv) Pre-training [33,34,35,36,17,18,19] provides another solution to decrease the demand for data. Some studies argued that generator pre-training introduces some image-diversity [37] and latent-smooth [38] priors, and discriminator pre-training [37] provides accurate gradient for DE-GANs.

According to the scale of training data, data-efficient generation can be roughly divided into two tasks: few-shot generation and limited-data generation. Few-shot generation is a challenging task due to the employed tiny amount of data, for instance, 10-shot and 100-shot. Compared to few-shot generation, limited-data generation employs more data, such as 1K or 2K. Generally, the employed technologies, apart from data augmentation, are notably different between few-shot and limited-data generation. Our proposed FakeCLR is a universal technology achieving remarkable performance on both tasks.

2.2 Contrastive Learning

Contrastive learning [39,25,24] is a kind of self-supervised method to extract representations. Generally, given an image \mathbf{x} , two random views (query: \mathbf{x}_q and

key: \mathbf{x}_q^+) are created by different data augmentations (T_1 and T_2), defined as:

$$\mathbf{x}_q = T_1(x), \quad \mathbf{x}_q^+ = T_2(x). \tag{3}$$

Positive pair is defined as such query-key pair, between \mathbf{x}_q and \mathbf{x}_q^+ , from the same image. Negative pairs are defined as pairs from different images, *i.e.* between \mathbf{x}_q and $\{\mathbf{x}_{k_i}^-\}_{i=1}^N$, where $\mathbf{x}_{k_i}^- = T_2(\mathbf{x}_i)$ and $\mathbf{x} \neq \mathbf{x}_i$, $i = 1, \dots, N$. All views are passed through the feature extractor $F(\cdot)$ to acquire representation \mathbf{v} :

$$\mathbf{v}_{q} = F\left(\mathbf{x}_{q}\right), \quad \mathbf{v}_{q}^{+} = F\left(\mathbf{x}_{q}^{+}\right), \quad \mathbf{v}_{k_{i}}^{-} = F\left(\mathbf{x}_{k_{i}}^{-}\right), \quad i = 1 \dots N.$$
(4)

Contrastive learning aims to maximize the similarity of positive pairs and makes negative pairs dissimilar. Therefore, the InfoNCE loss [39] is designed as:

$$\mathcal{C}_{F(\cdot),\phi(\cdot)}\left(\mathbf{x}_{q},\mathbf{x}_{q}^{+},\left\{\mathbf{x}_{k_{i}}^{-}\right\}_{i=1}^{N}\right) = \\ -\log\frac{\exp\left(\phi\left(\mathbf{v}_{q}\right)^{T}\phi\left(\mathbf{v}_{q}^{+}\right)/\tau\right)}{\exp\left(\phi\left(\mathbf{v}_{q}\right)^{T}\phi\left(\mathbf{v}_{q}^{+}\right)/\tau\right) + \sum_{i=1}^{N}\exp\left(\phi\left(\mathbf{v}_{q}\right)^{T}\phi\left(\mathbf{v}_{k_{i}}^{-}\right)/\tau\right)},$$
(5)

where $\phi(\cdot)$ is the head network to project the representation to different spaces, and τ is the temperature coefficient. Generally, contrastive learning strategies are employed in discriminative models and have achieved convincing performance on representation leraning [25,24,40,41]. There are still a few works [42,13] to apply contrastive learning into unconditional generative tasks. Specifically, Jeong *et al.* [42] proposed Contrastive Discriminator (ContraD), a way of training discriminators of GANs using improved SimCLR [25]. Yang *et al.* [13] proposed InsGen, a instance discrimination realized by MoCo-v2 [43], achieving the state-of-the-art performance on limited-data generation. Different from them, our FakeCLR only applies contrastive learning to the fake images and adopts three advanced strategies designed specifically for generative tasks, acquiring remarkable performance on both few-shot and limited-data generation.

3 How Contrastive Learning Benefits DE-GANs?

Previous works tend to attribute the success of contrastive learning to extra discrimination tasks. For example, in [13], contrastive learning is considered as an instance discrimination task applying to both real and fake images by two additional heads $\phi^r(\cdot)$ and $\phi^f(\cdot)$ besides the original bi-classification head $\phi^d(\cdot)$. In this section, we explore the question of *How contrastive learning benefits Data-Efficient GANs?* by conducting comprehensive comparisons among different contrastive learning strategies.

3.1 Comparisons and Analyses on Contrastive Learning Strategies

Following existing study [13], three main contrastive learning strategies can be applied in Data-Efficient GANs, *i.e.* Instance-real, Instance-fake, and Instanceperturbation. Specifically, *Instance-real* introduces an independent real head 6



Fig. 2: Part of the ablation experiments of different contrastive learning strategies with various queue sizes and datasets.

