

A Closer Look at Few-shot Image Generation

Yunqing Zhao¹Henghui Ding²Houjing Huang³Ngai-Man Cheung^{1*}¹ Singapore University of Technology and Design (SUTD)² ETH Zürich³ ByteDance Ltd.

yunqing.zhao@mymail.sutd.edu.sg, ngaiman.cheung@sutd.edu.sg

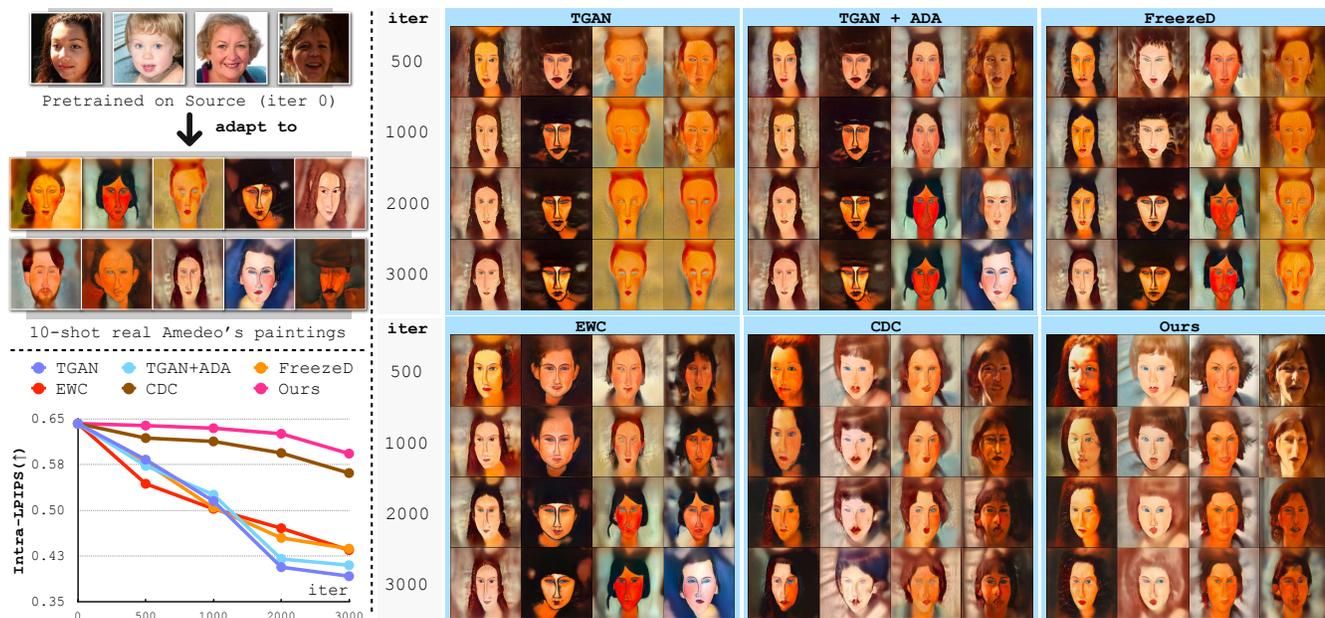


Figure 1. **Upper Left:** Transferring a GAN pretrained on a large-scale source domain (FFHQ [23]) to 10-shot Amedeo Modigliani’s paintings [53]. **Right:** During adaptation, we randomly select four fixed noise input and visualize the generated images with different existing methods. We observe that all methods including those with disproportionate focus on diversity preserving achieve similar quality after convergence (evaluated by our proposed realismness classifier in Sec. 4.2). Therefore, the better methods are those that can slow down diversity degradation (**Bottom Left**, evaluated by intra-LPIPS in Sec. 4.3). We perform a rigorous study and show quantitative results in Sec. 4.5.

Abstract

Modern GANs excel at generating high quality and diverse images. However, when transferring the pretrained GANs on small target data (e.g., 10-shot), the generator tends to replicate the training samples. Several methods have been proposed to address this few-shot image generation task, but there is a lack of effort to analyze them under a unified framework. As our first contribution, we propose a framework to analyze existing methods during the adaptation. Our analysis discovers that while some methods have disproportionate focus on diversity preserving which impede quality improvement, all methods achieve similar quality after convergence. Therefore, the better methods are those that can slow down diversity degradation. Furthermore, our analysis reveals that there is still plenty of room to further slow down diversity degradation.

Informed by our analysis and to slow down the diver-

sity degradation of the target generator during adaptation, our second contribution proposes to apply mutual information (MI) maximization to retain the source domain’s rich multi-level diversity information in the target domain generator. We propose to perform MI maximization by contrastive loss (CL), leverage the generator and discriminator as two feature encoders to extract different multi-level features for computing CL. We refer to our method as Dual Contrastive Learning (DCL). Extensive experiments on several public datasets show that, while leading to a slower diversity-degrading generator during adaptation, our proposed DCL brings visually pleasant quality and state-of-the-art quantitative performance.

1. Introduction

Powerful Generative Adversarial Networks (GANs) [4, 14, 24] have been built in recent years that can generate images with high fidelity and diversity [6, 44]. Unfortunately,

*Corresponding Author

these GANs often require large-scale datasets and computational expensive resources to achieve good performance. For example, StyleGAN [23] is trained on Flickr-Faces-HQ (FFHQ) [23], which contains 70,000 images, with almost 56 GPU days. When the dataset size is decreased, however, the generator often tends to replicate the training data [11].

Would it be possible to generate sufficiently diverse images, given only limited training data? For example, with 10-shot sketch style human faces [45], could we generate diverse face sketch paintings? This *few-shot image generation* task is important in many real-world applications with limited data, *e.g.*, artistic domains. It can also benefit some downstream tasks, *e.g.*, few-shot image classification [3].

