

# Deformable One-shot Face Stylization via DINO Semantic Guidance

Yang Zhou    Zichong Chen    Hui Huang\*  
 Visual Computing Research Center, Shenzhen University




Figure 1. We propose a **deformation-aware face stylization** framework trained on a single real-style image pair (left). Our framework can generate diverse, high-quality, stylized faces with desired deformations, while maintaining the input identity well (right).

## Abstract

This paper addresses the complex issue of one-shot face stylization, focusing on the simultaneous consideration of appearance and structure, where previous methods have fallen short. We explore deformation-aware face stylization that diverges from traditional single-image style reference, opting for a real-style image pair instead. The cornerstone of our method is the utilization of a self-supervised vision transformer, specifically DINO-ViT, to establish a robust and consistent facial structure representation across both real and style domains. Our stylization process begins by adapting the StyleGAN generator to be deformation-aware through the integration of spatial transformers (STN). We then introduce two innovative constraints for generator fine-tuning under the guidance of DINO semantics: i) a directional deformation loss that regulates directional vectors in DINO space, and ii) a relative structural consistency constraint based on DINO token self-similarities, ensuring diverse generation. Additionally, style-mixing is employed to align the color generation with the reference, minimiz-

ing inconsistent correspondences. This framework delivers enhanced deformability for general one-shot face stylization, achieving notable efficiency with a fine-tuning duration of approximately 10 minutes. Extensive qualitative and quantitative comparisons demonstrate our superiority over state-of-the-art one-shot face stylization methods. Code is available at <https://github.com/zichongc/DoesFS>.

## 1. Introduction

Face stylization is a stylish and eye-catching application favored by users in social media and the virtual world. Recent advances [22, 31, 36–38] mainly benefit from the capacity of generative models, e.g., StyleGAN [16–18]. When coping with artistic styles of extremely limited examples (e.g., one), it is necessary to prevent the training from over-fitting and mode collapse. Many one-shot methods [6, 21, 40, 45] have addressed this issue with different strategies. However, these approaches mainly focus on color and texture transfer, seldom exploring the geometric deformation potential. As exaggeration is an important characteristic of artistic style, structural deformations should also be emphasized in the stylization.

\*Corresponding author.

Can we stylize facial images according to only one style example while simultaneously considering *appearance change* and *structure exaggeration*? Previous methods [6, 21, 45] solve this problem by building “fake” guidance across the source and target domain based on the inversion of the style reference. However, as shown in Fig. 2, current GAN inversion techniques, such as [2] and [32], still can not produce faithful mappings for artistic images, which will definitely mislead the geometric deformation, especially when strong exaggerations occur. We argue that we can build reliable deformation guidance across domains using a *real-style* image pair instead of only the style example, thus reducing the difficulty; see, *e.g.*, Fig. 1. Nonetheless, existing one-shot methods still failed to capture the deformation pattern because of the weaker structural guidance they employed; see Sec. 5 for more details.

In this work, we propose a novel stylization network trained on a single real-style image pair. The network is built upon a pre-trained StyleGAN, with additional spatial transformers appended to make the generator deformation-aware. To overcome the huge domain gap between the given paired references, we dive into the feature space of DINO [5], a self-supervised vision transformer. We present evidence for DINO in capturing robust and consistent structure semantics across real and style face domains, superior to the other popular ViTs. Therefore, we optimize our generator based on the DINO semantic guidance to learn the cross-domain structural deformation. Specifically, we compute a novel directional deformation loss that regularizes the direction vector in DINO space between the real and the style faces. Meanwhile, a new relative structure consistency upon self-similarities of DINO features is calculated to ensure the diversity of the target (style) domain, preventing overfitting. Finally, color alignment based on style-mixing techniques is applied to further guarantee the correctness of semantic correspondence. Experiments show that our method can accurately stylize facial images into artistic styles with strong exaggerations, both in appearance change and structure deformation, yet still maintaining a faithful identity to the input.

The main contributions are summarized as follows:

- We explore the feature space of DINO and discover its powerful structural/semantic representation both in real and style face domains.
- Based on DINO features, we propose two novel cross-domain losses to constrain the geometric deformation from real faces to artistic styles.
- We propose a novel deformable face stylization network, trained with only a single-paired real-style example. Extensive qualitative and quantitative comparisons with existing state-of-the-art demonstrate the effectiveness and superiority of our framework.



Figure 2. **GAN inversion.** Previous one-shot stylization methods like [45] and [6], build the cross-domain guidance by inverting the artistic style image into real face domain. But compared with the ground truth real face, current GAN inversion techniques [2, 32] still cannot give out a faithful mapping across unseen domains, which may mislead the structure deformation in the stylization.

## 2. Related Work

**Face Stylization.** Driven by deep learning techniques, face stylization has evolved rapidly. Since the seminal work of [10], methods upon neural style transfer have been developed for face stylization [13, 19, 20, 30, 34]. However, these general style transfer methods did not make use of the learned prior of generative models, *e.g.*, GANs [12]. In contrast, image translation-based methods [4, 11, 15, 22, 26, 28, 31, 37] usually train dedicated stylization GANs, or fine-tune pretrained StyleGANs [17, 18] over large collections of artistic facial images (from hundreds to thousands). Although some of these methods are deformation-aware, specifically for caricature portraits, the requirement of data amount makes them fail in styles defined by extremely limited examples, which is yet the most cases in real applications. Recently, researchers have addressed the few data challenge on domain adaption by introducing a series of regularizations in the fine-tuning [23, 24, 35, 41–43]. On face stylization, more specifically, a few works [6, 21, 40, 45] further use a single example as the style reference. Although they can generate well-stylized appearances, the one-shot methods tend to strictly preserve the face structures, while overlooking the non-local deformations in the style example. In this work, we present a novel framework for deformable face stylization based on a single paired data.

