Generalized One-shot Domain Adaption of Generative Adversarial Networks

Anonymous Author(s)
Affiliation
Address
email

Figure 1: Illustrations of the generalized one-shot GAN adaptation task. Target images with the corresponding binary entity masks are provided to adapt the pre-trained GAN to the domain of similar style and entity as target images. The first three masks are all zeros and the others have value 1 in the areas bounded by green lines. Our adapted model is of strong cross-domain correspondence to realize style and entity transfer.

Abstract

The adaption of Generative Adversarial Network (GAN) aims to transfer a pre-trained GAN to a given domain with limited training data. In this paper, we focus on the one-shot case, which is more challenging and rarely explored in previous works. We consider that the adaptation from source domain to target domain can be decoupled into two parts: the transfer of global style like texture and color, and the emergence of new entities that do not belong to the source domain. While previous works mainly focus on the style transfer, we propose a novel and concise framework[1] to address the *generalized one-shot adaption* task for both style and entity transfer, in which a reference image and its binary entity mask are provided. Our core objective is to constrain the gap between the internal distributions of the reference and syntheses by sliced Wasserstein distance. To better achieve it, style fixation is used at first to roughly obtain the exemplary style, and an auxiliary network is introduced to the original generator to disentangle entity and style transfer. Besides, to realize cross-domain correspondence, we propose the variational Laplacian regularization to constrain the smoothness of the adapted generator. Both quantitative and qualitative experiments demonstrate the effectiveness of our method in various scenarios.

1 Introduction

Benefiting from numerous excellent pre-trained GANs (*e.g.,* styleGAN [15] and BigGAN [6]), GAN adaption has become an important research topic, which leverages the knowledge of GANs pre-trained on large-scale datasets to mitigate the lack of data and speed up the training on a new domain.

[1] `https://github.com/thevoidname/Generalized-One-shot-GAN-Adaption`

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

It can be divided into few-shot, one-shot and zero-shot [12] cases. The few-shot (usually ≥ 10) case has been extensively studied in the past years [39, 38, 24, 22]. The latest works [27, 40, 43] seek the cross-domain correspondence that keeps the shapes or contents of syntheses before and after adaption invariant. In recent times, the more challenging one-shot case has been explored by several works [44, 50, 19], which mainly focus on adapting GANs to the target domain of similar style as the given image. By utilizing GAN inversion [42], the adapted models can be applied to image style transfer and manipulation tasks.

However, there exist limitations in previous one-shot task settings, since an exemplar contains rich information more than general color and texture style. As shown in Fig. 1, where the source domain refers to natural faces, and the exemplars of target domains are artistic portraits with some *entities* (*e.g.,* glasses or accessories), previous works are only able to transfer the artistic style while the entities are neglected. This may lead to two problems: 1) In some cases, the entity (*e.g.,* hat), considered as an important part of the style, should be transferred with the color style simultaneously. 2) For the entities of large size, it easily causes artifacts by previous works due to the influence of entities on the whole color space. To tackle the above problems, we describe the target domain with the aid of an extra binary mask, which labels the entities we are interested in and helps to define the target domain more exactly and flexibly. We name the new task definition as *generalized one-shot GAN adaption*. The previous scenarios can be taken as the special case when the binary mask is full of zeros. We believe the exploration of the novel task is quite meaningful, which not only allows customization for more complex scenarios, but also facilitates further discovery and utilization of the prior knowledge in pre-trained models.

In specific, we propose a concise manifold regularized GAN adaption framework to deal with the novel task. Following previous works [44, 50, 19], StyleGAN2 [16] is selected as our base model. We modify the architecture of the original generator to decouple the adaption of style and entity, where the new generator comprises an additional auxiliary network to facilitate the generation of entity, and the original generator is dedicated to generating stylized images. To train efficiently, we first compute the latent code of the reference by GAN inversion, and mix it with the random latent codes for rough style transfer. Then we employ the sliced Wasserstein distance [5] to minimize the divergence of the internal distributions between the exemplar and syntheses for learning the domain knowledge More importantly, to prevent the content distortion during training, inspired by the classic manifold learning [13, 2] we propose the variant Laplacian regularization to smooth the network changes before and after adaptation and preserve the geometric structure of the source generative manifold. We conduct extensive experiments on various references with and without entities. The results show that our framework can fully exploit the cross-domain correspondence and achieve high transfer quality. Moreover, the adapted model can handily perform various image manipulations beyond previous works.

2 Related works

Few-shot GAN adaption [39] is the first to study and evaluate different transferring strategies of GANs, it shows that transferring both pre-trained generator and discriminator can efficiently improve the convergence time and the quality of the generated images. [26] finds that merely tuning the parameters for batch statistics of the source generator is also valid. [38] discovers that fixing the source generator and only transforming the noise space into a new space works for most cases. [24] tunes GANs by freezing the low layers of the discriminator. [49] demonstrates that transferring low-level filters of pre-trained GANs can facilitate generation. [22] proposes an Elastic Weight Consolidation loss to regularize the change of GANs. [30] tries to adapt the singular values while freezing the singular vectors of the pre-trained weights. As an important breakthrough, [27] proposes the cross-domain consistency loss to maintain the diversity of the source generator, which makes the samples from the adapted generator and source generator have a corresponding relation like style transfer. Some works [20, 40, 43, 17] are based on and improve [27] to get better performance.

Zero-shot GAN adaption Recently many works [28, 18, 10] have explored the zero-shot domain transfer, they define the target domain by the text prompts, and align the semantic features of image and text predicted by the pre-trained image-text model CLIP [29]. StyleGAN-nada [12] proposes the CLIP directional loss to represent the domain gap by the difference of textual embeddings from CLIP.

One-shot GAN adaption MTG [50] is the first to tackle the one-shot GAN adaption task, it follows [12] to preserve the CLIP direction extracted by CLIP image encoder for adaption. OSGD [44] introduces extra modules to the original architecture to better reuse the prior knowledge. [19] proposes a contrastive regularization for patchwise relationships in the CLIP space. However, these models perform unstably in practice, and they mainly consider the style transfer without entity.