 $\phi^r(\cdot)$ to conduct additional instance discrimination on real images, *i.e.*,

$$\mathcal{C}^{r} = \mathcal{C}_{d(\cdot),\phi^{r}(\cdot)}\left(T_{q}(\mathbf{x}_{q}), T_{k+}(\mathbf{x}_{q}), \{T_{k_{i}}(\mathbf{x}_{k_{i}})\}_{i=1}^{N}\right),\tag{6}$$

where T_q , T_{k^+} , and T_{k_i} are different augmentations, \mathbf{x}_q and \mathbf{x}_{k_i} are sampled from the real distribution. Similarly, *Instance-fake* introduces an independent fake head $\phi^f(\cdot)$ to conduct instance discrimination on fake images through contrastive learning and whose objective is formalized as:

$$\mathcal{C}^{f} = \mathcal{C}_{d(\cdot),\phi^{f}(\cdot)}\left(T_{q}(G(\mathbf{z}_{q})), T_{k^{+}}(G(\mathbf{z}_{q})), \{T_{k_{i}}(G(\mathbf{z}_{k_{i}}))\}_{i=1}^{N}\right),$$
(7)

where \mathbf{z}_q and \mathbf{z}_{k_i} are sampled from the latent distribution. Furthermore, *Instanceperturbation* introduces a noise perturbation strategy [44] to further improve the power of Instance-fake. The objective is formalized as:

$$\mathcal{C}_p^f = \mathcal{C}_{d(\cdot),\phi^f(\cdot)} \left(T_q(G(\mathbf{z}_q)), T_{k^+}(G(\mathbf{z}_q + \epsilon_q)), \{T_{k_i}(G(\mathbf{z}_{k_i}))\}_{i=1}^N \right),$$
(8)

where ϵ_q stands for the perturbation of latent augmentation, which is sampled from a Gaussian distribution with small variance.

In the following, we adopt StyleGAN2-ADA [15] and InsGen [13] as reference and conduct ablation studies⁴ to explore "*How Contrastive Learning Benefits Data-Efficient GANs?*". Related results are reported in Table 1 and Fig. 2.

Instance-real causes performance drop on small datasets. Although on large dataset such as FFHQ-5K, Instance-real⁵ is effective, it would cause performance drop on smaller datasets. As shown in Fig. 2, the performance with different queue size is consistently decreased on small datasets, *e.g.* FFHQ-100, FFHQ-1K, and FFHQ-2K. Such a degradation is much severer when the queue size is 100 or 300 on FFHQ-100 dataset.

Instance-fake improves performance marginally on small datasets. When adding Instance-fake⁵ strategy, compared to the baseline StyleGAN2-ADA-Linear, the best synthesis quality for different queue sizes is consistently improved in

⁴ In InsGen [13], ADA-Linear rather than ADA is used as augmentation T when the number of images in the dataset is less than 5K. For fairness, ADA-Linear, rather than ADA and InsGen, should be regraded as the baseline for comparison.

⁵ The objectives of discriminator are: i) Instance-real: $\mathcal{L}_D' = \mathcal{L}_D + \lambda^r \mathcal{C}^r$; ii) Instance-fake: $\mathcal{L}_D' = \mathcal{L}_D + \lambda^f \mathcal{C}^f$; iii) Instance-perturbation: $\mathcal{L}_D' = \mathcal{L}_D + \lambda^f \mathcal{C}_p^f$, Respectively.

Table 1: The whole ablation experiments of different contrastive learning strategies. We demonstrate the best FID metric for different queue sizes on various datasets. Results with * are best on four different queue sizes, which are better than those reported in [13].

256×256 Resolution	FFHQ-100	FFHQ-1K	FFHQ-2K	FFHQ-5K
StyleGAN2-ADA	85.8	21.29	15.39	10.96
StyleGAN2-ADA-Linear	82	19.86	13.01	9.39
Instance-real	107	20.68	14.15	8.50
Instance-fake	76.6	19.45	12.69	8.31
Instance-perturbation	43.35	17.46	11.29	8.01
InsGen [13]	53.93	19.58	11.92	-
InsGen*	45.75^{*}	18.21^{*}	11.47^{*}	7.83^{*}

Table 1, regardless of the number of training data. However, on small datasets such as FFHQ-100, FFHQ-1K, and FFHQ-2K, the increase is marginal, and the performance is even decreased when queue size is small (as shown in Fig. 2).

Instance-perturbation is the key point for DE-GANs. After using Instanceperturbation⁵, FID obtains significant improvement regardless of queue size and the number of training images. Specifically, compared to Instance-fake, the best FID for different queue size as shown in Table 1 acquires consistent improvement. Furthermore, Instance-perturbation also obtains FID improvements of -2.4, -0.75, and -0.18 relative to InsGen (*i.e.* employs instance-real at the same time) on FFHQ-100, FFHQ-1K, and FFHQ-2K datasets, respectively.