1.1. A Closer Look at Few-shot Image Generation

To address this few-shot image generation task, instead of training from scratch [42, 51], recent literature focus on transfer learning [2, 35, 50] based ideas, *i.e.*, leveraging the prior knowledge of a GAN pretrained on a large-scale, diverse dataset of the source domain and adapting it to a small target domain, without access to the source data. The early method is based on fine-tuning [47]. In particular, starting from the pretrained generator G_s , the original GAN loss [14] is used to adapt the generator to the new domain:

$$\min_{G_t} \max_{D_t} \mathcal{L}_{adv} = \mathbb{E}_{x \sim p_{data}(x)} [\log D_t(x)] \quad (1)$$

$$+ \mathbb{E}_{z \sim p_z(z)} [\log (1 - D_t(G_t(z)))],$$

where z is sampled from a Gaussian noise distribution $p_z(z)$, $p_{data}(x)$ is the probability distribution of the real target domain data x , G_t and D_t are generator and discriminator of the target domain, and G_t is initialized by the weights of G_s . This GAN loss in Eqn. 1 forces G_t to capture the statistics of the target domain data, thereby to achieve both good quality (realisticness w.r.t. target domain data) and diversity, the criteria for a good generator.

However, for few-shot setup (*e.g.* only 10 target domain images), such approach is inadequate to achieve diverse target image generation as very limited samples are provided to define $p_{data}(x)$. Recognizing that, recent methods [28, 33] have focused disproportionately on improving diversity by preserving diversity of the source generator during the generator adaptation. In [28], Elastic Weight Consolidation (EWC) [25] is proposed to limit changes in some important weights to preserve diversity. In [33], an additional Cross-domain Correspondence (CDC) loss is introduced to preserve the sample-wise distance information of source to maintain diversity, and the whole model is trained via a multi-task loss with the diversity loss \mathcal{L}_{dist} as an auxiliary task to regularize the main GAN task with loss \mathcal{L}_{adv} :

$$\min_{G_t} \max_{D_t} \mathcal{L}_{adv} + \mathcal{L}_{dist}. \quad (2)$$

In [33], a patch discriminator [20, 59] is also used to further improve the performance in \mathcal{L}_{adv} . Details of \mathcal{L}_{dist} in [33].

Diversity preserving methods [28, 33] have demonstrated impressive results based on Fréchet Inception Distance (FID) [18] which measures the quality and diversity of the generated samples simultaneously. However, second thoughts about these methods reveal some questions:

- With disproportionate focus on diversity preserving in recent works [28, 33], will quality of the generated samples be compromised? For example, in Eqn. 2, \mathcal{L}_{adv} is responsible for quality improvement during adaptation, but \mathcal{L}_{dist} may compete with \mathcal{L}_{adv} as it has been observed in multi-task learning [12, 39]. We note that this has not been analyzed thoroughly.
- With recent works' strong focus on diversity preserving [28, 33], will there still be room to further improve via diversity preserving? How could we know when the gain of diversity preserving approaches become saturated (without excessive trial and error)?

1.2. Our Contributions

In this paper, we take the first step to address these research gaps for few-shot image generation. Specifically, as our first contribution, we propose to independently analyze the quality and diversity during the adaptation. Using this analysis framework, we obtain insightful information on quality/diversity progression. In particular, on one hand, it is true that strong diversity preserving methods such as [33] indeed impede the progress of quality improvement. On the other hand, interestingly, we observe that these methods can still reach high quality rather quickly, and after quality converges they have no worse quality compared to other methods such as [47] which uses simple GAN loss (Eqn. 1). Therefore, methods with disproportionate focus on preserving diversity [33] stand out from the rests as they can produce slow diversity-degrading generators, maintaining good diversity of generated images when their quality reaches the convergence. Furthermore, our analysis reveals that there is still plenty of room to further slow down the diversity degradation across several source \rightarrow target domain setups.

Informed by our analysis, our second contribution is to propose a novel strong regularization to take a further step in slowing down the diversity degradation, with the understanding that it is unlikely to compromise quality as observed in our analysis. Our proposed regularization is based on the observation that rich diversity exists in the source images at different semantic levels: diversity in middle levels such as hair style, face shape, and that in high levels such as facial expression (smile, grin, concentration). However, such source diversity can be easily ignored in the images produced by target domain generators. Therefore, to preserve source diversity information, we propose to maximize the mutual information (MI) between the source/target image features originated from the same latent code, via contrastive loss [34] (CL). To compute CL, we leverage the generator and discriminator as two feature encoders to extract image features

at multiple feature scales, such that we can preserve diversity at various levels. By combining the two feature encoders (generator and discriminator), we gain additional feature diversity. We show that our proposed Dual Contrastive Learning (DCL) outperforms the previous work in slowing down the diversity degradation without compromising the image quality on the target domain, and hence achieving the state-of-the-art performance.

2. Related Work

Conventional few-shot learning [10, 13, 38] aims at learning a discriminative classifier [1, 16, 29, 38, 40]. In contrast, generative few-shot learning [28, 33, 49, 52] often follows a transfer learning [35] pipeline to adapt a pretrained GAN on a small target domain, without access to the source data. Specifically, our analysis and comparison in this paper mainly focus on the following recent baseline models:

- Transferring GAN [47] (TGAN): directly fine-tune all parameters of both the generator and the discriminator on the target domain;
- Adaptive Data Augmentation (ADA) [22]: apply data augmentation [41, 57, 58] during adaptation which does not leak to the generator;
- BSA [32]: only update the learnable scale and shift parameters of the generator during adaptation;
- FreezeD [31]: freeze a few high-resolution layers of the discriminator, during the adaptation process;
- MineGAN [46]: use additional modules between the noise input and the generator. It aims at matching the target distribution before input to the generator;
- Elastic Weight Consolidation (EWC) [28]: apply EWC loss to regularize GAN, preventing the important weights from drastic changes during adaptation;
- Cross-domain Correspondence (CDC) [33]: preserve the distance between different instances in the source.

We firstly take a rigorous study of the above methods under the same setup. Then, motivated by our findings, we aim to address this few-shot image generation task.