3 Preliminaries

StyleGAN A canonical StyleGAN [16] can be summarized in two parts: Firstly, given a noise distribution $P(\boldsymbol{z})$, a mapping network transforms the noise from space $\mathcal{Z} = \{\boldsymbol{z}|P(\boldsymbol{z}) \neq 0\} \subseteq \mathbb{R}^{1\times 512}$ into the replicated latent space $\mathcal{W} = \{[F(\boldsymbol{z}); \ldots; F(\boldsymbol{z})]|\boldsymbol{z} \in \mathcal{Z}\} \subseteq \mathbb{R}^{18\times 512}$, F is a fully-connected network. Secondly, a synthesis network G composed of convolutional blocks transfers $\mathcal{W}$ space into image space $\mathcal{M} = \{G(\boldsymbol{w})|\boldsymbol{w} \in \mathcal{W}\} \subseteq \mathbb{R}^{H\times W\times 3}$. Another important concept is the extended latent space $\mathcal{W}^+ = \{[F(\boldsymbol{z}_1); \ldots; F(\boldsymbol{z}_{18})]|z_i \in \mathcal{Z}\}$. For $\forall\, \boldsymbol{w}_1, \boldsymbol{w}_2 \in \mathcal{W}$, the latent code manipulation merges their information by the linear combination $diag(\boldsymbol{\alpha}_1)\boldsymbol{w}_1 + diag(\boldsymbol{\alpha}_2)\boldsymbol{w}_2 \in \mathcal{W}^+$, $\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2 \in \mathbb{R}^{18}$. Empirically, $\mathcal{W}^+$ space is inferior to $\mathcal{W}$ in generation fidelity and editability, but superior in GAN inversion that finds a code $\boldsymbol{w}_{ref}$ to reconstruct a reference $\boldsymbol{y}_{ref} \in \mathbb{R}^{H\times W\times 3}$. In this paper we only focus on tuning G, and fix F for keeping the distribution of source latent spaces.

Manifold learning Considering the sample matrix $\boldsymbol{X} = [\boldsymbol{x}_1^T; \ldots; \boldsymbol{x}_n^T]$ from source manifold $\mathcal{M}_s$, $\boldsymbol{W} = [w_{i,j}]_{i,j=1}^n$ is the weight matrix, where $w_{i,j}$ is the weight of $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$ usually defined as $e^{-\|\boldsymbol{x}_i - \boldsymbol{x}_j\|^2/\sigma}$. Given a task specific function f and $\boldsymbol{y} = f(\boldsymbol{x})$, manifold learning requires the transformed samples $\boldsymbol{Y} = [\boldsymbol{y}_1^T; \ldots; \boldsymbol{y}_n^T]$ on the new manifold $\mathcal{M}_t$ should preserve the geometric structure of $\mathcal{M}_s$, which means the data nearby on $\mathcal{M}_s$ are also nearby on $\mathcal{M}_t$, and vice versa. The requirement is mostly realized by minimizing the manifold regularization term [4] $\int_{\mathcal{M}_s} \|\nabla_{\mathcal{M}_s}\boldsymbol{y}\|^2\, dP(\boldsymbol{x})$, $\nabla_{\mathcal{M}_s}$ is the gradient along manifold, and $P(\boldsymbol{x})$ is the probability measure. If $\mathcal{M}_t$ is a submanifold of $\mathbb{R}$, the former is equalized with $\int_{\mathcal{M}_s} \boldsymbol{y}\Delta_{\mathcal{M}_s}\boldsymbol{y}dP(\boldsymbol{x})$ by Stokes theorem. Δ called the Laplace–Beltrami operator, is a key geometric object in the Riemannian manifold, thus the integral is also called Laplacian regularization. In practice, it is estimated by the discrete form $2tr(\boldsymbol{Y^T L Y}) = \sum_{i,j} w_{i,j} \left\|\boldsymbol{y}_i - \boldsymbol{y}_j\right\|^2$, $\boldsymbol{L} = \boldsymbol{D} - \boldsymbol{W}$ is known as the graph Laplacian operator, $\boldsymbol{D}$ is the degree matrix where the diagonal element $d_{i,i} = \sum_{j=1}^n w_{i,j}$ and other elements are zero. The convergence between $\boldsymbol{L}$ and Δ can be found in [3]. Intuitively, the regularization encourages the local smoothness of f. If data distribute densely in an area, the function should be smooth to avoid local oscillations. If $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$ are close, in the discrete form $w_{i,j}$ will impose a strong penalty to the distance between $\boldsymbol{y}_i$ and $\boldsymbol{y}_j$.

Task definition The generalized one-shot adaption can be formally described as: Given a reference $\boldsymbol{y}_{ref}$, a mask $\boldsymbol{m}_{ref} \in \{0,1\}^{H\times W}$ is provided to label the location of entity contained in $\boldsymbol{y}_{ref}$. The target domain $\mathcal{T}$ is defined to contain images of similar style and entities as $\boldsymbol{y}_{ref}$. With the knowledge stored in a generative model G_s pre-trained on source domain $\mathcal{S}$, a generative model G_t is learned to generate diverse images belonging to domain $\mathcal{T}$. Note that we can set $\boldsymbol{m}_{ref} = \boldsymbol{0}$ when there is no interested entities existed in $\boldsymbol{y}_{ref}$. Moreover, the adaptation should satisfy the cross-domain correspondence. *i.e.*, $\forall\, \boldsymbol{w} \in \mathcal{W}$, $G_s(\boldsymbol{w})$ and $G_t(\boldsymbol{w})$ should have similar shape or content in visual. With the correspondence, adapted models can be applied for not only synthesis, but also the transfer and manipulation of source images.

4 Methodology

4.1 Overall

Our framework for tackling the generalized one-shot GAN adaption task is illustrated in Fig. 2. For the sake of description, in the following part, we will take the StyleGAN pre-trained on FFHQ as an