The optimal queue size is related to the diversity of negative samples. Following MoCo-v2 [24], negative samples of InsGen are stored in a queue to reduce the computational complexity. Two queues, fake instance queue and real instance queue, are adopted in InsGen. The size of them are defined as fqs and rqs respectively. Empirically, length of the queue tends to be the 5% number of the dataset in MoCo-v2, indicating rqs is related to the number of real data. According to our experiments (Fig. 2), we find that the optimal fqs in Instancefake are larger than the optimal rqs in Instance-real under the same setting. We argue that the optimal queue size is related to the diversity of negative samples in queues. In Instance-real, the diversity of negative samples is low with limited training data. Thus long rqs leads to the degradation of contrastive learning.

3.2 Instance-perturbation and Latent Space Continuity

The above analyses demonstrate that Instance-perturbation is the key point of contrastive learning in DE-GANs and better performance than InsGen can be obtained by using Instance-perturbation alone. Moreover, we argue that Instance-perturbation mitigates the discontinuity of latent space effectively. Intuitively, Instance-perturbation introduces latent similarity prior that images synthesized from the latent codes within a neighbourhood are close to each other, which makes the latent space continuous. Experimentally, Fig. 1 illustrates that

Matha d	PI	PL (z)	PPL (w)		
Method	Mean	Std. Dev	Mean	Std. Dev	
StyleGAN2-ADA-Linear	2163	1819	232	100.3	
Instance-fake	2399	2082	217	95.6	
Instance-perturbation	1729	1626	167	90.5	
FakeCLR	1411	1490	136	65.4	

Table 2: Perceptual path length (PPL) of z and w spaces on FFHQ-2K (256×256)

Instance-perturbation owns more continuous interpolation, more accurate inversion, better FID, and better PPL compared to StyleGAN2-LDA, StyleGAN2-APA [23], and other contrastive learning strategies on Obama dataset. Furthermore, Table 2 shows that PPL [45] of Instance-perturbation outperform other methods in Mean and Std. We also show some similar results of FFHQ-100 dataset in Sec. A.1 of the supplementary materials.

4 Methodology

As discussed above, real instance discrimination is not conducive to the DE-GANs' training. Hence, we only adopt Instance-perturbation, which is illustrated on the top part of Fig. 3. Furthermore, bottom part of Fig. 3 introduces three innovative strategies, that are Noise-related Latent Augmentation, Diversity-aware Queue, and Forgetting Factor of Queue, for improving the efficiency of contrastive learning on fake images. In summary, the proposed complete objective function⁶ of FakeCLR is optimized using:

$$\hat{\mathcal{L}}_D = \mathcal{L}_D + \lambda^f \hat{\mathcal{C}}_p^f, \quad \hat{\mathcal{L}}_G = \mathcal{L}_G + \lambda_G \mathcal{C}^f,$$
where $\hat{\mathcal{C}}_p^f = \hat{\mathcal{C}}_{d(\cdot),\phi^f(\cdot)} \left(T_q(G(\mathbf{z}_q)), T_{k+}(G(\mathbf{z}_q + \hat{\boldsymbol{\epsilon}}_q)), \{T_{k_i}(G(\mathbf{z}_{k_i})), \mathbf{m}_i\}_{i=1}^{\hat{N}} \right).$
(9)

Noise-related Latent Augmentation. From Sec. 3.2, noise perturbation is the key point for contrastive learning in DE-GANs, which alleviates the discontinuity of latent space effectively. From this perspective, noise perturbation introduces latent similarity prior to the discriminator, which provides latentcontinuity guidance for the generator. However, the original strategy holds the same perturbation radius (ϵ_q), which is not desirable due to the nonuniform latent distribution. More sensibly, a noise-related perturbation with different latent sampling z_q should be adopted. Generally, $\epsilon_q \propto |z_q|$ when $z_q \sim \mathcal{N}(0, 1)$. It can be interpreted as stronger similarity prior should be required in low-density regions of the latent distribution. In this paper, we adopt a simple implementation that is $\hat{\epsilon}_q = l_1 \cdot |z_q|$, where l_1 is a coefficient.

⁶ Followed InsGen [13], Instance-fake (\mathcal{C}^{f}) is added into the generator loss to improve the generated diversity.



Fig. 3: Overview of FakeCLR. Overall, the pipeline only adopting fake instance discrimination with latent augmentation is illustrated on the top of the figure, where the representation is learned from bi-classification loss (\mathcal{L}_D) and Instanceperturbation loss (\mathcal{C}_p^f) . Here, N is the number of negative samples (queue size), $T(\cdot)$ is the data augmentation, d is the backbone of discriminator, ϕ^d and ϕ^f are two independent heads. Furthermore, three strategies proposed in this paper are illustrated at the bottom of the figure.