Our proposed method is related to contrastive learning [5, 8, 15, 17, 34, 55]. Contrastive learning for unsupervised instance discrimination aims at learning invariant embeddings with *different* transformation functions of the *same* object [17, 21, 30], which can benefit the downstream tasks. To slow down the diversity degradation during few-shot GAN adaptation, we present a simple and novel method via contrastive loss [34, 48]: we hypothesize that, the same noise input could be mapped to fake images in the source and target domains with shared semantic information [37]. In experiments, we demonstrate convincing and competitive results of our approach for few-shot GAN adaptation, with sufficiently high quality and diversity.

3. Preliminary

Given a GAN pretrained on the source domain, we denote the generator as G_s . In the adaptation stage, the target generator G_t and the discriminator D_t can be obtained by fine-tuning the pretrained GAN on the target domain via Eqn. 1. We follow the settings of the prior work [33, 46, 47]: during adaptation, there is *no* access to the abundant source data, and we can only leverage the pretrained GAN and target data for transfer learning.

For the target data size, some recent work [22, 47] applies more than 1,000 images in adaptation, and show satisfied results. In contrast, we focus on the challenging case: fine-tune the pretrained GAN with only few-shot (*e.g.*, 10-shot) data.

4. Revisit Few-Shot GAN Adaptation

4.1. Motivation

While a few recent works [28, 33] propose different ideas, including to inherit the diversity during adaptation, to our knowledge, no effort attends to analyze the underlying mechanism of this problem. The specific issues concerning this study are as follows: will quality be compromised due to the disproportionate focus on diversity preserving? will there still be room to further improve via diversity preserving? We study these concerns from a novel perspective: compare different methods by visualizing and quantifying the few-shot adaptation process. In particular, we decouple their performance by evaluating the quality and diversity on the target domain separately, under a unified framework.

4.2. Binary Classification for Quality Evaluation

The probability output of a classifier indicates the confidence of how likely an input sample belongs to a specific class [7, 19]. Therefore, we employ the probability output of a binary classifier to assess to what extent the generated images belong to the *target domain*. In particular, we train a convolutional network C on two sets of real images (from source and target, excluded during adaptation). Then, we apply C to the synthetic images from the adapted generator G_t during adaptation. The soft output of C are p_t , predicted probability of input belonging to the target domain, and $(1 - p_t)$, that of not belonging to the target domain. Therefore, we take p_t as an assessment of the realisticness of the synthetic images on the target domain:

$$p_t = \mathbb{E}_{z \sim p_z(z)}[C(G_t(z))]. \quad (3)$$

z is a batch of random noise input (size=1,000) fixed during adaptation, which is a *proxy* to indicate the quality and realisticness evolution process of different methods.

4.3. Intra-Cluster Diversity Evaluation

To evaluate the diversity, we introduce a Learned Perceptual Image Patch Similarity (LPIPS) [56] based metric, intra-cluster LPIPS (*intra*-LPIPS) [33], to identify the similarity

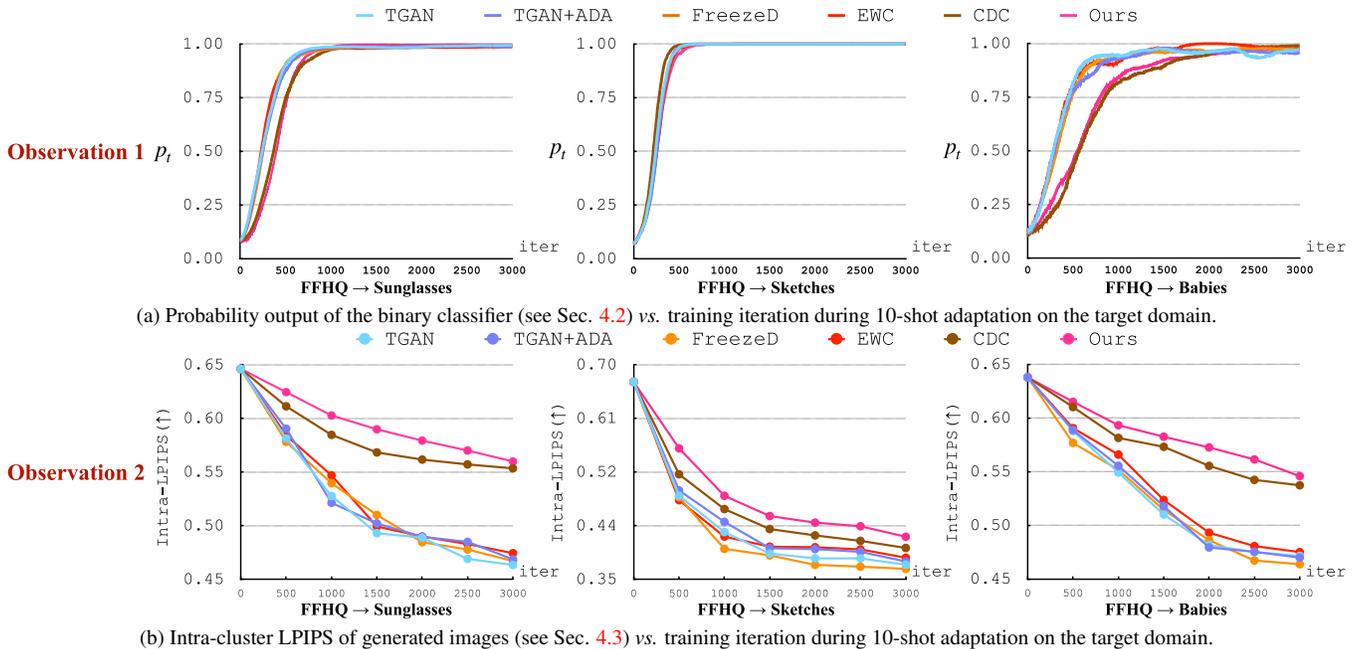


Figure 2. Analysis of existing few-shot GAN adaptation methods under a unified setup (see Sec. 4.4). **Observation 1:** While methods such as CDC [33] have disproportionate focus on diversity preserving which impedes quality improvement, all methods we have evaluated achieve *similar* quality/realisticness after convergence. **Observation 2:** Different methods exhibit substantially *dissimilar* diversity degrading rates. Combined the results in Figure 2a and 2b: since the achieved quality is similar, the better methods are those that can slow down diversity degradation. We summarize our findings in Sec. 4.5, which motivate our proposed method in Sec. 5. Best viewed in color.

between the generated images and the few-shot target samples during adaptation. The *standard* LPIPS evaluates the perceptual distance between two images, and it is empirically shown to align with human judgements [56]. Intra-LPIPS is a variation of the standard LPIPS: We firstly generate abundant (size=1,000) images, then assign each of them to one of the M target samples (with lowest standard LPIPS), and form M clusters. The intra-LPIPS is obtained by computing the average standard LPIPS for random paired images within each cluster, then average over M clusters. We provide the pseudo-code in Supp.