**ViT Feature Representation.** Features from vision transformer (ViT) [8] are powerful and versatile visual representations. Researchers have demonstrated that ViT trained in specific manners can serve well for numerous downstream tasks in both vision and graphics [3, 5, 27, 44]. Among them, CLIP [27] are most widely used for text-guided image editing [25], generation [9], and stylization [21, 40, 45]. Then, DINO [5], a self-supervised ViT model, exhibits striking properties in capturing high-level semantic information [3]. Amir *et al.* [3] use the keys of DINO as ViT features and apply them to many challenging vision tasks in unconstrained settings. Tumanyan *et al.* [33] further splice the DINO features as disentangled appearance and structure representa-

tions, realizing a semantic-aware appearance transfer from one natural image to another. Given the success achieved, we believe DINO has the potential for deformable face stylization. Next, we will explore the DINO features and compare them with two other ViT features.

### 3. DINO Semantic Guidance

In this section, we first briefly review the DINO-ViT [5]. Then, we explore and analyze the properties of different DINO features, and further compare them with two weakly supervised ViTs, CLIP [27] and FaRL [44].

DINO is a ViT model trained in self-distillation [5]. During training, an input image is randomly transformed into two different inputs for a student and a teacher network (with the same architecture but different parameters). The student is optimized by cross-entropy loss that measures the output similarity, and the exponential moving average of the student updates the teacher. DINO shows emerging properties in encoding high-level semantic information at high spatial resolution [3, 5]. We expect DINO to extract a consistent semantic representation across the source and target domains (*i.e.*, the real and the style faces). Thus, we can build reliable deformation guidance despite the huge domain gap.

Previously, Amir *et al.* [3] demonstrated that DINO is superior to traditional CNN-based models and supervised ViTs in semantic information capture. But they lack the comparison with alternative weakly-supervised ViTs, such as CLIP [27], which has been widely adopted in many style transfer tasks [9, 25]. To investigate comprehensively, we choose two recently popular weakly-supervised ViTs, CLIP and FaRL [44], which are both text-guided and learn a multi-modal embedding to estimate the semantic similarity between a given text and an image. In particular, FaRL is specially targeted for human facial representation.

So, are DINO features better than those of existing weakly-supervised ViTs (*i.e.*, CLIP and FaRL) for semantic representation in the task of face stylization? To this end, we visualize their intermediate features for comparison. As shown in Fig. 3, we input both real and artistic face images to these ViTs. Keys and tokens from different layers are visualized by principal component analysis (PCA). From the PCA results, high-level DINO features show cleaner, more precise, and, more importantly, more consistent facial segments than CLIP and FaRL features. Lower-level features of DINO contain semantic and positional information simultaneously. In contrast, features from CLIP and FaRL are relatively chaotic, demonstrating the superiority of DINO.

We conclude that DINO can well capture the facial semantics across different facial domains, whereas others cannot. We attribute the advantages to DINO’s self-distillation. The pair-wised augmentation on the training input makes DINO naturally focus on structural semantics. On the con-




Figure 3. **Visualization of the hierarchical features** from CLIP [27], FaRL [44], and DINO [5], where the same color represents the same semantics shared. Note that the same architecture ViT (ViT-B/16) is employed fairly for image encoding. We choose layers 3, 6, and 12 to represent different levels (L, M, H) of features. Following [3], we only use keys and tokens in ViTs, while discarding the queries and values. The [CLS] token is also discarded as it mainly encodes the visual appearance [33].

trary, text-guided ViTs lose precision in semantic capture, probably due to the ambiguity of natural language. Furthermore, different from [3, 33], we found using keys will also miss some structural semantics through a simple over-fitting test; see supplementary for more details and a comprehensive comparison. Therefore, we use DINO tokens as our feature representation, named DINO semantic guidance.

### 4. Method

As illustrated in Fig. 4, at the core of our framework is a deformation-aware generator  $G^t$  fine-tuned by two novel DINO-based losses and an adversarial loss. Specifically, we design  $G^t$  upon StyleGANv2 [18], with spatial transformers [14] appended for structural deformation. During training, we first sample a latent code  $w \in \mathbb{R}^{18 \times 512}$  in  $W+$  space, followed by style mixing for color alignment. Then, we feed the code to  $G^t$  and  $G^s$ , respectively, where  $G^s$  is

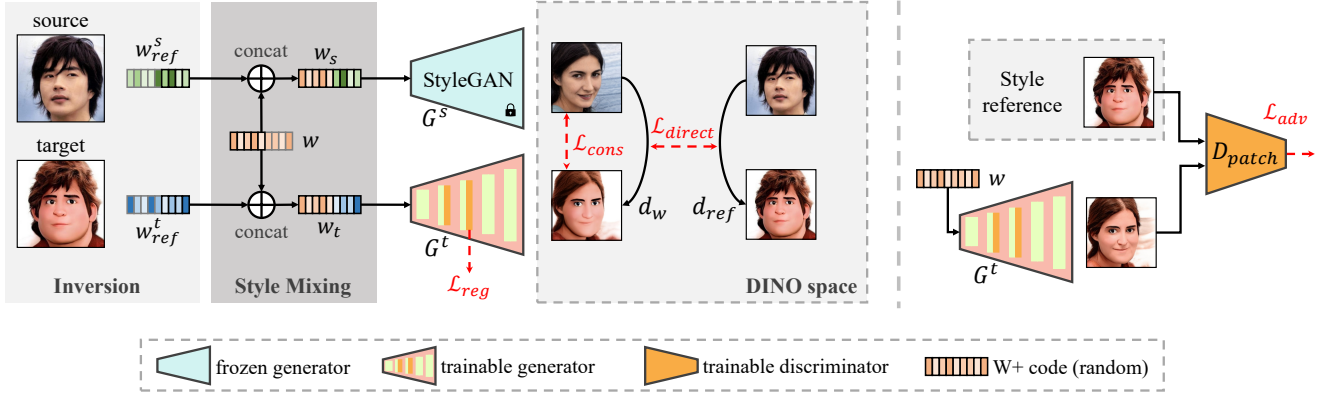


Figure 4. **Framework Overview.** Given a single real-style paired reference, we fine-tune a deformation-aware generator  $G^t$  that simultaneously realizes geometry deformation and appearance transfer. To learn the cross-domain deformation, we design a directional deformation loss  $\mathcal{L}_{direct}$  and a relative structural consistency loss  $\mathcal{L}_{cons}$ , both computed in DINO feature space (middle). Inversion and style mixing further ensure a consistent DINO semantic representation aligned with the given reference (left). In addition, we involve adversarial training using a patch-level discriminator to enhance the transferred style and fidelity (right).