Diversity-aware Queue. The number of negative samples is critical in contrastive learning. Short queue leads to inhibiting convergence [25], while long queue suffers from redundancy and inappropriate negative samples. Empirically, length of the queue is related to the size of dataset, such as 5% number of the dataset in MoCo-v2 [24]. However, this empirical hypothesis does not seem to hold up in GANs due to the continuously changing generated images over training iteration. As described in Sec. 3.1, the optimal queue size may be related to the diversity of samples in queue. Furthermore, the diversity of fake queues is indeed higher than that of real queues and decreases gradually accompanying with training due to the convergence and stability of generator. Therefore, a diversity-aware queue is required for contrastive learning in fake images. Generally, $N \propto DE(G(z))$, where $DE(\cdot)$ is a diversity estimator. In this paper, we also adopt a simple implementation that is $\hat{N} = l_2 \cdot t$ to reduce the computational complexity, where l_2 is a negative coefficient and t is the iteration of training. Therefore, queue size in FakeCLR decreases linearly with respect to the training.

Forgetting Factor of Queue. The choice of negative pairs has drawn much attention in contrastive learning. [46,47] demonstrate that hard negative samples capture desirable generalization properties and improve the performance of downstream task. Different from real images, fake images are updated with training iteration. As a result, varying levels of generated samples exist in the

9

10 Ziqiang Li, Chaoyue Wang, Heliang Zheng, Jing Zhang, and Bin Li

negative queue. However, vanilla contrastive learning considers them to have the same importance. Some iteration-based prior, *i.e.*, the negative samples generated by the current iteration should be given more attention, should be introduced. In this paper, we apply a forgetting factor to negative samples of the queue. Specifically, the contribution of previous samples should be smaller than current samples to contrastive loss. This makes the contrastive term focus more on recently generated samples. The iteration-based InfoNCE is formulated as:

$$\hat{\mathcal{C}}_{F(\cdot),\phi(\cdot)}\left(\mathbf{x}_{q},\mathbf{x}_{k}^{+},\left\{\mathbf{x}_{k_{i}}^{-},\mathbf{m}_{i}\right\}_{i=1}^{N}\right) = \exp\left(\phi\left(\mathbf{v}_{q}\right)^{T}\phi\left(\mathbf{v}_{k}^{+}\right)/\tau\right) \\ -\log\frac{\exp\left(\phi\left(\mathbf{v}_{q}\right)^{T}\phi\left(\mathbf{v}_{k}^{+}\right)/\tau\right)}{\exp\left(\phi\left(\mathbf{v}_{q}\right)^{T}\phi\left(\mathbf{v}_{k}^{+}\right)/\tau\right) + \sum_{i=1}^{N}\exp\left(\left(\phi\left(\mathbf{v}_{q}\right)^{T}\phi\left(\mathbf{v}_{k_{i}}^{-}\right)+\mathbf{m}_{i}\right)/\tau\right)},$$
(10)

where \mathbf{m}_i is the introduced forgetting factor to adjust the importance of different negative samples, as defined by:

$$\mathbf{m}_{i} = \frac{\exp\left(\hat{\mathbf{t}}_{i}/\tau_{\mathbf{m}}\right)}{\sum_{j=1}^{N} \exp\left(\hat{\mathbf{t}}_{j}/\tau_{\mathbf{m}}\right)}, \quad \text{where} \quad \hat{\mathbf{t}}_{i} = \frac{\mathbf{t}_{i} - \min(\mathbf{t})}{\max(\mathbf{t}) - \min(\mathbf{t})}, \qquad (11)$$
$$\mathbf{t} = \{\mathbf{t}_{1}, \mathbf{t}_{2}, \cdots, \mathbf{t}_{N}\}, \quad i = 1, 2, \cdots, N,$$

 $\tau_{\mathbf{m}}$ is a temperature coefficient that controls the distribution $\mathbf{m} := {\mathbf{m}_1, \dots, \mathbf{m}_N}$, and t_i means that the negative sample $\mathbf{v}_{k_i}^-$ is generated by the generator of the t_i -th iteration. The distribution of \mathbf{m} under different temperature $\tau_{\mathbf{m}}$ and queue size N has be shown in Sec. A.2 of supplementary materials. The proposed forgetting factor (\mathbf{m}) indeed improves the importance of current negative samples to contrastive term. The proof and discussion are in Sec. A.3 of supplementary materials. Furthermore, a PyTorch-like pseudocode of iteration-based contrastive learning has also be shown in Sec. A.4 of supplementary materials.

5 Experiments

We evaluate the advancement of our proposed FakeCLR on multiple datasets. After introducing implementation details, we present the comparison on FFHQ [45] dataset. Sec. 5.3 presents the comparison on few-shot datasets (Obama [14], Grumby Cat [14], Panda [14], and AnimalFace [50]). Furthermore, ablation studies showing the importance of different components are demonstrated in Sec. 5.4.

5.1 Implementation Details

Similar to InsGen [13], our FakeCLR also can be easily implemented on any GAN framework. In this paper, we take the state-of-the-art GANs model, StyleGAN2 [2], as an example to demonstrate how FakeCLR is implemented.