Ideally, we hypothesize that the highest diversity knowledge is achieved by the generator pretrained on the large source domain, and it will degrade during the few-shot adaptation. In the worst case, the adapted generator simply replicates the target images, and intra-LPIPS will be zero. To justify our conjecture, we sample the target generator at different adaptation iterations, then apply intra-LPIPS to assess the diversity of the generated images.

4.4. Experiment settings

Basic setup: For all methods: We follow the previous work [28, 33] to use StyleGAN-V2 [24] with default hyperparameters. We use AlexNet [27] as the binary classifier. We fine-tune the pretrained GAN on the target domain with batch size 4 and 3,000 iterations, on a Tesla V100 GPU.

Source \rightarrow target adaptation: Training the binary clas-

sifier requires rich data on both source and target domains. Therefore, we transfer the model pretrained on FFHQ [23] (70,000 images) to 10-shot samples from target domains that contain abundant data originally: Sketches, Babies, and Sunglasses (roughly 300, 2500, 2700 images, respectively). To obtain unbiased classifiers, we form each dataset by keeping the source and target training data balanced.

Evaluation: The baseline methods for evaluation are introduced in Sec. 2. In this section, we do not include BSA [32] and MineGAN [46]: BSA applies a different GAN architecture (BigGAN [4]), and MineGAN employs additional training modules during adaptation, and it is time-consuming for training. We include the results of these two models in Sec. 6. To compute the intra-LPIPS, we follow [33] to use the base implementation from [56]¹.

4.5. Empirical Analysis Results

The analysis results are shown in Figure 1 and Figure 2. Generally, our observations can be stated as follows:

- **Image quality/realisticness (Observation 1):** As shown in Figure 2a, strong diversity preserving methods such as CDC [33] (and ours, to be discussed in Sec. 5) indeed impede the progress of quality improvement, most noticeable in FFHQ \rightarrow Babies. This is because additional regularization for diversity preserving (e.g.

¹<https://github.com/richzhang/PerceptualSimilarity>

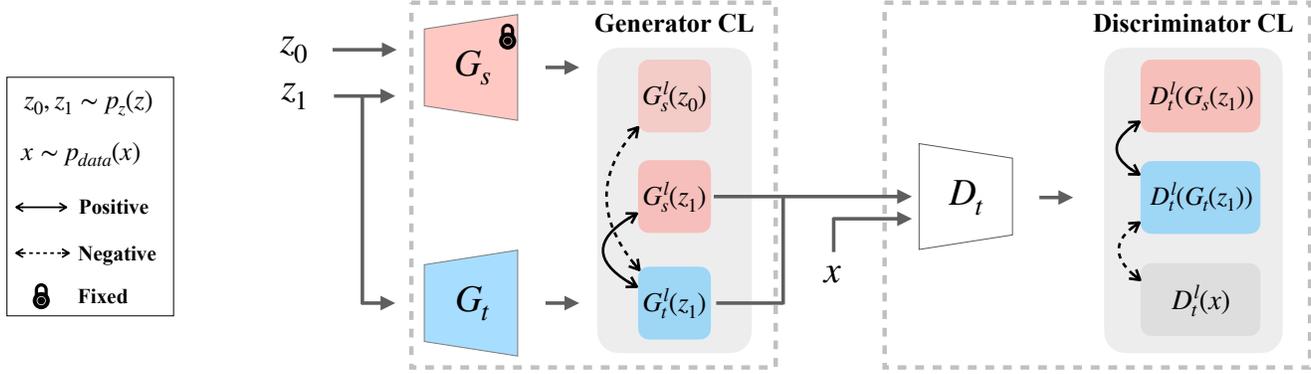


Figure 3. Our proposed Dual Contrastive Learning (DCL) for few-shot GAN adaptation. To slow down diversity degradation during adaptation, in both Generator CL (Eqn. 5) and Discriminator CL (Eqn. 7), we maximize the similarity between positive pairs generated using the same noise input, and push away those that are negative pairs. See details in Sec. 5.1.

\mathcal{L}_{dist} in [33], see Eqn. 2) competes with main GAN loss \mathcal{L}_{adv} which is responsible for quality improvement. Note that such observation is consistent with findings in multi-task learning [12, 39]. On the other hand, interestingly, we observe that these methods can still reach high quality on the target domain rather quickly, and they achieve similar quality as other methods after quality converges. We conjecture that, due to the informative prior knowledge on the source domain, it is not difficult to obtain adequate realism when transferring to the small target domain.

- **Image diversity (Observation 2):** In Figure 1 and Figure 2b, we demonstrate the qualitative and quantitative results of the diversity change during few-shot adaptation. While different methods can achieve similar quality on the target domain, their diversity-degrading rates vary drastically. For example, in Figure 1, TGAN [47] immediately loses the diversity and the generator replicates the target samples at the early adaptation stage. Since the achieved quality is similar for all these methods, the better methods are those that do not suffer from rapid diversity degradation.

As shown in Figure 2b, for all existing methods, the loss of diversity is inevitable in adaptation. On the other hand, since all methods achieve similar realism after the quality converges, the better generators are those that can slow down the diversity degradation, and the worse ones are those that lose diversity rapidly before the convergence of the quality improvement. Recalling the concerns we raise in Sec. 4.1: we show that, there is still plenty of room to further reduce the rate of diversity degradation, besides, the gain of diversity preservation becomes saturated only until the rate of diversity degradation is much reduced. These findings motivate our proposed method in the next section to further slow down diversity degradation.