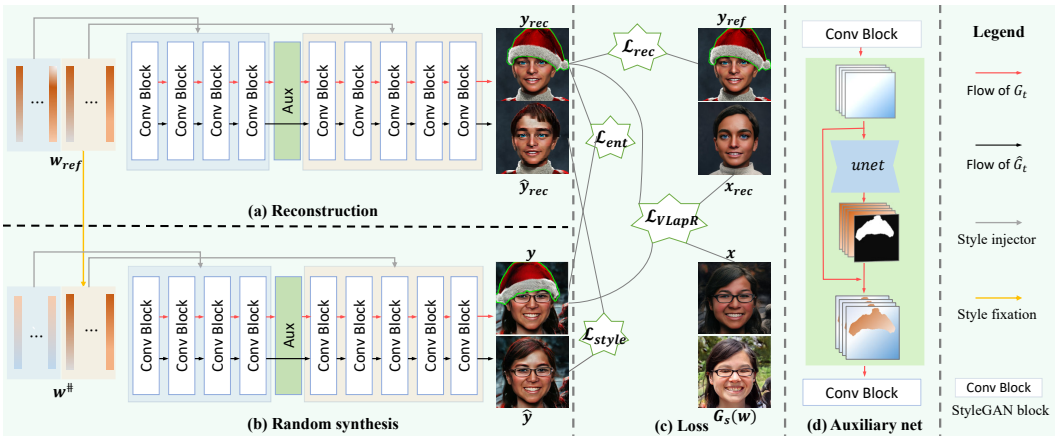

Figure 2: We propose a novel framework for generalized one-shot GAN adaption. The convolutional blocks are inherited from source StyleGAN. The black arrows represent the flow of $\hat{G}_t$, and the red arrows represent the flow of G_t. For other details please refer to the text in Sec. 4.1.

example to describe each component of the framework in detail. We will demonstrate our framework can be applied to other source domains in the experiments.

Network architecture Different from previous works [50, 19] that directly adapt the source generator G_s, our target generator G_t (as Fig. 2 shows) consists of a generator $\hat{G}_t$ inherited from G_s, and an auxiliary network (aux) trained from scratch. The design is inspired by the fact that G_s can approximate arbitrary real faces but cannot handle most entities since G_s captures the prevalent elements (clear natural faces) of FFHQ rather than the tails (faces with entities like hats or ornaments) of the distribution [46, 25]. We make aux serve to entities, and $\hat{G}_t$ just focus on stylizing clear face to benefit from the prior knowledge stored in G_s. As shown in Fig. 2(d), for each $\boldsymbol{w}$, its feature map generated by the convolutional block of StyleGAN is fed into aux, where a UNet [31] predicts the feature map $\boldsymbol{f}_{ent}$ of entity and the mask $\boldsymbol{m}$ that exactly labels the position of the entity, then they are merged with $\boldsymbol{f}_{in}$ by Eq. (1) to get the feature map of face covered with entity.

$$aux(\boldsymbol{f}_{in}) = (\boldsymbol{1} - \boldsymbol{m}) \odot \boldsymbol{f}_{in} + \boldsymbol{m} \odot \boldsymbol{f}_{ent}, \tag{1}$$

$\odot$ denotes the Hadamard product. We hope the shape of the entity is only influenced by the content of synthesis, hence we place aux after the fourth block of StyleGAN, where the 32×32 resolution $\boldsymbol{f}_{in}$ contains the most content information and the least style information of synthesis (see Sec. 4.2). Because the global style and local entity are usually independent, the disentangled design also prevents the stylization of the face from being contaminated by the style of entity (see Sec. 4.3).

Training process Our training process can be divided into three main parts: ① Style fixation and exemplar reconstruction. In Fig. 2(a), we get the latent code $\boldsymbol{w}_{ref} = [\boldsymbol{w}^c_{ref}; \boldsymbol{w}^s_{ref}]$ of $\boldsymbol{y}_{ref}$ by the GAN inversion technique [36], and train G_t to reconstruct $\boldsymbol{y}_{ref}$ with reconstruction loss $\mathcal{L}_{rec}$. In Fig. 2(b), each $\boldsymbol{w}$ will be transformed into the style-fixed code $\boldsymbol{w}^{\sharp} = [\boldsymbol{w}^c; \boldsymbol{w}^s_{ref}]$ to roughly obtain the exemplary style. ② Internal distribution learning. We minimize the divergence of internal patch distributions between syntheses and exemplar for style and entity transfer. Instead of using GAN loss that is hard to optimize, we take sliced Wasserstein distance (SWD) [5] as style loss $\mathcal{L}_{style}$ and entity loss $\mathcal{L}_{ent}$ to learning internal distribution efficiently. ③ Manifold Regularization. To suppress the content distortion during training, we propose the variational Laplacian regularization $\mathcal{L}_{VlapR}$ to smooth the change from G_s to G_t by preserving the geometric structure of the source manifold. As shown in Fig. 2(c), the overall loss is

$$\min_{G_t} \lambda_1 \mathcal{L}_{rec} + \lambda_2 \mathcal{L}_{style} + \lambda_3 \mathcal{L}_{ent} + \lambda_4 \mathcal{L}_{VLapR} \tag{2}$$

We will illustrate and discuss the training process in detail in the following sections. For convenience of description and self-consistency, we may use the abbreviations: $\boldsymbol{y} = G_t(\boldsymbol{w}^{\sharp})$, $\hat{\boldsymbol{y}} = \hat{G}_t(\boldsymbol{w}^{\sharp})$, $\boldsymbol{y}_{rec} = G_t(\boldsymbol{w}_{ref})$, $\hat{\boldsymbol{y}}_{rec} = \hat{G}_t(\boldsymbol{w}_{ref})$, $\boldsymbol{x} = G_s(\boldsymbol{w}^{\sharp})$, $\boldsymbol{x}_{rec} = G_s(\boldsymbol{w}_{ref})$.

4.2 Style fixation and exemplar reconstruction

It is well known that StyleGAN is equipped with the disentangled latent spaces $\mathcal{W}$ and $\mathcal{W}^+$ [1], where the latter vectors of code mainly determine the style of synthesis whereas the earlier vectors determine the coarse-structure or content of the synthesis. Taking advantage of this trait, we first roughly transfer the exemplary style to other syntheses by fixing the style code. The exemplar $\boldsymbol{y}_{ref}$ is fed into the pre-trained GAN inversion encoder e4e [36] to get the latent code $\boldsymbol{w}_{ref} = [\boldsymbol{w}^c_{ref}; \boldsymbol{w}^s_{ref}]$, where $\boldsymbol{w}^c_{ref} \in \mathbb{R}^{l\times 512}$ and $\boldsymbol{w}^s_{ref} \in \mathbb{R}^{(18-l)\times 512}$ respectively encode content and style information of exemplar. For each $\boldsymbol{w} \in \mathcal{W}$, it will be transformed to $\boldsymbol{w}^\sharp$ before fed into networks,

$$\boldsymbol{w}^\sharp = diag(\boldsymbol{\alpha})\boldsymbol{w} + diag(\boldsymbol{1} - \boldsymbol{\alpha})\boldsymbol{w}_{ref},\ \alpha_i = \mathbb{1}_{i<=l}(i),\ i = 1, \ldots, 18. \tag{3}$$

The content part of $\boldsymbol{w}$ is reserved and style part is replaced by $\boldsymbol{w}^s_{ref}$, and we denote the new space as $\mathcal{W}^\sharp$. Because style fixation should not change the content of the original synthesis, by using the pre-trained Arcface model [11] to compare the identity similarity before and after style fixation, we find setting $l = 8$ is acceptable which gets 65% average similarity and obtains the adequate exemplary style in visual. In MTG [50], style mixing assists as a post-processing step to improve the style quality, but the experiments (Fig. 3) show that the style of syntheses is underfitting and not consistent. In our framework, it is used as a pre-processing step to facilitate the following learning.