InsGen reuses the network structure and learning hyper-parameters of Style-GAN2 [2], the augmentation pipeline of ADA [15], and contrastive learning settings of MoCo-v2 [24] achieving state-of-the-art results with limited training

Table 3: Comparison with the state of the arts over FFHQ (256×256): Training with 100, 1K, 2K, and 5K samples. FID (lower is better) is reported as the evaluation metric. our method performs the best consistently. All methods adopt the StyleGAN2 architecture. Results with * are best on four different queue sizes, which are better than those reported in [13].

Method	Augmentation	FFHQ-100	FFHQ-1K	FFHQ-2k	FFHQ-5K
StyleGAN2 [2]	No	179	100.16	54.3	49.68
GenCo [48]	No	148	65.31	47.32	27.96
DISP $[49]$	Yes	-	-	21.06	-
LeCam-GAN [21]	Yes	-	21.70	-	-
ADA [15]	Yes	85.8	21.29	15.39	10.96
ADA-Linear [13]	Yes	82	19.86	13.01	9.39
APA [23]	Yes	65	18.89	16.90	8.38
InsGen [13]	Yes	53.93	19.58	11.92	-
$InsGen^*$	Yes	45.75^{*}	18.21^{*}	11.47^{*}	7.83^{*}
FakeCLR (ours)	Yes	42.56	15.92	9.90	7.25

data. Additionally, a linear hyper-parameter that controls the strength of augmentation (ADA-Linear) has also been proposed to replace the adaptive hyperparameter in ADA [15]. Therefore, we exactly reuse the above settings in InsGen for a fair comparison. Besides, some new hyper-parameter in FakeCLR are setted as $l_1 = 0.1$ and $\tau_{\mathbf{m}} = 0.01$. All experiments are conducted on two NVIDIA Tesla A100 GPUs with 64 batch size.

5.2 Results on FFHQ

FFHQ contains 70K high-resolution images of human faces. To reduce the computational complexity, we resize images to 256×256 . For the experiments of limited data, we collect a subset of training data by randomly sampling with 100, 1K, 2K, and 5K images. Regardless of the number of the training data, the Fréchet Inception Distance (FID) [51] metric is calculated between 50K fake images and all 70K images, evaluating the performance in terms of both reality and diversity. Moreover, the implementation of our FakeCLR is based on the official implementation of StyleGAN2-ADA and InsGen.

Table 3 shows the comparison with previous studies on FFHQ-100, FFHQ-1K, FFHQ-2K, and FFHQ-5k. We compare to most recent works, and our method achieves the new state-of-the-art consistently. Specifically, compared to the baseline method InsGen, our method improves the FID on FFHQ by 3.19, 2.29, 1.57, and 0.58 with 100, 1K, 2K, and 5K training images, respectively. Such results show the effectiveness of our proposed FakeCLR, especially for data-efficient scenarios. Some generated examples and the corresponding FID on FFHQ are presented on the top part of Fig. 4. More examples generated by InsGen [13] and our FakeCLR are presented in Sec. A.5 of the supplementary materials. Furthermore, we also analyze the computation cost in Sec. A.6 of



Fig. 4: Generated images and the corresponding FID under various datasets. All images are synthesized randomly without truncation.

the supplementary materials. Note that the three strategies adopted in Fake-CLR only introduce a fraction of the cost. More importantly, our FakeCLR achieves better results even with 8.8% less computational cost than InsGen [13]. Some qualitative analysis of nearest neighbor test as done in DiffAugment [14] in the LPIPS and pixel space, and comparisons on PPL and KID metrics on the FFHQ-100 dataset have also been presented in Sec. A.7 and Sec. A.8 of the supplementary materials, respectively.

5.3 Results on Few-shot Generation

We also evaluate the advancement of FakeCLR on few-shot generation that contains five datasets (*i.e.*, Obama, Grumpy Cat, Panada, AnimalFace Cat, and AnimalFace Dog) with 100, 100, 100, 160, and 389 training images, respectively. All models in this section are trained with the resolution of 256×256 , and the datasets can be found at the link. FID [51] is calculated between 50K fake images and the training images. Other settings are the same as experiments on FFHQ.

Table 4 demonstrates the quantitative comparison on few-shot generation. Top five baselines in Table 4 pre-train the model with large FFHQ-70K. It can be observed that FakeCLR can even achieve better FID by using only 100-400 training samples. The synthesis quality is substantially improved by our method on most datasets except for Grumpy Cat datasets. Specifically, we obtain an FID improvements over GenCo [48] by 16.3%, -9.9%, 11.3%, 14.7%, and 15.3% on five datasets, respectively. Compared to InsGen [13], the FID improvements are 16.9%, 11.1%, 14.5%, 20.2%, and 15.3%, respectively. Moreover, some qualitative results are shown on the second row of Fig. 4. More examples generated by InsGen [13] and FakeCLR can be found in Sec. A.9 of the supplementary. Furthermore, nearest neighbor test in the LPIPS and pixel space, and comparisons on PPL and KID metrics on the Obama dataset have also been presented in Sec. A.7 and Sec. A.8 of the supplementary, respectively.