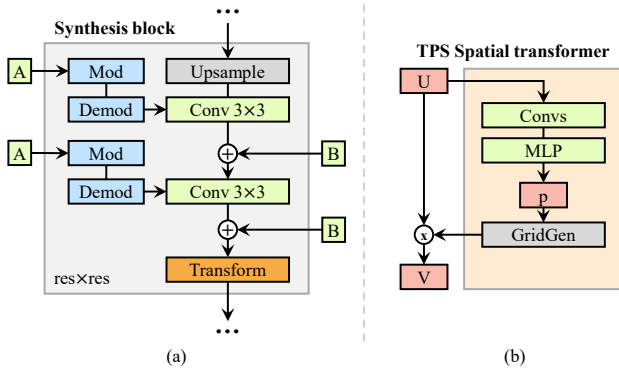


Figure 5. **Deformation-aware generator.** (a) A StyleGAN synthesis block plugged in with a Transform module. (b) A thin-plate-spline spatial transformer (TPS-STN) for feature warping, which shows the detail of the Transform module in (a).

a pre-trained StyleGAN generator frozen during training, and  $G^t$  is to be fine-tuned by the new proposed losses. In inference time, we first input the given facial images to a pre-trained e4e encoder [32], obtaining their inversion code  $w^*$  in W+ space. Then, we can generate the corresponding stylized faces  $G^t(w^*)$  by a single forward pass.

Below, we first describe the details of our deformation-aware generator, followed by the DINO-based domain adaption, where two novel constraints based on DINO semantic guidance are introduced. In the end, we describe the color alignment enhancing the stability.

#### 4.1. Deformation-aware generator

StyleGAN [18] pretrained on large-scale facial datasets, such as FFHQ [17]), is powerful in various facial image generation tasks. However, it is hard to synthesize faces with exaggerated components given the prior learned from

real face domain. Inspired by [38], we warp the intermediate features of StyleGAN to facilitate the generator to output the desired deformation pattern.

As depicted in Fig. 5, we insert a simple differentiable spatial transformer network (STN) [14], denoted as Transform, after the synthesis block of StyleGAN generator. The Transform module transforms a feature map by a single forward pass, where the transformation could be translation, rotation, thin-plate-spline warping (TPS), etc. In our case, we chose TPS-STN for warping feature maps, as well as a basic STN that shares the same architecture as TPS-STN to perform translation, rotation, and scaling. Empirically, we only add Transform in the resolutions of  $32 \times 32$  and  $64 \times 64$  with a grid size of 10.

We are not the first to make StyleGAN deformation-aware. StyleCariGAN [15] presented a layer-mixed StyleGAN with CNN-based exaggeration blocks. However, these deformation blocks require additional training with massive training data. Our STN-based Transform, instead, only performs simple transformations on feature maps, which can be directly plugged into the pretrained StyleGAN generator. According to [11], we regularize the TPS warping field to be smooth by

$$\mathcal{L}_{reg} = \sum_{i,j \in \mathbf{F}} \left( 2 - \text{sim}(\mathbf{F}_{i,j-1}, \mathbf{F}_{i,j}) - \text{sim}(\mathbf{F}_{i-1,j}, \mathbf{F}_{i,j}) \right), \quad (1)$$

where  $\text{sim}$  refers to the cosine similarity,  $\mathbf{F}$  represents the warping field, and  $i, j$  are pixel indices.

#### 4.2. DINO-based domain adaption

Our deformation-aware generator is initialized by a StyleGANv2 model pre-trained on FFHQ. To adapt it to the target domain defined by the style reference, we fine-

tune the generator with three criteria: 1) a novel directional deformation loss guiding the cross-domain structural deformation, 2) a newly designed relative structural consistency transferring the generation diversity, and 3) an adversarial loss facilitating the style synthesis.

**Directional Deformation Loss.** Based on the DINO semantic guidance in Sec. 3, we construct a directional deformation loss guiding the cross-domain structural deformation. As in Fig. 4, given a single paired reference, we first project them into DINO space and calculate the structural changing direction from source to target, represented by  $\mathbf{d}_{ref} = E_D(I_{ref}^t) - E_D(I_{ref}^s)$ , as the directional deformation reference. We expect the deformation direction of generated images across domains to align with the deformation reference. Akin to [9], we compute the DINO-based directional loss measured by cosine similarity between the cross-domain deformation directions, which is

$$\mathcal{L}_{direct} = 1 - \text{sim}\left(\frac{\mathbf{d}_w}{\|\mathbf{d}_w\|}, \frac{\mathbf{d}_{ref}}{\|\mathbf{d}_{ref}\|}\right), \quad (2)$$

where  $\mathbf{d}_w = E_D(G^t(w)) - E_D(G^s(w))$ , with  $E_D$  as the DINO encoder, and  $w$  is the latent code in  $W+$  space.