5. Dual Contrastive Learning

5.1. Overall Framework

Rich diversity exists in the source images $G_s(z)$ at different semantic levels: middle levels such as different hair style, high levels such as different facial expression. To preserve source images’ diversity information in the target domain generator, we propose to maximize the mutual information (MI) between the source/target image features originated from the same input noise. In particular, for an input noise z_i , we seek to maximize:

$$\text{MI}(\pi^l(G_t(z_i)); \pi^l(G_s(z_i))), \quad (4)$$

where $G_t(z_i), G_s(z_i)$ are generated images by the target / source generators respectively, and $\pi^l(\cdot)$ is a feature encoder to extract l -th level features (to be further discussed). As G_s is fixed during adaptation, the MI maximization enables G_t to learn to generate images that include source images’ diversity at various semantic levels (with the use of multiple l during adaptation). Furthermore, in GAN, we have G_t and D_t , and we take advantage both of them to use as the feature encoders $\pi^l(\cdot)$ to extract different features. As directly maximizing MI is challenging [36], we apply contrastive learning [34] to solve Eqn. 4.

Specifically, we propose Dual Contrastive Learning (DCL) for few-shot GAN adaptation, as Figure 3. There are several goals for DCL: i) Maximize the MI between generated images on the source/target domain originated from the same noise input; ii) push away the generated images on the source and target domain that use different noise input; iii) push away the generated target images and the real target images, to prevent collapsing to the few-shot target set. To achieve these goals, we let the generating and the discriminating views using the same noise input, on source and target, as the positive pair, and maximize the agreement between them. Concretely, DCL includes two parts.

Generator CL: Given a batch of noise input $\{z_0, z_1, \dots, z_{N-1}\}$, we obtain images $\{G_s(z_0), G_s(z_1), \dots, G_s(z_{N-1})\}$

and $\{G_t(z_0), G_t(z_1), \dots, G_t(z_{N-1})\}$, generated on the source and the target domain, respectively. Considering an *anchor* image $G_t(z_i)$, we optimize the following object:

$$\mathcal{L}_{CL_1} = -\log \frac{f(G_t^l(z_i), G_s^l(z_i))}{\sum_{j=0}^{N-1} f(G_t^l(z_i), G_s^l(z_j))} \quad (5)$$

which is an N -Way categorical cross-entropy loss to classify the positive pair $(G_t^l(z_i), G_s^l(z_i))$ at l -th layer correctly. $\frac{f(G_t^l(z_i), G_s^l(z_i))}{\sum_{j=0}^{N-1} f(G_t^l(z_i), G_s^l(z_j))}$ is the prediction probability and

$$f(G_t^l(z_i), G_s^l(z_i)) = \exp(\text{CosSim}(G_t^l(z_i), G_s^l(z_i))/\tau) \quad (6)$$

for $i \in \{0, 1, \dots, N-1\}$ is the exponential of the cosine similarity between l -th layer features of the generated images on source and target, scaled by a hyperparameter temperature $\tau = 0.07$ and passed as logits [17, 37].

Discriminator CL: We focus on the view of the discriminator D_t . Given the generated images by the source and target generator and real data x as input to D_t , we maximize the agreement between the discriminating features of the generated images on source and target, using the same noise z_i . To prevent replicating the target data, we regularize the training process by pushing away discriminating features of the generated target images and the real target data at different scales. We optimize the following object:

$$\mathcal{L}_{CL_2} = -\log \frac{f(D_t^l(G_t(z_i)), D_t^l(G_s(z_i)))}{f(D_t^l(G_t(z_i)), D_t^l(G_s(z_i))) + \Delta}, \quad (7)$$

where $\Delta = \sum_{j=1}^M f(D_t^l(G_t(z_i)), D_t^l(x_j))$ and M is the number of real target images. The final objective of DCL in our work is simply:

$$\min_G \max_D \mathcal{L}_{adv} + \lambda_1 \mathcal{L}_{CL_1} + \lambda_2 \mathcal{L}_{CL_2}. \quad (8)$$

In practice, we find $\lambda_1 = 2$ and $\lambda_2 = 0.5$ work well. In each iteration, we randomly select different layers of G_t and D_t to perform DCL with multi-level features.

5.2. Design Choice

The design choice of DCL is intuitive: DCL takes the advantage of $G_s=G_t$ before adaptation. Our idea is to re-use the highest diversity knowledge from the source as the strong regularization to slow down the *inevitable* diversity degradation, before the quality improvement converges. The Generator CL and Discriminator CL are similar in idea, while they constraint different parts for GAN adaptation: Generator CL requires the adapted generator G_t to generate features at different scales that are similar to that of the source generator (with the same noise input). Differently, Discriminator CL regularizes the adversarial training: while fitting to the target domain, the generated target images are encouraged to be pushed away from the real data at different feature scales.

	FFHQ \mapsto Babies	FFHQ \mapsto Sunglasses	FFHQ \mapsto Sketches
TGAN [47]	104.79 \pm 0.03	55.61 \pm 0.04	53.42 \pm 0.02
TGAN+ADA [22]	102.58 \pm 0.12	53.64 \pm 0.08	66.99 \pm 0.01
BSA [32]	140.34 \pm 0.01	76.12 \pm 0.01	69.32 \pm 0.02
FreezeD [31]	110.92 \pm 0.02	51.29 \pm 0.05	46.54 \pm 0.01
MineGAN [46]	98.23 \pm 0.03	68.91 \pm 0.03	64.34 \pm 0.02
EWC [28]	87.41 \pm 0.02	59.73 \pm 0.04	71.25 \pm 0.01
CDC [33]	74.39 \pm 0.03	42.13 \pm 0.04	45.67 \pm 0.02
DCL (Ours)	52.56 \pm 0.02	38.01 \pm 0.01	37.90 \pm 0.02

Table 1. For target domains that contain rich data, we follow [33] to compare FID \downarrow [18] with different adaptation setups. We firstly generate 5,000 fake images using the adapted generator, then compare them with the real target data (excluded during adaptation). The standard derivations is computed by 5 different runs.