Table 4: Comparison with previous works over 100-shot and AFHQ: Training with 100 (Obama, Grumpy Cat, and Panda), 160 (AFHQ Cat), and 389 (AFHQ Dog) samples. FID (lower is better) is reported as the evaluation metric. Our method performs the best in most datasets. All methods adopt the StyleGAN2 architecture except for FastGAN+DA.

Mathad	Massive	Pre-training	aining 100-shot		AnimalFace		
Method	Augmentation	w/ 70K images	Obama	Grumpy Cat	Panda	Cat	Dog
Scale/shift [33]	No	Yes	50.72	34.20	21.38	54.83	83.04
MineGAN [34]	No	Yes	50.63	34.54	14.84	54.45	93.03
TransferGAN [35]	No	Yes	48.73	34.06	23.20	52.61	82.38
TransferGAN + DA [14]	Yes	Yes	39.85	29.77	17.12	49.10	65.57
FreezeD [36]	No	Yes	41.87	31.22	17.95	47.70	70.46
StyleGAN2 [2]	No	No	80.20	48.90	34.27	71.71	130.19
DA [14]	Yes	No	46.87	27.08	12.06	42.44	58.85
FastGAN+DA [31]	Yes	No	41.05	26.65	10.03	35.11	50.66
ADA [15]	Yes	No	45.69	26.62	12.90	40.77	56.83
ADA-Linear [13]	Yes	No	34.41	25.67	10.23	36.30	54.39
SGLP [30]	Yes	No	45.37	26.52	-	-	-
LeCam-GAN [21]	Yes	No	33.16	24.93	10.16	34.18	54.88
GenCo [48]	Yes	No	32.21	17.79	9.49	30.89	49.63
InsGen [13]	Yes	No	32.42	22.01	9.85	33.01	44.93
FakeCLR (Ours)	Yes	No	26.95	19.56	8.42	26.34	42.02

Table 5: Ablation study of FakeCLR. The proposed strategies in FakeCLR all improve the generation over the baseline. Here, the first line is baseline StyleGAN2-ADA-Linear, and the other lines mean adding the proposed strategies.

\mathcal{C}_p^f	$\hat{\epsilon}_q$	$\hat{\mathcal{C}}_{d(\cdot),\phi^f(\cdot)}$	\hat{N}	FFHQ-100	FFHQ-1K	FFHQ-2K	FFHQ-5K
				82.00	19.86	13.01	9.39
\checkmark				43.35	17.46	11.29	8.01
\checkmark	\checkmark			44.10	17.09	10.65	7.89
\checkmark	\checkmark	\checkmark		41.62	16.05	10.40	7.48
\checkmark	\checkmark	\checkmark	\checkmark	42.56	15.92	9.90	7.25

5.4 Ablation Study

Component-level investigation. We investigate the effectiveness of each component on the FFHQ dataset. As shown in Table 5, all components are effective and can improve the performance Among them, $\hat{C}_{d(\cdot),\phi^{f}(\cdot)}$ achieves the most significant improvement. In addition, combining all components performs the best on FFHQ-1K, FFHQ-2K, and FFHQ-5K. On FFHQ-100, Diversity-aware Queue does not further improve the performance because the training data size is extremely small. Detailed ablations on FFHQ-2K are introduced as follows and more results can be found in Sec. A.10, A.11, and A.12 of the supplementary. **Noise-related Latent Augmentation**. Perturbation Radius (ϵ_q) is an important hyper-parameter, and we compare different ϵ_q to our Noise-related latent augmentation ($\hat{\epsilon}_q$) in Fig. 5 (a). The result shows that our strategy outperforms the best fixed strategy with a clear margin. Furthermore, negative prior that is $\epsilon_q \propto 1/|z_q|$ has been introduced in Negative Noise-related Latent Augmentation.



Fig. 5: Ablation studies of different components in strategies of FakeCLR. (a): Ablation study of perturbation radius (ϵ_q) in Noise-related Latent Augmentation. (b): Ablation study of batch size in Forgetting Factor of Queue. (c): Ablation study of temperature ($\tau_{\mathbf{m}}$) in Forgetting Factor of Queue.