**Relative Structural Consistency.** As a few-shot training, adapting the generator only with the above directional deformation guidance will easily result in over-fitting. Inspired by [24], we introduce a relative structural consistency to preserve the structural diversity across domains. Specifically, we measure the difference/similarity between every two generated batch samples of the same domain. The self-similarity vector\* computed with DINO tokens is used as the structure representation for each sample, which is  $s_i = \frac{E_D(I_i)E_D(I_i)^T}{\|E_D(I_i)\|_2^2}$ . For the ground truth pair, we involve them in the relative structural consistency as well. Accordingly, two relative similarity matrices can be formed for the source and target domain, respectively, as depicted in Fig. 6. We then transform them into probability distributions by

$$\begin{aligned} C^s &= \text{softmax}\left(\left\{\text{sim}(s_i^s, s_j^s)\right\}_{i,j=1,\dots,N}^{\forall i>j}\right) \\ &\cup \text{softmax}\left(\left\{\text{sim}(s_i^s, s_{ref}^s)\right\}_{\forall i=1,\dots,N}\right), \\ C^t &= \text{softmax}\left(\left\{\text{sim}(s_i^t, s_j^t)\right\}_{i,j=1,\dots,N}^{\forall i>j}\right) \\ &\cup \text{softmax}\left(\left\{\text{sim}(s_i^t, s_{ref}^t)\right\}_{\forall i=1,\dots,N}\right) \end{aligned} \quad (3)$$

where  $\text{sim}$  is cosine similarity,  $N$  is the batch-size, superscripts  $s$  and  $t$  indicate the samples generated from  $G^s$  and  $G^t$ , respectively. To encourage the generated images of

\*The self-similarity between DINO tokens forms a large similarity matrix (in size of 783\*783). We flatten it into a self-similarity vector.




Figure 6. **Relative structural consistency.** We compute a relative structural similarity matrix for each domain, where pairs of generated-reference (in blue) and generated-generated (in green) are both considered. A well-designed loss  $\mathcal{L}_{cons}$  is dedicated to preserving similar structural diversity between the two domains.

target domain to have a similar structural diversity as the source, we compute the MSE loss between them, which is

$$\mathcal{L}_{cons} = \frac{1}{|C|} \sum_{i=1}^M \|C_i^t - C_i^s\|_2^2. \quad (4)$$

**Adversarial Style Transfer.** Given the DINO features employed, the above two components mainly correspond to structural deformation. To guarantee a correct color transfer for face stylization, we introduce a patch-level discriminator  $D_{patch}$  according to [24]. Unlike image-level discriminators, the patch-level discriminator focuses more on local texture and color. Meanwhile, we found that a discriminator also helps to improve the fidelity of the generated stylized faces, thanks to the adversarial training. Following [24], we finetune a StyleGANv2 discriminator pre-trained on FFHQ, and read off the layers with effective patch size of  $22 \times 22$ . The adversarial losses for  $D_{patch}$  and our deformation-aware generator  $G^t$  are given by:

$$\begin{aligned} \mathcal{L}_{adv}^D &= \log(1 - D_{patch}(I_{ref}^t)) + \mathbb{E}_w[\log(D_{patch}(G^t(w)))], \\ \mathcal{L}_{adv}^G &= -\mathbb{E}_w[\log(D_{patch}(G^t(w)))]. \end{aligned} \quad (5)$$

The total loss function can be finalized as:

$$\mathcal{L}_{total} = \mathcal{L}_{adv} + \lambda_{direct}\mathcal{L}_{direct} + \lambda_{cons}\mathcal{L}_{cons} + \lambda_{reg}\mathcal{L}_{reg}. \quad (6)$$

### 4.3. Color alignment

Even though DINO well disentangles the structure and appearance of images in its feature space, the vast color variations in the real facial domain would bring undesired errors in structure matching. To alleviate this disturbance, we use the inversion and style mixing technique of StyleGAN to align the samples' color. As shown in the leftmost of Fig. 4, we first inverse the paired reference ( $I_{ref}^s, I_{ref}^t$ ) into the latent space of source domain by optimizing from

a mean latent  $\bar{w}$  using L1 and perceptual loss [39], obtaining their latents  $w_{ref}^s, w_{ref}^t \in \mathbb{R}^{18 \times 512}$ . Then, for a random sample  $w$  in  $W+$  space, unlike [6, 21, 45], we swap the fine-level (9-18th) codes of  $w$  with the corresponding codes of  $w_{ref}^s$  and  $w_{ref}^t$ , resulting in the latent codes  $w_s, w_t$  with aligned color to the references. The generated images after color alignment share the same structure/identity as the original random sample but have aligning colors to the corresponding reference. See supplementary for more details.

## 5. Experiments

**Implementation Details.** Our framework is built upon the StyleGANv2 [18] and initialized by the model pre-trained on FFHQ [17]. We use  $\lambda_{direct} = 6$ ,  $\lambda_{cons} = 5e4$  and  $\lambda_{reg} = 1e-6$ , empirically. For training, we set the batch size to 4 and used ADAM optimizer with a learning rate 0.002 for the generator following the existing one-shot methods [6, 21, 40, 45]. For the multi-resolution STN modules, we set the learning rate of TPS-STN and basic-STN to  $5e-6$  and  $1e-4$ , respectively, which are different from the generator, but they are fine-tuned together. All the STNs appended consist of two convolution layers followed by two linear layers. We use a mixture of M- and H-level features of DINO for the computation of the directional deformation guidance and only M-level features for the relative structural consistency, as we expect to keep more positional information while changing the semantic contents. All experiments are performed using a single NVIDIA RTX 3090.

For simplicity, the paired references we used are mainly obtained from existing generative models that are trained with massive training data, such as [1, 15, 31, 37]. These models can stylize facial images with convincing quality. Their produced results are enough to represent real pairs.

### 5.1. Evaluation and comparisons

To validate the effectiveness of our framework, we conducted qualitative and quantitative comparisons to several baselines, including MTG [45], JoJoGAN [6], DiFa [40], and OneshotCLIP [21]. For fairness, we also convert MTG and JoJoGAN into paired one-shot methods for comparison.