Because $\boldsymbol{x}_{rec}$ is just the projection of $\boldsymbol{y}_{ref}$ on the source domain, there often exists obvious distinction between them in visual as shown in Fig. 2. We use the reconstruction loss $\mathcal{L}_{rec}$ to narrow the gap,

$$\mathcal{L}_{rec} = dssim(\boldsymbol{y}_{rec}, \boldsymbol{y}_{ref}) + lpips(\boldsymbol{y}_{rec}, \boldsymbol{y}_{ref}) + \lambda_5 mse(\boldsymbol{m}^\uparrow_{rec}, \boldsymbol{m}_{ref}). \tag{4}$$

$dssim$ denotes the negative structural similarity metric [41] and $lpips$ is the perceptual loss [47], $\boldsymbol{m}_{rec}$ is the upsampled reconstructed mask from aux. It is worth noting $\mathcal{L}_{rec}$ acts on G_t rather than the vanilla architecture $\hat{G}_t$. If $\boldsymbol{m}_{ref} = \boldsymbol{0}$, aux does not work and G_t be equivalent to $\hat{G}_t$.

4.3 Internal distribution learning

Recent works about internal learning [32, 34, 48, 21] prove that the internal patch distribution of a single image contains rich meaning. Minimizing the divergence of internal distributions works out many tasks like image generation, style transfer, and super-resolution. It is usually achieved by the adversarial patch discriminator [14]. However, it is well known that the training of GAN is time-consuming, unstable, and requires large GPU memory. Previous few-shot GAN adaption works [22, 27] using GAN loss almost have the serious model collapse phenomenon. We find that slice Wasserstein distance (SWD) [5] can lead to the same destination but with greater efficiency. For two tensor $\boldsymbol{A}, \boldsymbol{B} \in \mathbb{R}^{H\times W\times d}$, the SWD of their empirical internal distributions is defined as

$$SWD(\boldsymbol{A}, \boldsymbol{B}) = \int_{\mathbf{S}^d} \|sort(proj(\boldsymbol{A}, \boldsymbol{\theta})) - sort(proj(\boldsymbol{B}, \boldsymbol{\theta}))\|^2 \, d\boldsymbol{\theta}, \tag{5}$$

where $\mathbf{S}^d = \{\boldsymbol{\theta} \in \mathbb{R}^d : \|\boldsymbol{\theta}\| = 1\}$, $proj : \mathbb{R}^{H\times W\times d} \to \mathbb{R}^{H\times W}$ that projects each pixel from d channels to a scalar by $\boldsymbol{\theta}$, $sort$ denotes the sort operator ordering all values. SWD does not introduce extra parameter, and can be readily realized by a 1×1 convolutional kernel with quicksort algorithm.

To realize style and entity transfer, we use SWD to minimize the divergence of the empirical internal distribution of syntheses and reference. The style loss is defined as

$$\mathcal{L}_{style} = \frac{1}{n|\Phi_{style}|} \sum_{\varphi\in\Phi_{style}} \sum_{i=1}^{m} SWD(\varphi(\hat{\boldsymbol{y}}_i), \varphi(\hat{\boldsymbol{y}}_{rec})). \tag{6}$$

Φ_{style} is a set of convolutional layers from pre-trained $lpips$ network to extract spatial features. Note that the generator learns the style of $\hat{\boldsymbol{y}}_{rec}$ rather than that of $\boldsymbol{y}_{ref}$ to prevent the style of the entity from leaking into the other area. The entity loss is defined as

$$\mathcal{L}_{ent} = \frac{1}{n|\Phi_{ent}|} \sum_{\varphi\in\Phi_{ent}} \sum_{i=1}^{m} SWD(\varphi(\boldsymbol{m}^\uparrow_i \odot \boldsymbol{y}_i), \varphi(\boldsymbol{m}^\uparrow_{rec} \odot \boldsymbol{y}_{rec})). \tag{7}$$

Φ_{ent} is also a set of convolutional layers from pre-trained $lpips$ net, $\boldsymbol{m}^\uparrow$ is upsampled by $\boldsymbol{m}$ to get the same resolution as the final synthesis. Taking the reconstructed entity $\boldsymbol{m}^\uparrow_{rec} \odot \boldsymbol{y}_{rec}$ as the target can be viewed as an implicit augmentation technique to avoid overfitting.

Figure 3: Comparison of our framework with other SOTA adaption frameworks. The source domain contains the 1024×1024 natural face images. Please zoom in for better visual effects.

4.4 Manifold regularization

The training losses, particularly the style loss, inevitably distort the contents of source syntheses, and change their relative relationship. For achieving the cross-domain correspondence, we use the variant Laplacian regularization $\mathcal{L}_{VLapR}$ to alleviate the distortion. Since we do not tune the mapping network of StyleGAN, the $\mathcal{W}$ space coupled with source generator G_s is not changed, we define

$$\mathcal{L}_{VLapR} = \int_{\mathcal{W}} \left\| \nabla_{\mathcal{W}^\sharp} \phi(G_t(\boldsymbol{w}^\sharp)) - \nabla_{\mathcal{W}^\sharp} \phi(G_s(\boldsymbol{w}^\sharp)) \right\|^2 dP(\boldsymbol{w}), \tag{8}$$