The results demonstrate that negative prior severely compromise performance, which further supports our motivation of Noise-related Latent Augmentation. Forgetting Factor of Queue. The queue of negative samples is always a key factor in contrastive learning. An effective negative queue is supposed to be with diverse and up-to-date generated samples. Simply tuning batch size would cause a trade-off, *i.e.*, small batch size involves old samples, and large batch size reduces data diversity. As shown in Fig. 5 (b), 256 is the optimal batch size. And proposed Forgetting Factor of Queue provides a solution to break such trade-off, which suppresses old samples involved by small batch size and can achieve better performance. Note that for fair comparisons, we keep the training batch size constant and only adopt gradient accumulation to change the number of negative samples entering the queue for each iteration in contrastive learning (i.e.,batch size in Forgetting Factor of Queue). Fig. 5 (c) further shows the impact of Temperature $(\tau_{\mathbf{m}})$ in our method. As shown in Sec. A.2 of the supplementary materials, $\tau_{\mathbf{m}}$ controls the distribution of \mathbf{m} and smaller $\tau_{\mathbf{m}}$ indicate higher weights for samples in current iteration. The results show that proposed method outperform the original method, and $\tau_{\rm m} = 0.001$ achieves the best performance.

6 Conclusions

In this paper, we revisit the principle of different contrastive learning strategies in DE-GANs. According to experiments, we identify that latent similarity prior introduced by Instance-perturbation makes the main contribution to DE-GANs' performance. Furthermore, we also explore and emphasize a major bottleneck of existing DE-GAN methods, discontinuous of the latent space and demonstrate that Instance-perturbation can mitigate it effectively. Based on these, we propose a new contrastive learning method, FakeCLR, in DE-GANs, which acquires a new state of the art on both few-shot generation and limited-data generation.

Acknowledgements The work is partially supported by the National Natural Science Foundation of China under Grand No.U19B2044 and No.61836011.

References

- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in Neural Information Processing Systems. (2014) 2672–2680
- Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2020) 8110–8119
- Wang, C., Xu, C., Wang, C., Tao, D.: Perceptual adversarial networks for imageto-image transformation. IEEE Transactions on Image Processing 27(8) (2018) 4066–4079
- 4. Wenlong, Z., Yihao, L., Dong, C., Qiao, Y.: Ranksrgan: Generative adversarial networks with ranker for image super-resolution. IEEE Transactions on Pattern Analysis and Machine Intelligence (2021)
- Yang, Y., Wang, C., Liu, R., Zhang, L., Guo, X., Tao, D.: Self-augmented unpaired image dehazing via density and depth decomposition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). (June 2022) 2037–2046
- Qiu, J., Wang, X., Maybank, S.J., Tao, D.: World from blur. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2019) 8493–8504
- Yang, C., Wang, Z., Zhu, X., Huang, C., Shi, J., Lin, D.: Pose guided human video generation. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018) 201–216
- Tao, R., Li, Z., Tao, R., Li, B.: Resattr-gan: Unpaired deep residual attributes learning for multi-domain face image translation. IEEE Access 7 (2019) 132594– 132608
- Jiang, L., Zhang, C., Huang, M., Liu, C., Shi, J., Loy, C.C.: Tsit: A simple and versatile framework for image-to-image translation. In: European Conference on Computer Vision, Springer (2020) 206–222
- Qiu, J., Yang, Y., Wang, X., Tao, D.: Scene essence. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2021) 8322– 8333
- 11. Yaniv, J., Newman, Y., Shamir, A.: The face of art: landmark detection and geometric style in portraits. ACM Transactions on graphics (TOG) 38(4) (2019) 1–15
- 12. Li, Z., Wu, X., Xia, B., Zhang, J., Wang, C., Li, B.: A comprehensive survey on data-efficient gans in image generation. arXiv preprint arXiv:2204.08329 (2022)
- Yang, C., Shen, Y., Xu, Y., Zhou, B.: Data-efficient instance generation from instance discrimination. arXiv preprint arXiv:2106.04566 (2021)
- Zhao, S., Liu, Z., Lin, J., Zhu, J.Y., Han, S.: Differentiable augmentation for data-efficient gan training. Advances in Neural Information Processing Systems 33 (2020)
- Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., Aila, T.: Training generative adversarial networks with limited data. Advances in Neural Information Processing Systems 33 (2020)
- Arjovsky, M., Bottou, L.: Towards principled methods for training generative adversarial networks. arXiv preprint arXiv:1701.04862 (2017)
- Yang, C., Shen, Y., Zhang, Z., Xu, Y., Zhu, J., Wu, Z., Zhou, B.: One-shot generative domain adaptation. arXiv preprint arXiv:2111.09876 (2021)