**Qualitative Comparison.** Fig. 7 shows the results of stylizing two real portraits into three different styles using different adaption methods. Due to the unfaithful inversion of style reference, the compared one-shot methods fail to capture the exaggerated components, *e.g.*, the face contour in all three examples. In contrast, our method stylizes the faces with plausible exaggeration and correct color transfer.

We also convert MTG and JoJoGAN to accept paired references in training; see supplementary for the details of these two variants (MTG-pair and JoJoGAN-pair). Fig. 8 shows the qualitative comparison. Again, our results have

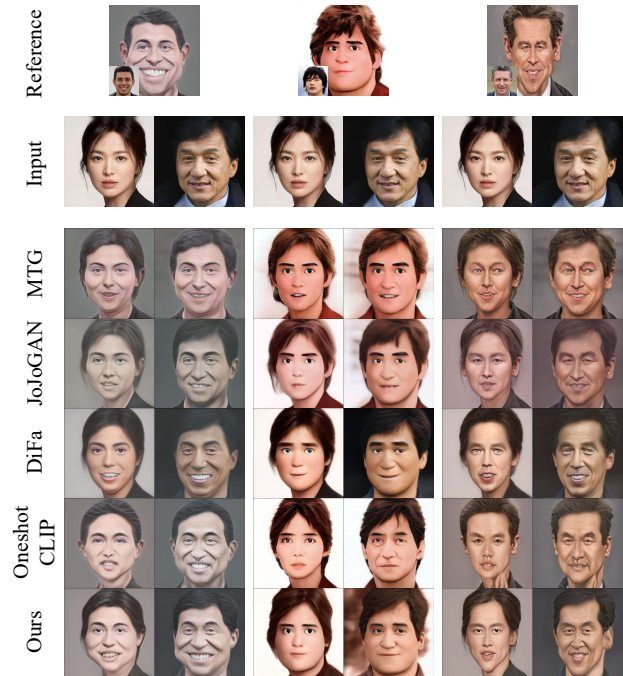


Figure 7. **Qualitative comparison with existing one-shot face stylization methods**, MTG [45], JoJoGAN [6], DiFa [40] and Oneshot-CLIP [21]. Note all these competitors were trained using only the style example, whereas we used paired reference.

better color style and finer geometry deformation. The two variants are still limited to generating deformed faces due to the lack of precise cross-domain structure guidance. More comparisons are included in the supplementary.

**Quantitative Comparison.** We evaluate the generated results from three aspects: perception, deformation, and identity. For perceptual evaluation, we compute the widely used LPIPS distance [39]. To evaluate structural deformation and identity, we designed two new metrics: directional content consistency (dir-CC) and directional identity similarity (dir-ID). For both metrics, we first compute a directional feature vector for each pair between the source and the target domain. As the reference pair provides the ground truth directional vector, we measure the cosine similarity between the ground truth pair and the generated pairs. Given that we have already used DINO features in our training, for fairness, we employ VGG [29] and ArcFace [7] features for these two metrics.

Table 1 lists the full quantitative evaluation. Specifically, our method shows an obvious advantage in deformation and identity compared with the one-shot methods trained with only the style reference. While compared with the “paired” variants, our method still leads ahead on most metrics, especially style perception and structure deformation.

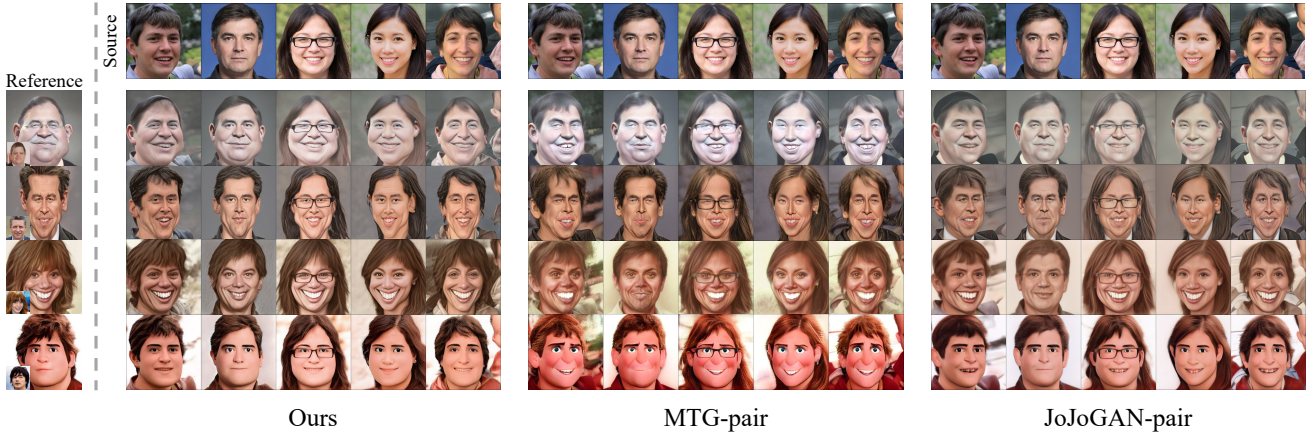


Figure 8. **Qualitative comparison with the “paired” variants** of MTG [45] and JoJoGAN [6]. We replace the inversion of the style example they used with the ground-truth real face. Our method surpasses these two variants both in color style and geometry deformation.