	FFHQ \mapsto Otto's Paintings	Church \mapsto Haunted House	Cars \mapsto Abandoned Cars
TGAN [47]	0.51 \pm 0.02	0.52 \pm 0.04	0.46 \pm 0.03
TGAN+ADA [22]	0.54 \pm 0.02	0.57 \pm 0.03	0.48 \pm 0.04
BSA [32]	0.46 \pm 0.02	0.43 \pm 0.02	0.41 \pm 0.03
FreezeD [31]	0.54 \pm 0.03	0.45 \pm 0.02	0.50 \pm 0.05
MineGAN [46]	0.53 \pm 0.04	0.44 \pm 0.06	0.49 \pm 0.02
EWC [28]	0.56 \pm 0.03	0.58 \pm 0.06	0.43 \pm 0.02
CDC [33]	0.63 \pm 0.03	0.60 \pm 0.04	0.52 \pm 0.04
DCL (Ours)	0.66 \pm 0.02	0.63 \pm 0.01	0.53 \pm 0.02

Table 2. For target domains containing only 10-shot data, we evaluate the diversity of the adapted generator. We firstly generate 5,000 fake images, then we compute intra-cluster LPIPS \uparrow with few-shot target samples. The standard derivation is computed across 10-shot clusters (see details in Sec. 4.3).

Under mild assumptions, DCL maximizes the MI (Eqn. 4) between the generated samples using the same input noise, on source and target, *e.g.*, for \mathcal{L}_{CL_1} : $\text{MI}(G_t^l(z_i); G_s^l(z_i)) \geq \log[N] - \mathcal{L}_{CL_1}$ [34]. In the next, we show the effectiveness of DCL in experiments.

6. Experiments

6.1. Implementation Details

In this section, we discuss our main experiments. Additional analysis and discussion are included in Supp.

Basic setups: We mostly follow the experiment setups as Ojha *et al.* [33], including the non-saturating GAN loss \mathcal{L}_{adv} in Eqn. 2, the same pretrained GAN and 10-shot target samples, with no access to the source data. We employ StyleGAN-V2 [24] as the GAN architecture for pretraining and few-shot adaptation, with an image/patch level discriminator [33]. We fine-tune with batch size 4 on a Tesla V100 GPU. The baseline methods are introduced in Sec. 2.

Datasets: We use models pretrained on different source domains: i) FFHQ [23] ii) LSUN Church and iii) LSUN Cars [54]. We transfer the pretrained model to the following target domains: i) Sketches [45] ii) Face paintings by Amedeo and Otto Dix [53] iii) FFHQ-Babies iv) FFHQ-Sunglasses v) Haunted houses vi) Wrecked cars. Images for training and evaluation are interpolated to the resolution of 256×256 .

²Our implementation is based on <https://github.com/mseitzer/pytorch-fid>



Figure 4. **Left:** Transferring a GAN pretrained on FFHQ to 10-shot target samples. We fix the noise input by each column to observe the relationship of generated images before and after adaptation. **Mid:** We observe that, most of existing methods lose diversity quickly before the quality improvement converges, and tend to replicate the training data. Our method, in contrast, slows down the loss of diversity and preserves more details. For example: In **red** frames (**Upper**), the hair style and hat are better preserved. In **pink** frames (**Bottom**), the smile teeth are well inherited from the source domain. We also outperform others in quantitative evaluation (**Right**). See details in Sec. 6.2.

6.2. Comparison with State-of-the-art Methods

Qualitative results. In Figure 4, we visualize the generated images with different methods after adaptation on few-shot target samples. In Figure 5, we show more results with different source \rightarrow target adaptation setups. Regularized by DCL, the adapted target generator needs to retain the connection to the generated images on source, a slow diversity degradation is thus achieved. We show that, the refined details, *e.g.*, the smile teeth, hairline or structure appearance are

well preserved, while the style and texture features are fitted to the target domain. Compared with recent state-of-the-art methods, these semantic meaningful features in the source domain are well inherited by DCL.

Quantitative comparison. We mainly focus on two metrics to evaluate our method. i) For datasets which contain a lot of real data, *e.g.*, FFHQ-Babies, FFHQ-Sunglasses, we apply the widely used Fréchet Inception Distance (FID) [18] to evaluate the generated fake images. ii) For those datasets