- 16 Ziqiang Li, Chaoyue Wang, Heliang Zheng, Jing Zhang, and Bin Li
- Sauer, A., Chitta, K., Müller, J., Geiger, A.: Projected gans converge faster. Advances in Neural Information Processing Systems 34 (2021)
- Kumari, N., Zhang, R., Shechtman, E., Zhu, J.Y.: Ensembling off-the-shelf models for gan training. arXiv preprint arXiv:2112.09130 (2021)
- Chen, T., Cheng, Y., Gan, Z., Liu, J., Wang, Z.: Data-efficient gan training beyond (just) augmentations: A lottery ticket perspective. Advances in Neural Information Processing Systems 34 (2021)
- Tseng, H.Y., Jiang, L., Liu, C., Yang, M.H., Yang, W.: Regularizing generative adversarial networks under limited data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2021) 7921–7931
- 22. Tran, N.T., Tran, V.H., Nguyen, N.B., Nguyen, T.K., Cheung, N.M.: On data augmentation for gan training. IEEE Transactions on Image Processing **30** (2021) 1882–1897
- Jiang, L., Dai, B., Wu, W., Loy, C.C.: Deceive d: Adaptive pseudo augmentation for gan training with limited data. In: Thirty-Fifth Conference on Neural Information Processing Systems. (2021)
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2020) 9729–9738
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning, PMLR (2020) 1597–1607
- Zhao, Z., Zhang, Z., Chen, T., Singh, S., Zhang, H.: Image augmentations for gan training. arXiv preprint arXiv:2006.02595 (2020)
- Wang, C., Xu, C., Yao, X., Tao, D.: Evolutionary generative adversarial networks. IEEE Transactions on Evolutionary Computation 23(6) (2019) 921–934
- Tran, N.T., Tran, V.H., Nguyen, N.B., Nguyen, T.K., Cheung, N.M.: Towards good practices for data augmentation in gan training. arXiv preprint arXiv:2006.05338 (2020)
- Li, Z., Wu, X., Usman, M., Tao, R., Xia, P., Chen, H., Li, B.: A systematic survey of regularization and normalization in gans. arXiv preprint arXiv:2008.08930 (2020)
- Kong, C., Kim, J., Han, D., Kwak, N.: Smoothing the generative latent space with mixup-based distance learning. arXiv preprint arXiv:2111.11672 (2021)
- Liu, B., Zhu, Y., Song, K., Elgammal, A.: Towards faster and stabilized gan training for high-fidelity few-shot image synthesis. In: International Conference on Learning Representations. (2020)
- Zhang, Z., Chen, X., Chen, T., Wang, Z.: Efficient lottery ticket finding: Less data is more. In: International Conference on Machine Learning, PMLR (2021) 12380–12390
- Noguchi, A., Harada, T.: Image generation from small datasets via batch statistics adaptation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. (2019) 2750–2758
- Wang, Y., Gonzalez-Garcia, A., Berga, D., Herranz, L., Khan, F.S., Weijer, J.v.d.: Minegan: effective knowledge transfer from gans to target domains with few images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2020) 9332–9341
- Wang, Y., Wu, C., Herranz, L., van de Weijer, J., Gonzalez-Garcia, A., Raducanu, B.: Transferring gans: generating images from limited data. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018) 218–234
- Mo, S., Cho, M., Shin, J.: Freeze the discriminator: a simple baseline for fine-tuning gans. arXiv preprint arXiv:2002.10964 (2020)

- 37. Grigoryev, T., Voynov, A., Babenko, A.: When, why, and which pretrained GANs are useful? In: International Conference on Learning Representations. (2022)
- Ojha, U., Li, Y., Lu, J., Efros, A.A., Lee, Y.J., Shechtman, E., Zhang, R.: Fewshot image generation via cross-domain correspondence. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2021) 10743–10752
- Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
- Chen, X., He, K.: Exploring simple siamese representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2021) 15750–15758
- 41. Xiao, T., Wang, X., Efros, A.A., Darrell, T.: What should not be contrastive in contrastive learning. arXiv preprint arXiv:2008.05659 (2020)
- 42. Jeong, J., Shin, J.: Training gans with stronger augmentations via contrastive discriminator (2021)
- 43. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. arXiv preprint arXiv:2003.04297 (2020)
- Cheung, T.H., Yeung, D.Y.: Modals: Modality-agnostic automated data augmentation in the latent space. In: International Conference on Learning Representations. (2020)
- Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. (2019) 4401–4410
- Robinson, J., Chuang, C.Y., Sra, S., Jegelka, S.: Contrastive learning with hard negative samples. arXiv preprint arXiv:2010.04592 (2020)
- Kalantidis, Y., Sariyildiz, M.B., Pion, N., Weinzaepfel, P., Larlus, D.: Hard negative mixing for contrastive learning. Advances in Neural Information Processing Systems 33 (2020) 21798–21809
- Cui, K., Huang, J., Luo, Z., Zhang, G., Zhan, F., Lu, S.: Genco: Generative co-training for generative adversarial networks with limited data. arXiv preprint arXiv:2110.01254 (2021)
- Mangla, P., Kumari, N., Singh, M., Krishnamurthy, B., Balasubramanian, V.N.: Data instance prior (disp) in generative adversarial networks. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. (2022) 451–461
- 50. Si, Z., Zhu, S.C.: Learning hybrid image templates (hit) by information projection. IEEE Transactions on pattern analysis and machine intelligence 34(7) (2011) 1354– 1367
- 51. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems **30** (2017)