where ϕ is a pre-trained semantic extractor, here we use the pre-trained CLIP image encoder [29] that is sensitive to both style and content variation. $\mathcal{L}_{VLapR}$ can be interpreted from two perspectives. The first one is the smoothness of function: Because the integrand also represents the smoothness of the residual $\phi(G_t(\boldsymbol{w}^\sharp)) - \phi(G_s(\boldsymbol{w}^\sharp))$, minimizing $\mathcal{L}_{VlapR}$ expects syntheses from G_s and G_t across $\mathcal{W}$ have a smooth semantic difference. In the idea case, the integrand is nearly zero, thus for various $\boldsymbol{w}$ in a local area, $G_t(\boldsymbol{w}^\sharp)$ and $G_s(\boldsymbol{w}^\sharp)$ have similar difference in terms of style and entity. The second perspective is the geometric structure preservation: According to the theory of Laplacian regularization [3], the above integral form in Eq. (8) can be estimated by a discrete form,

$$\mathcal{L}_{VLapR} = \sum_{i,j=1}^{n} w_{i,j} \left\| \phi(\boldsymbol{y}_i) - \phi(\boldsymbol{y}_j) + \phi(\boldsymbol{x}_j) - \phi_j(\boldsymbol{x}_i) \right\|^2 \tag{9a}$$

$$= 2tr(\boldsymbol{R}^T \boldsymbol{L} \boldsymbol{R}), \tag{9b}$$

where $\boldsymbol{R} = [\phi(\boldsymbol{y}_1)^T - \phi(\boldsymbol{x}_1)^T; \ldots; \phi(\boldsymbol{y}_n)^T - \phi(\boldsymbol{x}_n)^T]$, $w_{i,j} = e^{-\|\boldsymbol{w}_i^\sharp - \boldsymbol{w}_j^\sharp\|/\sigma}$, $\boldsymbol{L}$ is the Laplacian matrix (see Sec. 3). In Eq. (9a), if $\boldsymbol{w}_i$ and $\boldsymbol{w}_j$ are close in latent space, $w_{i,j}$ will be a large scalar encouraging $\left\|\phi(\boldsymbol{y}_i) - \phi(\boldsymbol{y}_j)\right\|$ and $\|\phi(\boldsymbol{x}_i) - \phi(\boldsymbol{x}_j)\|$ to be same. It means the relative distance of adjacent syntheses before and after adaption shall be isometric in the feature space. In practice, the loss can be efficiently computed by Eq. (9b), the derivation is provided in supplementary materials.

Relation between $\mathcal{L}_{VlapR}$ and $\mathcal{L}_{CDC}$ Recently the influential work FSGA [27] introduces the cross-domain distance consistency loss $\mathcal{L}_{CDC}$ to achieve cross-domain correspondence. For arbitrary $\boldsymbol{w}_i$ and $\boldsymbol{w}_j$, without loss of generality FSGA defines the conditional probability $p_{j|i}$ and $q_{j|i}$ in Eq. (10) to form the similarity distributions, then minimizes their Kullback-Leibler divergence Eq. (11) to encourage syntheses of G_t have the identical similarity property as that of G_s.

$$p_{j|i} = \frac{e^{-\|\phi(\boldsymbol{x}_j) - \phi(\boldsymbol{x}_i)\|^2}}{\Sigma_{j \neq i}^{n} e^{-\|\phi(\boldsymbol{x}_j) - \phi(\boldsymbol{x}_i)\|^2}}, \; q_{j|i} = \frac{e^{-\|\phi(\boldsymbol{y}_j) - \phi(\boldsymbol{y}_i)\|^2}}{\Sigma_{j \neq i}^{n} e^{-\|\phi(\boldsymbol{y}_j) - \phi(\boldsymbol{y}_i)\|^2}}. \tag{10}$$

$$\mathcal{L}_{CDC} = \Sigma_{i=1}^{n} KL(p_{\cdot|i} || q_{\cdot|i}) = \Sigma_{i=1}^{n} \Sigma_{j \neq i}^{n} p_{j|i} \log \frac{p_{j|i}}{q_{j|i}}. \tag{11}$$

In the above equation, ϕ denotes the feature map from the corresponding generator. From the perspective of manifold learning, $p_{\cdot|i}$ and $q_{\cdot|i}$ depict the neighborhood structure of $\boldsymbol{x}_i$ and $\boldsymbol{y}_i$ respectively,

Table 1: Quantitative results. The target images are showed in Fig. 3. US denotes the user preference. NME and ID are not counted on the Mask since the masked face cannot be identified by [7] and [11].

	Sketch			Disney			Arcane			Hat			Zelda			Mask
	NME↓	ID↑	US↑	NME↓	ID↑	US↑	NME↓	ID↑	US↑	NME↓	ID↑	US↑	NME↓	ID↑	US↑	US↑
FSGA	0.14	0.11	0.09	0.26	0.00	0.09	0.17	0.04	0.09	0.19	0.03	0.10	0.28	0.00	0.09	0.10
MTG	0.08	0.31	0.09	0.19	0.14	0.11	0.08	0.24	0.10	0.09	0.29	0.10	0.15	0.05	0.09	0.10
JoJoGAN	0.09	0.24	0.13	0.11	0.17	0.11	0.10	0.17	0.11	0.08	0.21	0.10	0.09	0.09	0.11	0.10
OSCLIP	0.11	0.20	0.09	0.11	0.19	0.11	0.10	0.19	0.09	0.13	0.23	0.10	0.21	0.13	0.10	0.10
Ours	**0.08**	**0.40**	**0.60**	**0.09**	**0.26**	**0.58**	**0.07**	**0.25**	**0.61**	**0.08**	**0.36**	**0.60**	**0.09**	**0.25**	**0.61**	**0.60**
w/o $\mathcal{L}_{style}$	0.07	0.45	-	0.07	0.40	-	0.07	0.32	-	0.08	0.40	-	0.08	0.35	-	-
w/o $\mathcal{L}_{VLapR}$	0.09	0.35	-	0.10	0.11	-	0.08	0.21	-	0.10	0.29	-	0.12	0.15	-	-

$\mathcal{L}_{CDC}$ tries to place $\boldsymbol{x}_i$ in a new space so as to optimally preserve neighborhood identity. Actually, we find that this idea is equivalent to the Stochastic Neighbor Embedding [13, 37], which is also a popular method to preserve the manifold structure. $\mathcal{L}_{CDC}$ has the following disadvantages. Firstly, it needs a large batch that is nearly unacceptable for high-resolution GAN to estimate $p_{\cdot|i}$ and $q_{\cdot|i}$. $\mathcal{L}_{VLapR}$ directly constrains the distance and performs well even when the batch is 2. Secondly, $\mathcal{L}_{CDC}$ is non-convex and hard to optimize, in practice the loss value is usually around 1e-5, implying $\mathcal{L}_{CDC}$ does not provide an effective penalty to the differences. Finally, because the $softmax$ form of $p_{j|i}$ and $q_{j|i}$ is scale-invariant about the inputs, $\mathcal{L}_{CDC}$ cannot guarantee the isometric relationship of source and adapted syntheses like $\mathcal{L}_{VLapR}$, thus it is hard to avoid mode collapse. In the Supplementary, we show replacing $\mathcal{L}_{CDC}$ with $\mathcal{L}_{VLapR}$ can hugely remedy the collapse of FSGA.