Table 1. **Quantitative comparison.** We choose the three styles shown in Fig. 7 for evaluation, denoted as Example-1, Example-2 and Example-3. Part I shows the comparison with existing one-shot methods trained with only the style reference, whereas Part II lists the comparison with the “paired” variants of MTG [45] and JoJoGAN [6]. Note the best results are **bold**, and the second-bests are underlined.

|    | Method           | Example-1    |              |              | Example-2    |              |              | Example-3    |              |              |
|----|------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
|    |                  | LPIPS ↓      | dir-CC ↑     | dir-ID ↑     | LPIPS ↓      | dir-CC ↑     | dir-ID ↑     | LPIPS ↓      | dir-CC ↑     | dir-ID ↑     |
| I  | MTG [45]         | <b>0.340</b> | 0.123        | 0.264        | <u>0.361</u> | 0.111        | 0.301        | <b>0.328</b> | 0.095        | 0.254        |
|    | JoJoGAN [6]      | 0.357        | 0.149        | 0.290        | 0.385        | <u>0.138</u> | <u>0.378</u> | 0.355        | 0.097        | <b>0.382</b> |
|    | DiFa [40]        | 0.362        | 0.142        | 0.223        | <b>0.293</b> | 0.137        | 0.328        | 0.346        | 0.094        | 0.123        |
|    | OneshotCLIP [21] | 0.452        | <u>0.157</u> | <u>0.333</u> | 0.449        | 0.119        | 0.113        | 0.460        | <u>0.136</u> | 0.201        |
|    | Ours             | <u>0.353</u> | <b>0.196</b> | <b>0.468</b> | 0.379        | <b>0.194</b> | <b>0.379</b> | <u>0.342</u> | <b>0.191</b> | <u>0.330</u> |
| II | MTG-pair         | 0.423        | <u>0.195</u> | 0.392        | 0.447        | <u>0.174</u> | 0.335        | <u>0.379</u> | 0.159        | 0.280        |
|    | JoJoGAN-pair     | <u>0.381</u> | 0.180        | <u>0.405</u> | <u>0.395</u> | 0.155        | <b>0.393</b> | 0.383        | <u>0.173</u> | <b>0.400</b> |
|    | Ours             | <b>0.353</b> | <b>0.196</b> | <b>0.468</b> | <b>0.379</b> | <b>0.194</b> | <u>0.379</u> | <b>0.342</b> | <b>0.191</b> | <u>0.330</u> |

Table 2. **User preference score.** Each user answered 60 questions, and each question asked to compare our method to an opponent on 3 real faces. Finally, 30 participants were involved, resulting in 1800×3 comparisons. See supplementary for more details.

|          |          |              |                  |
|----------|----------|--------------|------------------|
| Ours vs. | MTG [45] | JoJoGAN [6]  | OneshotCLIP [21] |
| Rate     | 73.0%    | 77.0%        | 83.7%            |
| Ours vs. | MTG-pair | JoJoGAN-pair | DiFa [40]        |
| Rate     | 76.7%    | 70.7%        | 71.3%            |

**User Study.** We further conducted a user study to investigate human evaluation of the results. Each time, we show the user the paired reference, three real face images, three results from our method, and three results from a competitor. The user is asked to choose the better results group or “comparable”. Table 2 reports the user preference score of our method against the opponents.

## 5.2. Ablation Analysis

**Effect of loss terms.** As shown in Fig. 9, the directional deformation loss  $\mathcal{L}_{direct}$  brings reliable deformation guidance but, meanwhile, leads to over-fitting. The relative structural

consistency  $\mathcal{L}_{cons}$  ensures a correct correspondence across domains but introduces color variations too. The adversarial loss stabilizes the style quality. Therefore, our full loss terms guarantee the deformation and style simultaneously.

**Effect of STN blocks.** As shown in Fig. 11, the generator with STN blocks brings a more accurate deformation effect regarding the desired exaggerations shown in the style reference, thus improving the stylization quality, as exaggeration is a key component in artistic styles.

**Effect of color alignment.** Fig. 12 verifies the effect of color alignment. When a huge color gap exists in the paired reference, color alignment ensures a consistent DINO semantic representation, alleviating structural mismatching.

## 5.3. Facial deformation control

Since the STNs are plug-in blocks that warp the intermediate features of StyleGAN, we can use them to control the feature warping degree by interpolation. Specifically, we interpolate the warping field of TPS-STNs by a weight  $\alpha$ :

$$\mathbf{F}(\alpha) = (1 - \alpha) * \mathbf{F}_0 + \alpha * \mathbf{F}, \quad (7)$$

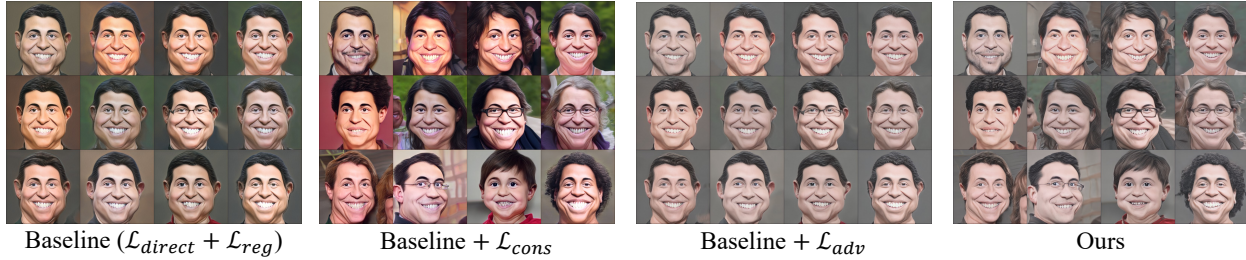


Figure 9. **Ablation of different loss terms.** Using only  $\mathcal{L}_{direct} + \mathcal{L}_{reg}$  makes the generated results all look similar to the reference. Adding  $\mathcal{L}_{cons}$  improves the diversity but still lacks a correct color style regarding the reference, which is rectified after  $\mathcal{L}_{adv}$  is introduced.