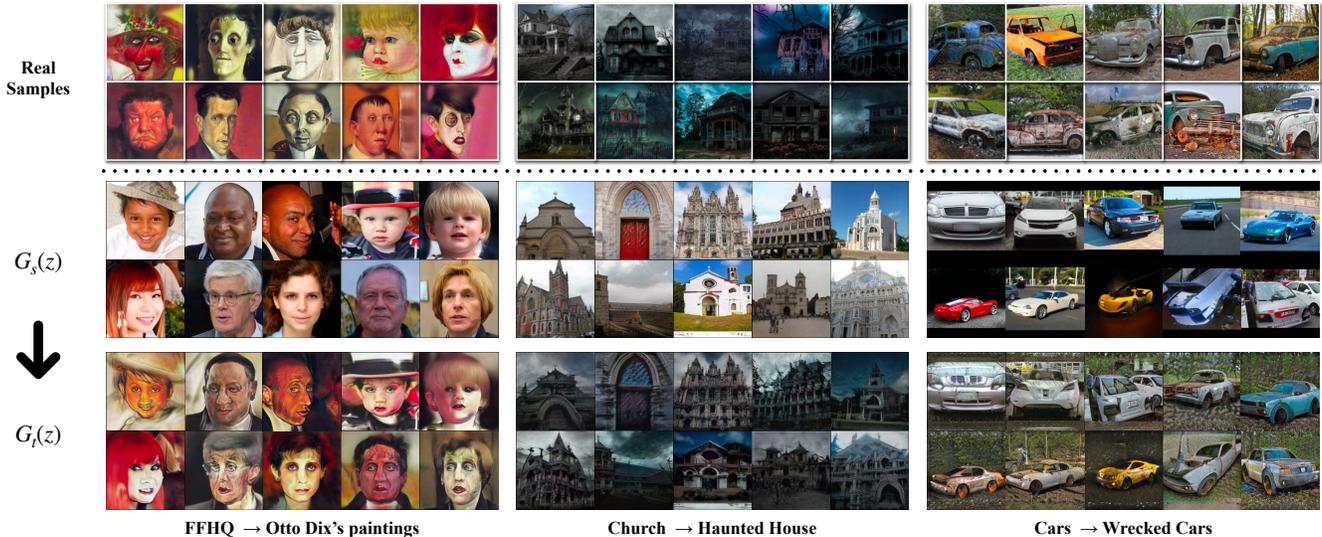


Figure 5. Generated images with different source \rightarrow target adaptation setups. We show that, the target generator after adaptation with our proposed method can generate visually pleasant images *without much loss of diversity*. The facial expressions of the characters, the structure appearance of the building, and the form/color of vehicles are well preserved and corresponded. See quantitative comparison in Table 2.

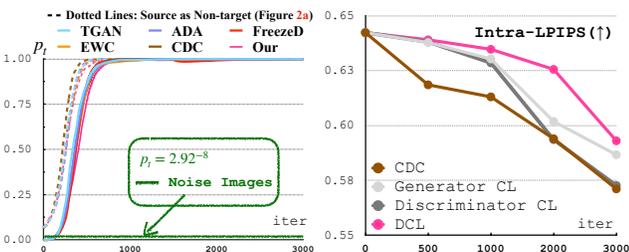


Figure 6. **Left:** Ablation of the realisticness classifier. We show that 1): Using alternative datasets as non-target to train the classifier give similar results. 2): Moreover, if we input randomly generated noise images, the classifier will give low confidence of being target. **Right:** Ablation study of DCL, with the same setup as Figure 1.

which contain only few-shot real data, FID becomes tricky and unstable, since it summarizes the quality and diversity to a single score. Therefore, we use intra-LPIPS (see Sec. 4.3) to measure the diversity of generated images. The better diversity comes with a higher score. As shown in Table 1 and Table 2, our method outperforms the baseline methods. This indicates that the target generator we obtain can cover a wide range of modes, and the loss of diversity is further reduced.

6.3. Analysis

Effect of our methods. The goal of our proposed binary classifier is to detect if an input image is “non-target” or “target”. In Figure 6 (left), we replace the source with equally sampled images from {FFHQ, ImageNet [9], CUB [43], Cars [26]} as “non-target” and Sketches as “target”, and we observe the similar results, compared to Figure 2a. We also show that both Generator CL and Discriminator CL can slow down the diversity degradation, and the better results are achieved when combined together. This further confirms our observations in Sec. 4: the slower diversity degradation

leads to the better target generator for few-shot adaptation.

Effect of target data size. In this work, we mainly focus on 10-shot adaptation, similar to the prior literature [28, 33]. We further perform ablation study on 5-shot and 1-shot adaptation, use the same setup as Sec. 4 and Sec. 6. As results in Supp., we have the similar qualitative and quantitative observations, which confirms our findings stated in Sec. 4.5.

7. Conclusion

Focusing on few-shot image generation, our first contribution is to analyze existing few-shot image generation methods in a unified framework to gain insights on quality/diversity progress during adaptation. Of surprise is that the achieved quality is similar for different existing methods despite some of them have disproportionate focus on preserving diversity. Of interest is that the different rates of diversity degradation is the main factor for different performances. Informed by our analysis, as our second contribution we propose a mutual information based method to slow down degradation of diversity during adaptation. We connect the source and the target generators through the latent code, and construct positive-negative pairs to facilitate the contrastive learning. Our proposed method shows successful results both visually and quantitatively. Our study provides concrete information that future work could continue to focus on reducing the rate of diversity degradation in order to further improve few-shot image generation.

Acknowledgments: This research is supported by the National Research Foundation, Singapore under its AI Singapore Programmes (AISG Award No.: AISG2-RP-2021-021; AISG Award No.: AISG-100E2018-005). This project is also supported by SUTD project PIE-SGP-AI-2018-01. We thank anonymous reviewers and Chandrasegaran Keshigeyan for their fruitful discussion and kind help.

References

- [1] Milad Abdollahzadeh, Toubia Malekzadeh, and Ngai-Man Man Cheung. Revisit multimodal meta-learning through the lens of multi-task learning. *Advances in Neural Information Processing Systems*, 34, 2021. 3
- [2] Abulikemu Abuduweili, Xingjian Li, Humphrey Shi, Cheng-Zhong Xu, and Dejing Dou. Adaptive consistency regularization for semi-supervised transfer learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [3] Antreas Antoniou, Amos Storkey, and Harrison Edwards. Data augmentation generative adversarial networks. *arXiv preprint arXiv:1711.04340*, 2017. 2
- [4] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *International Conference on Learning Representations*, 2018. 1, 4
- [5] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33:9912–9924, 2020. 3
- [6] Keshigeyan Chandrasegaran, Ngoc-Trung Tran, and Ngai-Man Cheung. A closer look at fourier spectrum discrepancies for cnn-generated images detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7200–7209, 2021. 1
- [7] Keshigeyan Chandrasegaran, Ngoc-Trung Tran, Yunqing ZHAO, and Ngai man Cheung. To smooth or not to smooth? on compatibility between label smoothing and knowledge distillation, 2022. 3
- [8] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021. 3
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255, 2009. 8
- [10] Li Fei-Fei, Rob Fergus, and Pietro Perona. One-shot learning of object categories. *IEEE transactions on pattern analysis and machine intelligence*, 28(4):594–611, 2006. 3
- [11] Qianli Feng, Chenqi Guo, Fabian Benitez-Quiroz, and Aleix M Martinez. When do gans replicate? on the choice of dataset size. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6701–6710, 2021. 2
- [12] Christopher Fifty, Ehsan Amid, Zhe Zhao, Tianhe Yu, Rohan Anil, and Chelsea Finn. Measuring and harnessing transference in multi-task learning. *arXiv preprint arXiv:2010.15413*, 2020. 2, 5
- [13] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017. 3
- [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 1, 2
- [15] Jean-Bastien Grill, Florian Strub, Florent Althché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in Neural Information Processing Systems*, 33:21271–21284, 2020. 3
- [16] Yiluan Guo and Ngai-Man Cheung. Attentive weights generation for few shot learning via information maximization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13499–13508, 2020. 3
- [17] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. 3, 6
- [18] 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. *Advances in neural information processing systems*, 30, 2017. 2, 6, 7
- [19] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NeurIPS Deep Learning and Representation Learning Workshop*, 2015. 3
- [20] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. 2
- [21] Ashish Jaiswal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. A survey on contrastive self-supervised learning. *Technologies*, 9(1):2, 2021. 3
- [22] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *Advances in Neural Information Processing Systems*, 33:12104–12114, 2020. 3, 6
- [23] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 1, 2, 4, 6
- [24] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. 1, 4, 6
- [25] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. 2
- [26] J. Krause, M. Stark, J. Deng, and L. Fei-Fei. 3d object representations for fine-grained categorization. In *2013 IEEE International Conference on Computer Vision Workshops*, pages 554–561, 2013. 8
- [27] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012. 4