5 Experiments

Our implementation is based on the official code of StyleGAN [2]. If $\boldsymbol{m}_{ref} \neq \mathbf{0}$, the total epoch is 2000, we adopt the Adam optimizer with learning rate $1e-3$, $\beta_1 = 0$, $\beta_2 = 0.999$. The cosine annealing strategy of learning rate is adopted to reduce the learning rate to $1e-4$ gradually. In Eq. (2), $\lambda_1 = 10$, $\lambda_2 = 0.2$, $\lambda_3 = 2$, $\lambda_4 = 1$, and in Eq. (4) $\lambda_5 = 100$. We use the pre-trained $lpips$ network with vgg architecture [35] containing five convolutional blocks. In Eq. (6) and Eq. (7), $\Phi_{style} = \{vgg_{1:5}\}$, $\Phi_{ent} = \{vgg_{1:1}, vgg_{1:2}, vgg_{1:3}, vgg_{1:4}\}$, $m = 1$. In Eq. (9b), $n = 2$ including the samples from reconstruction and internal distribution learning, σ for calculating the graph weights is 128. We use the Monte Carlo simulation that randomly sampling 256 vectors on the unit sphere to compute the integral in Eq. (5). We only augment the reference by the horizontal flip. If $\boldsymbol{m}_{ref} = \mathbf{0}$, we change λ_2 to 2, λ_4 to 0.5, Φ_{style} to $\{vgg_{1:4}, vgg_{1:5}\}$, and epoch to 1000 for speeding up training.

5.1 Comparison with SOTA methods

We take the few-shot adaption method FSGA [27], one-shot adaption methods MTG [50] and OSCLIP [19], one-shot face style transfer method JoJoGAN [9] to compare on the commonly studied face domain. Since these works can not deal with the entity yet, for a fair and comprehensive comparison, as shown in Fig. 3 we select three portraits without entity (Sketch, Disney, Arcane) from previous works and three portraits with entities (hat, Zelda elf decorations, and mask) from AAHQ [23] .

Qualitative results Fig. 3 shows the qualitative results. The top row images are the source natural faces, they are inverted by e4e [36] to obtain the latent codes. From the results, we can draw the following conclusions. Firstly, our syntheses are of competitive style as the given references. Though shown in thumbnails, they have more obvious high-frequency details to look sharp and realistic, the results of JoJoGAN are most fuzzy. Secondly, our method achieves the strongest cross-domain correspondence that preserves the content and shape of the source image, which is a great advance compared with other methods without a doubt. It also implies our method can keep the diversity of the source model without content collapse. Thirdly, our method can generate high-quality entities, which are adaptive to the various face shapes and look harmonious with the rest area of the syntheses. Experiments prove that our method has sufficient visual advantages over the competition.

[2] `https://github.com/NVlabs/stylegan3`

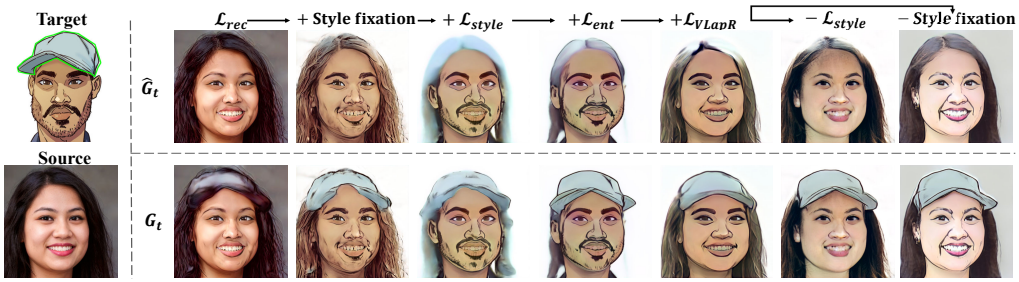

Figure 4: Ablation study. The arrows indicate how the current model has changed from the previous model. The upper and lower syntheses come from $\hat{G}_t$ and G_t respectively.

Quantitative results Previous works mainly conduct user studies to compare the visual quality. Following them, we survey user preferences for different methods. We first randomly sample 50 latent codes in $\mathcal{W}$ space and feed them into various models. We provide the references, source syntheses, and adapted syntheses to volunteers, they will consider both style and content and subjectively vote for their favorite adapted syntheses. Moreover, to objectively evaluate the cross-domain correspondence, we use the face alignment network [7] to extract facial landmarks and calculate the Normalized mean error (NME) of landmarks between source syntheses and adapted syntheses.

$$NME(G_s, G_t) = \frac{1}{n}\sum_{i=1}^{n}\frac{\|lmk(\boldsymbol{x}_i) - lmk(\boldsymbol{y}_i)\|^2}{\sqrt{HW}}. \tag{12}$$

lmk denotes the 68 points two-dimension landmarks extractor, n is set to 1000. NME reflects the face shape difference before and after adaption. We further report the identity (ID) similarity predicted by the Arcface [11] to metric the preservation of identity information. The model of the lowest NME and highest ID has the best cross-domain corresponding property. The quantitative results have been shown in Table 1. Our method undoubtedly obtains the best scores in terms of NME and ID. It performs stably and gets similar NME in different domains. FSGA and OSCLIP obtain relatively worse scores because of the overfitting as shown in Fig. 3, their common characteristic is that the GAN discriminator is used for training. MTG performs very unsteadily especially when the target domains have strong semantic style like Disney and Zelda. Compared with our approach JoJoGAN does similar well on NME, but poorly on ID. For the user study, we finally collect valid votes from 53 volunteers, and a user averagely spends 4 seconds on one question. It shows our method obtains about 60% of the votes on each adapted model, whereas other methods evenly acquire the rest of the votes, proving that our method is much more popular with users.