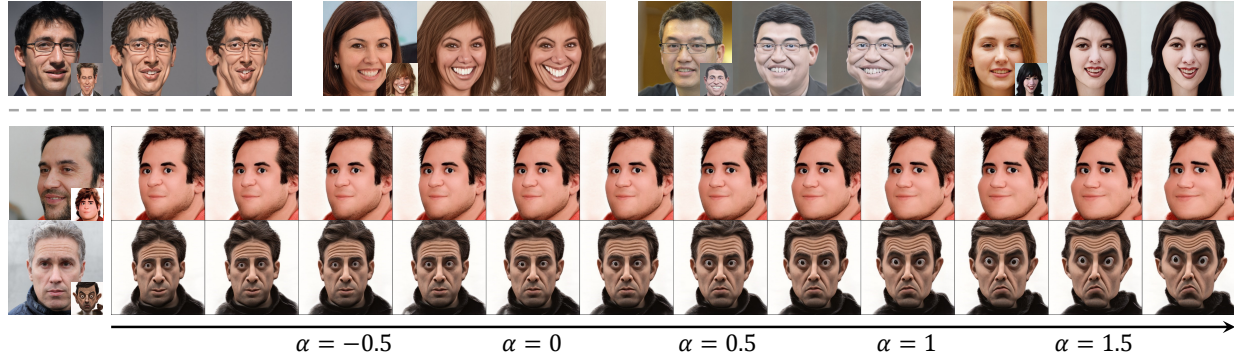


Figure 10. **Controllable face deformation.** The STN blocks are functioned as plug-ins. We linearly interpolate the warping field of TPS-STNs to control the deformation degree of stylized faces. The triplets in the top row show the real faces, the stylized faces without and with STN-deformation, respectively. At the bottom, we showcase two stylized faces the deformation control of different degrees.




Figure 11. **Ablation of the STNs appended in generator.** We fine-tune the generator with and without STN blocks. Some key exaggerated components, such as the eye contraction (left) and the cheek deformation (right), are missed when removing STN blocks.

where  $\mathbf{F}_0$  is the original warping field without any warping. Setting  $\alpha=1$  or  $\alpha=0$ , respectively, leads to results with or without the warping deformations. Fig. 10 shows examples of such controllable face deformation in different styles.

## 6. Conclusion

We have presented a deformable face stylization framework using only a single paired reference. By fine-tuning a deformation-aware generator under DINO semantic guidance, we can stylize facial images with high-quality appearance transfer and convincing structure deformation. Qual-

Figure 12. **Ablation of color alignment.** Color alignment (C.A.) plays an important role in ensuring correct semantic matching when color gaps exist in the paired reference. Note how the lower lips (highlighted) are influenced when there is no color alignment.

itative and quantitative comparisons demonstrate that our method surpasses the existing one-shot methods.

**Acknowledgements** This work was supported in parts by NSFC (U21B2023, U2001206, 62161146005), GD Natural Science Foundation (2022A1515010221), DEGP Innovation Team Program (2022KCXTD025), Shenzhen Science and Technology Program (KQTD20210811090044003, RCJC20200714114435012), Guangdong Laboratory of Artificial Intelligence and Digital Economy (SZ) and Scientific Development Funds of Shenzhen University.



## References

- [1] Rameen Abdal, Hsin-Ying Lee, Peihao Zhu, Menglei Chai, Aliaksandr Siarohin, Peter Wonka, and S. Tulyakov. 3davatargan: Bridging domains for personalized editable avatars. In *Proc. of IEEE Conf. on Computer Vision & Pattern Recognition*, pages 4552–4562, 2023. 6
- [2] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space?, 2019. 2
- [3] Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. Deep vit features as dense visual descriptors. *ECCVW What is Motion For?*, 2022. 2, 3
- [4] Kaidi Cao, Jing Liao, and Lu Yuan. Carigans: Unpaired photo-to-caricature translation, 2018. 2
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proc. of Int. Conf. on Computer Vision*, 2021. 2, 3
- [6] Min Jin Chong and David Forsyth. Jojogan: One shot face stylization. In *Proc. of Euro. Conf. on Computer Vision*, 2022. 1, 2, 6, 7
- [7] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proc. of IEEE Conf. on Computer Vision & Pattern Recognition*, pages 4690–4699, 2019. 6
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proc. of Int. Conf. on Learning Representations*, 2021. 2
- [9] Rinon Gal, Or Patashnik, Haggai Maron, Amit H. Bermano, Gal Chechik, and Daniel Cohen-Or. Stylegan-nada: Clip-guided domain adaptation of image generators. *ACM Trans. on Graphics (Proc. of SIGGRAPH)*, 41(4), jul 2022. 2, 3, 5
- [10] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proc. of IEEE Conf. on Computer Vision & Pattern Recognition*, pages 2414–2423, Las Vegas, NV, USA, June 2016. IEEE. 2
- [11] Julia Gong, Yannick Hold-Geoffroy, and Jingwan Lu. Autotoon: Automatic geometric warping for face cartoon generation. *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 349–358, 2020. 2, 4
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. 2
- [13] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proc. of Int. Conf. on Computer Vision*, 2017. 2
- [14] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and koray kavukcuoglu. Spatial transformer networks. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. 3, 4
- [15] Wonjong Jang, Gwangjin Ju, Yucheol Jung, Jiaolong Yang, Xin Tong, and Seungyong Lee. Stylecarigan: Caricature generation via stylegan feature map modulation. *ACM Trans. on Graphics (Proc. of SIGGRAPH)*, 40(4), 2021. 2, 4, 6
- [16] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *Advances in Neural Information Processing Systems*, 2020. 1
- [17] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks, 2019. 1, 2, 4, 6
- [18] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proc. of IEEE Conf. on Computer Vision & Pattern Recognition*, 2020. 1, 2, 3, 4, 6
- [19] Sunnie S. Y. Kim, Nicholas Kolkin, Jason Salavon, and Gregory Shakhnarovich. Deformable style transfer. In *Proc. of Euro. Conf. on Computer Vision*, 2020. 2
- [20] N. Kolkin, J. Salavon, and G. Shakhnarovich. Style transfer by relaxed optimal transport and self-similarity. In *Proc. of IEEE Conf. on Computer Vision & Pattern Recognition*, pages 10043–10052, Los Alamitos, CA, USA, jun 2019. IEEE Computer Society. 2
- [21] Gihyun Kwon and Jong Chul Ye. One-shot adaptation of gan in just one clip. *IEEE Trans. Pattern Analysis & Machine Intelligence*, 45(10):12179–12191, 2023. 1, 2, 6, 7
- [22] Yifang Men, Yuan Yao, Miaomiao Cui, Zhouhui Lian, and Xuansong Xie. Dct-net: Domain-calibrated translation for portrait stylization. *ACM Trans. on Graphics*, 41(4):1–9, 2022. 1, 2
- [23] Arnab Kumar Mondal, Piyush Tiwary, Parag Singla, and Prathosh AP. Few-shot cross-domain image generation via inference-time latent-code learning. In *Proc. of Int. Conf. on Learning Representations*, 2023. 2
- [24] Utkarsh Ojha, Yijun Li, Cynthia Lu, Alexei A. Efros, Yong Jae Lee, Eli Shechtman, and Richard Zhang. Few-shot image generation via cross-domain correspondence. In *Proc. of IEEE Conf. on Computer Vision & Pattern Recognition*, 2021. 2, 5
- [25] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *Proc. of Int. Conf. on Computer Vision*, pages 2085–2094, October 2021. 2, 3
- [26] Justin N. M. Pinkney and Doron Adler. Resolution dependent gan interpolation for controllable image synthesis between domains, 2020. 2
- [27] 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, 2021. 2, 3
- [28] Yichun Shi, Deb Debayan, and Anil K. Jain. Warpgan: Automatic caricature generation. In *Proc. of IEEE Conf. on Computer Vision & Pattern Recognition*, 2019. 2