- [28] Yijun Li, Richard Zhang, Jingwan (Cynthia) Lu, and Eli Shechtman. Few-shot image generation with elastic weight consolidation. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 15885–15896. Curran Associates, Inc., 2020. 2, 3, 4, 6, 8
- [29] Weide Liu, Chi Zhang, Guosheng Lin, and Fayao Liu. Crnet: Cross-reference networks for few-shot segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4165–4173, 2020. 3
- [30] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6707–6717, 2020. 3
- [31] Sangwoo Mo, Minsu Cho, and Jinwoo Shin. Freeze the discriminator: a simple baseline for fine-tuning gans. In *CVPR AI for Content Creation Workshop*, 2020. 3, 6
- [32] Atsuhiko Noguchi and Tatsuya Harada. Image generation from small datasets via batch statistics adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2750–2758, 2019. 3, 4, 6
- [33] Utkarsh Ojha, Yijun Li, Jingwan Lu, Alexei A Efros, Yong Jae Lee, Eli Shechtman, and Richard Zhang. Few-shot image generation via cross-domain correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10743–10752, 2021. 2, 3, 4, 5, 6, 8
- [34] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2, 3, 5, 6
- [35] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009. 2, 3
- [36] Liam Paninski. Estimation of entropy and mutual information. *Neural computation*, 15(6):1191–1253, 2003. 5
- [37] Taesung Park, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *European Conference on Computer Vision*, pages 319–345. Springer, 2020. 3, 6
- [38] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pages 4077–4087, 2017. 3
- [39] Trevor Standley, Amir Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Which tasks should be learned together in multi-task learning? In *International Conference on Machine Learning*, pages 9120–9132. PMLR, 2020. 2, 5
- [40] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, pages 1199–1208, 2018. 3
- [41] Ngoc-Trung Tran, Viet-Hung Tran, Ngoc-Bao Nguyen, Trung-Kien Nguyen, and Ngai-Man Cheung. On data augmentation for gan training. *IEEE Transactions on Image Processing*, 30:1882–1897, 2021. 3
- [42] Hung-Yu Tseng, Lu Jiang, Ce Liu, Ming-Hsuan Yang, and Weilong Yang. Regularizing generative adversarial networks under limited data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7921–7931, 2021. 2
- [43] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. 8
- [44] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. Cnn-generated images are surprisingly easy to spot... for now. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8695–8704, 2020. 1
- [45] Xiaogang Wang and Xiaoou Tang. Face photo-sketch synthesis and recognition. *IEEE transactions on pattern analysis and machine intelligence*, 31(11):1955–1967, 2008. 2, 6
- [46] Yaxing Wang, Abel Gonzalez-Garcia, David Berge, Luis Herranz, Fahad Shahbaz Khan, and Joost van de Weijer. 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*, pages 9332–9341, 2020. 3, 4, 6
- [47] Yaxing Wang, Chenshen Wu, Luis Herranz, Joost van de Weijer, Abel Gonzalez-Garcia, and Bogdan Raducanu. Transferring gans: generating images from limited data. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 218–234, 2018. 2, 3, 5, 6
- [48] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018. 3
- [49] Jiayu Xiao, Liang Li, Chaofei Wang, Zheng-Jun Zha, and Qingming Huang. Few shot generative model adaption via relaxed spatial structural alignment. *arXiv preprint arXiv:2203.04121*, 2022. 3
- [50] Renjun Xu, Pelen Liu, Liyan Wang, Chao Chen, and Jindong Wang. Reliable weighted optimal transport for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4394–4403, 2020. 2
- [51] Ceyuan Yang, Yujun Shen, Yinghao Xu, and Bolei Zhou. Data-efficient instance generation from instance discrimination. *Advances in Neural Information Processing Systems*, 34, 2021. 2
- [52] Ceyuan Yang, Yujun Shen, Zhiyi Zhang, Yinghao Xu, Jiapeng Zhu, Zhirong Wu, and Bolei Zhou. One-shot generative domain adaptation. *arXiv preprint arXiv:2111.09876*, 2021. 3
- [53] Jordan Yaniv, Yael Newman, and Ariel Shamir. The face of art: landmark detection and geometric style in portraits. *ACM Transactions on graphics (TOG)*, 38(4):1–15, 2019. 1, 6
- [54] Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 6
- [55] Han Zhang, Zizhao Zhang, Augustus Odena, and Honglak Lee. Consistency regularization for generative adversarial networks. *Proceedings of 6th International Conference on Learning Representations ICLR*, 2020. 3
- [56] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 3, 4

- [57] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable augmentation for data-efficient gan training. *Advances in Neural Information Processing Systems*, 33:7559–7570, 2020. [3](#)
- [58] Zhengli Zhao, Zizhao Zhang, Ting Chen, Sameer Singh, and Han Zhang. Image augmentations for gan training. *arXiv preprint arXiv:2006.02595*, 2020. [3](#)
- [59] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. [2](#)