Training and inference time Our method costs about 12 minutes for $\boldsymbol{m}_{ref} \neq \mathbf{0}$ and 4 minutes for $\boldsymbol{m}_{ref} = \mathbf{0}$ on NVIDIA RTX 3090. FSGA, MTG, JoJoGAN, OSCLIP cost about 48, 9, 2, 24 minutes respectively. All of them take about 30ms to generate an image.

5.2 Ablation study

We start from the basic reconstruction loss $\mathcal{L}_{rec}$ and continuously improve the framework to observe the qualitative changes of syntheses. As illustrated in Fig. 4, the target exemplar is a hard case depicting a man with a beard and cap, and the style fixation can bring a rough but remarkable improvement to the style of synthesis. Although $\mathcal{L}_{style}$ can enhance the style of the face, it deteriorates the hair and introduces the noise like the beard to the female face. $\mathcal{L}_{ent}$ plays a critical role in the generation of entity, without it the cap can not be synthesized. Notably, $\mathcal{L}_{VLapR}$ has a significant suppression of the noise in previous syntheses. Finally, we remove $\mathcal{L}_{style}$ and fixation style respectively and the syntheses are no longer of the target style. It means that both of them play vital roles in style transfer. We also provide the quantitative ablation in Table 1 to study the impacts of $\mathcal{L}_{style}$ and $\mathcal{L}_{VLapR}$. It shows that removing $\mathcal{L}_{style}$ has a positive effect on NME and ID while removing $\mathcal{L}_{VLapR}$ has an opposite effect. The quantitative results confirm that $\mathcal{L}_{VLapR}$ is necessary for preventing content from being distorted by style loss.

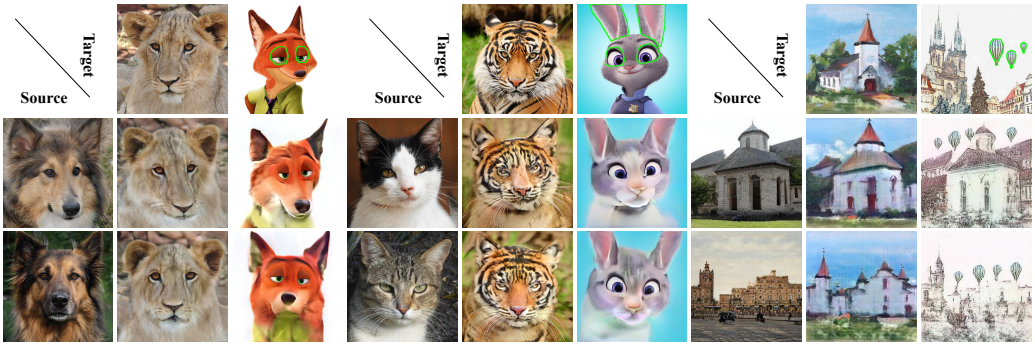

Figure 5: Our framework can transfer various source domains to various target domains.

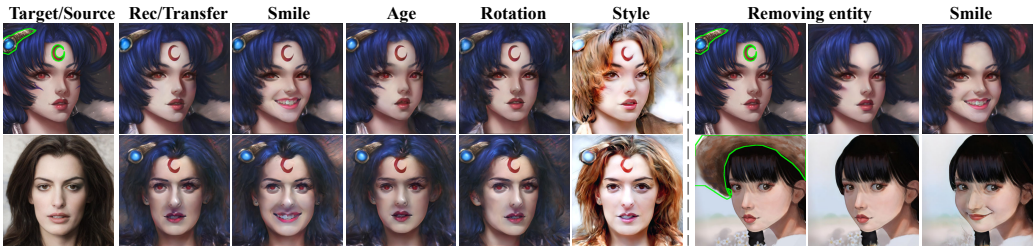

Figure 6: Applying the adapted model to manipulate target image and transferred image, including semantic editing, style transfer and entity removal.

5.3 More results

Adaption to other domains In addition to the success of the face domain, our framework can be applied to other source domains. We take StyleGANs pre-trained on 512×512 AFHQ dog, cat dataset [8] and 256×256 LSUN church dataset [45] as exemplars. As shown in Fig. 5, for each source domain, we provide the images with entities and without entity, the source images are randomly sampled from the source space. For the images without entity, the adapted syntheses are realistic and have a strong corresponding relationship in visual. For the other case, our approach also achieves satisfactory results, for example, the cartoon eyes of the dogs are in the right place and very similar to that of the reference. It proves that our framework is valid and has a wide universality.

Image manipulation Because our adapted model retains the geometric structure of the source generative space, it can efficiently perform semantic editing on image using the semantic direction found in the source space. As shown in Fig. 6, since our model is able to accurately reconstruct the target exemplar, we can edit both the exemplar and synthesis. We select three representative semantic directions [33] to change smile, age, and pose rotation, the results show that the adapted generator achieves exact control not only on the face but also on the entity. In addition, our model can also change the style by altering the style code, which is hard for previous works. Moreover, our framework can be applied to remove the entity with $\hat{G}_t$, and the result can be manipulated similarly. We consider these functions will be helpful for artistic creation.

6 Conclusion and limitations

In this paper, we construct a new GAN adaption framework for the generalized one-shot GAN adaption task. Because our framework relies heavily on the learning of internal distribution, it has some limitations inevitably. We think the most important one is that it cannot precisely control the position of the entity, which may lead to failures when the pose changes too much. Another limitation is that the entity cannot be too complex, otherwise it is difficult to learn by patch distribution. However, these extreme cases require additional consideration. In general, our framework is effective in various scenarios. We believe the exploration of the novel task and the introduction of manifold regularization are meaningful for the later work.

References

[1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *Proceedings of the IEEE international conference on computer vision*, pages 4432–4441, 2019.

[2] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15:1373–1396, 2003.

[3] Mikhail Belkin and Partha Niyogi. Towards a theoretical foundation for laplacian-based manifold methods. *J. Comput. Syst. Sci.*, 74:1289–1308, 2008.

[4] Mikhail Belkin, Partha Niyogi, and Vikas Sindhwani. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *J. Mach. Learn. Res.*, 7:2399–2434, 2006.

[5] Nicolas Bonneel, Julien Rabin, Gabriel Peyré, and Hanspeter Pfister. Sliced and radon wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 51:22–45, 2014.

[6] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019.

[7] Adrian Bulat and Georgios Tzimiropoulos. How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks). In *International Conference on Computer Vision*, 2017.