- [29] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *Proc. of Int. Conf. on Learning Representations*, 2015. 6
- [30] Chunjin Song, Zhijie Wu, Yang Zhou, Minglun Gong, and Hui Huang. Etnet: Error transition network for arbitrary style transfer. In *Advances in Neural Information Processing Systems*, pages 668–677, 2019. 2
- [31] Guoxian Song, Linjie Luo, Jing Liu, Wan-Chun Ma, Chunpong Lai, Chuanxia Zheng, and Tat-Jen Cham. Agilegan: Stylizing portraits by inversion-consistent transfer learning. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, jul 2021. 1, 2, 6
- [32] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *arXiv preprint arXiv:2102.02766*, 2021. 2, 4
- [33] Narek Tumanyan, Omer Bar-Tal, Shai Bagon, and Tali Dekel. Splicing vit features for semantic appearance transfer. In *Proc. of IEEE Conf. on Computer Vision & Pattern Recognition*, pages 10748–10757, 2022. 2, 3
- [34] Zhijie Wu, Chunjin Song, Yang Zhou, Minglun Gong, and Hui Huang. Efanet: Exchangeable feature alignment network for arbitrary style transfer. *Proc. AAAI Conf. on Artificial Intelligence*, pages 12305–12312, 4 2020. 2
- [35] Jiayu Xiao, Liang Li, Chaofei Wang, Zheng-Jun Zha, and Qingming Huang. Few shot generative model adaption via relaxed spatial structural alignment. In *Proc. of IEEE Conf. on Computer Vision & Pattern Recognition*, pages 11204–11213, June 2022. 2
- [36] Shuai Yang, Liming Jiang, Ziwei Liu, , and Chen Change Loy. Styleganex: Stylegan-based manipulation beyond cropped aligned faces. In *Proc. of Int. Conf. on Computer Vision*, 2023. 1
- [37] Shuai Yang, Liming Jiang, Ziwei Liu, and Chen Change Loy. Pastiche master: Exemplar-based high-resolution portrait style transfer. In *Proc. of IEEE Conf. on Computer Vision & Pattern Recognition*, 2022. 1, 2, 6
- [38] Shuai Yang, Liming Jiang, Ziwei Liu, and Chen Change Loy. Vtoonify: Controllable high-resolution portrait video style transfer. *ACM Trans. on Graphics (Proc. of SIGGRAPH Asia)*, 41(6):1–15, 2022. 1, 4
- [39] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. of IEEE Conf. on Computer Vision & Pattern Recognition*, 2018. 6
- [40] Yabo Zhang, mingshuai Yao, Yuxiang Wei, Zhilong Ji, Jinfeng Bai, and Wangmeng Zuo. Towards diverse and faithful one-shot adaption of generative adversarial networks. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. 1, 2, 6, 7
- [41] Yunqing Zhao, Keshigeyan Chandrasegaran, Milad Abdollahzadeh, and Ngai man Cheung. Few-shot image generation via adaptation-aware kernel modulation. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. 2
- [42] Yunqing Zhao, Henghui Ding, Houjing Huang, and Ngai-Man Cheung. A closer look at few-shot image generation. *Proc. of IEEE Conf. on Computer Vision & Pattern Recognition*, pages 9130–9140, 2022. 2
- [43] Yunqing Zhao, Chao Du, Milad Abdollahzadeh, Tianyu Pang, Min Lin, Shuicheng Yan, and Ngai-Man Cheung. Exploring incompatible knowledge transfer in few-shot image generation. In *Proc. of IEEE Conf. on Computer Vision & Pattern Recognition*, pages 7380–7391, 2023. 2
- [44] Yinglin Zheng, Hao Yang, Ting Zhang, Jianmin Bao, Dongdong Chen, Yangyu Huang, Lu Yuan, Dong Chen, Ming Zeng, and Fang Wen. General facial representation learning in a visual-linguistic manner. In *Proc. of IEEE Conf. on Computer Vision & Pattern Recognition*, pages 18676–18688, New Orleans, LA, USA, June 2022. IEEE. 2, 3
- [45] Peihao Zhu, Rameen Abdal, John Femiani, and Peter Wonka. Mind the gap: Domain gap control for single shot domain adaptation for generative adversarial networks. In *Proc. of Int. Conf. on Learning Representations*, 2022. 1, 2, 6, 7