[8] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8185–8194, 2020.

[9] Min Jin Chong and David Forsyth. Jojogan: One shot face stylization. *ArXiv*, abs/2112.11641, 2021.

[10] Katherine Crowson, Stella Rose Biderman, Daniel Kornis, Dashiell Stander, Eric Hallahan, Louis Castricato, and Edward Raff. Vqgan-clip: Open domain image generation and editing with natural language guidance. *ArXiv*, abs/2204.08583, 2022.

[11] Jiankang Deng, J. Guo, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4685–4694, 2019.

[12] Rinon Gal, Or Patashnik, Haggai Maron, Gal Chechik, and Daniel Cohen-Or. Stylegan-nada: Clip-guided domain adaptation of image generators, 2021.

[13] Geoffrey E. Hinton and Sam T. Roweis. Stochastic neighbor embedding. In *NIPS*, 2002.

[14] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5967–5976, 2017.

[15] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4401–4410, 2019.

[16] 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.

[17] Chaerin Kong, Jeesoo Kim, Dong Yan Han, and Nojun Kwak. Smoothing the generative latent space with mixup-based distance learning. *ArXiv*, abs/2111.11672, 2021.

[18] Gihyun Kwon and Jong-Chul Ye. Clipstyler: Image style transfer with a single text condition. *ArXiv*, abs/2112.00374, 2021.

[19] Gihyun Kwon and Jong-Chul Ye. One-shot adaptation of gan in just one clip. *ArXiv*, abs/2203.09301, 2022.

[20] Hyuk-Gi Lee, Gi-Cheon Kang, Changhoon Jeong, Han-Wool Sul, and Byoung-Tak Zhang. C: Contrastive learning for cross-domain correspondence in few-shot image generation. 2021.

[21] B. Li, Yuanbiao Gou, J. Liu, Hongyuan Zhu, Joey Tianyi Zhou, and Xi Peng. Zero-shot image dehazing. *IEEE Transactions on Image Processing*, 29:8457–8466, 2020.

[22] Yijun Li, Richard Zhang, Jingwan Lu, and E. Shechtman. Few-shot image generation with elastic weight consolidation. *ArXiv*, abs/2012.02780, 2020.

[23] Mingcong Liu, Qiang Li, Zekui Qin, Guoxin Zhang, Pengfei Wan, and Wen Zheng. Blendgan: Implicitly gan blending for arbitrary stylized face generation. In *Advances in Neural Information Processing Systems*, 2021.

[24] Sangwoo Mo, Minsu Cho, and Jinwoo Shin. Freeze discriminator: A simple baseline for fine-tuning gans. *ArXiv*, abs/2002.10964, 2020.

[25] Ron Mokady, Michal Yarom, Omer Tov, Oran Lang, Daniel Cohen-Or, Tali Dekel, Michal Irani, and Inbar Mosseri. Self-distilled stylegan: Towards generation from internet photos. *ArXiv*, abs/2202.12211, 2022.

[26] Atsuhiro Noguchi and Tatsuya Harada. Image generation from small datasets via batch statistics adaptation. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2750–2758, 2019.

[27] 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. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10738–10747, 2021.

[28] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and D. Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2065–2074, 2021.

[29] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021.

[30] Esther Robb, Wen-Sheng Chu, Abhishek Kumar, and Jia-Bin Huang. Few-shot adaptation of generative adversarial networks. *arXiv preprint arXiv:2010.11943*, 2020.

[31] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.

[32] Tamar Rott Shaham, Tali Dekel, and T. Michaeli. Singan: Learning a generative model from a single natural image. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4569–4579, 2019.

[33] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. Interfacegan: Interpreting the disentangled face representation learned by gans. *arXiv preprint arXiv:2005.09635*, 2020.

[34] Assaf Shocher, S. Bagon, Phillip Isola, and M. Irani. Ingan: Capturing and retargeting the "dna" of a natural image. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4491–4500, 2019.

[35] K. Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2015.

[36] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics (TOG)*, 40:1 – 14, 2021.

[37] Laurens van der Maaten and Geoffrey E. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605, 2008.

[38] Yaxing Wang, Abel Gonzalez-Garcia, David Berga, Luis Herranz, Fahad Shahbaz Khan, and Joost van de Weijer. Minegan: Effective knowledge transfer from gans to target domains with few images. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9329–9338, 2020.

[39] Yaxing Wang, Chenshen Wu, Luis Herranz, Joost van de Weijer, Abel Gonzalez-Garcia, and Bogdan Raducanu. Transferring gans: generating images from limited data. In *European Conference on Computer Vision*, 2018.

[40] Yue Wang, Ran Yi, Ying Tai, Chengjie Wang, and Lizhuang Ma. Ctlgan: Few-shot artistic portraits generation with contrastive transfer learning. 2022.

[41] Zhou Wang, Alan Conrad Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13:600–612, 2004.

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

[43] Jiayu Xiao, Liang Li, Chaofei Wang, Zhengjun Zha, and Qingming Huang. Few shot generative model adaption via relaxed spatial structural alignment. *ArXiv*, abs/2203.04121, 2022.

[44] Ceyuan Yang, Yujun Shen, Zhiyi Zhang, Yinghao Xu, Jiapeng Zhu, Zhirong Wu, and Bolei Zhou. One-shot generative domain adaptation. *ArXiv*, abs/2111.09876, 2021.

[45] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.

[46] Ning Yu, Ke Li, Peng Zhou, Jitendra Malik, Larry Davis, and Mario Fritz. Inclusive gan: Improving data and minority coverage in generative models. *ArXiv*, abs/2004.03355, 2020.

[47] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.

[48] Zicheng Zhang, Yinglu Liu, Congying Han, Hailin Shi, Tiande Guo, and Bo Zhou. Petsgan: Rethinking priors for single image generation. *ArXiv*, abs/2203.01488, 2022.

[49] Miaoyun Zhao, Yulai Cong, and Lawrence Carin. On leveraging pretrained gans for generation with limited data. In *ICML*, 2020.

[50] Peihao Zhu, Rameen Abdal, John C. Femiani, and Peter Wonka. Mind the gap: Domain gap control for single shot domain adaptation for generative adversarial networks. *ArXiv*, abs/2110.08398, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [No] The code and the data are proprietary.
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No] The images we use are widely used in previous works or from open dataset.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No] The images we use are widely used in previous works or from open dataset.
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [No]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]