

StyO: Stylize Your Face in Only One-Shot

Bonan Li^{1*}, Zicheng Zhang^{1*}, Xuecheng Nie², Congying Han^{1†},
Yinhan Hu¹, Xinmin Qiu¹, Tiande Guo¹

¹ University of Chinese Academy of Sciences

² MT Lab, Meitu Inc.

{libonan, hancy, tdguo}@ucas.ac.cn, {zhangzicheng19, huyinhan16, qiuxinmin21}@mails.ucas.ac.cn, nxc@meitu.com

Abstract

This paper focuses on face stylization with a single artistic target. Existing works for this task often fail to retain the source content while achieving geometry variation. Here, we present a novel StyO model, *i.e.* Stylize the face in only **One**-shot, to solve the above problem. In particular, StyO exploits a disentanglement and recombination strategy. It first disentangles the content and style of source and target images into identifiers, which are then recombined in a cross manner to derive the stylized face image. In this way, StyO decomposes complex images into independent and specific attributes, and simplifies one-shot face stylization as the combination of different attributes from input images, thus producing results better matching face geometry of target image and content of source one. StyO is implemented with latent diffusion models (LDM) and composed of two key modules: 1) Identifier Disentanglement Learner (IDL) for disentanglement phase. It represents identifiers as contrastive text prompts, *i.e.* positive and negative descriptions. And it introduces a novel triple reconstruction loss to fine-tune the pre-trained LDM for encoding style and content into corresponding identifiers; 2) Fine-grained Content Controller (FCC) for recombination phase. It recombines disentangled identifiers from IDL to form an augmented text prompt for generating stylized faces. In addition, FCC also constrains the cross-attention maps of latent and text features to preserve source face details in results. The extensive evaluation shows that StyO produces high-quality images on numerous paintings of various styles and outperforms the current state-of-the-art.

Introduction

Face stylization aims at automatically creating personalized artistic portraits from face photographs. It is a fundamental problem in vision and graphics communities (Akleman 1997; Kim et al. 2019a; Han et al. 2021; Song et al. 2021) and widely applied in creative industries such as social media avatars, films, advertising, etc. Recently, one-shot solution for this task has drawn lots of attention, due to its significant feature of data-efficiency, which makes the deployment easier in practice.

*Equal contribution

†Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Existing works (Chong and Forsyth 2022; Zhu et al. 2021; Zhang et al. 2022b,a) consider one-shot face stylization as a style transfer problem (Gatys, Ecker, and Bethge 2016), *e.g.*, transferring face style of a single target image into face of the source image, and solve this problem with Generative Adversarial Networks (GANs) (Goodfellow et al. 2014; Karras, Laine, and Aila 2019). Despite of their success, we observe that even state-of-the-art methods still fail to achieve reasonable style (*e.g.*, geometry variations) while maintain source content (*e.g.*, facial and hair colors), as shown in Figure 1. Note that these features are delineated as fundamental attributes warranting consideration (Cao, Liao, and Yuan 2018; Gong, Hold-Geoffroy, and Lu 2020; Shi, Deb, and Jain 2019; Men et al. 2022; Han et al. 2021; Song et al. 2021; Liu et al. 2021a). This is caused by two main reasons. First, these methods heavily relied on pre-trained GAN models on large real face dataset (Karras, Laine, and Aila 2019), thus hard to match styles of faces, especially geometry changes, from another domain with one-shot guidance, due to capability limitation of GANs for producing distribution variations. Second, they entangle the style and content information of facial images together in the latent space, making current encoding and inversion techniques difficult to derive accurate latent code representing the stylized faces. It is important to address above issues for facilitating usage of relevant techniques in real world.

Motivated by this, we propose a novel StyO model in this paper for pushing forward the frontier of the task of Stylizing the face in only **One**-shot. For improving the generation variety, StyO exploits Diffusion Denoising Probabilistic Models (DDPMs) (Ho, Jain, and Abbeel 2020; Song, Meng, and Ermon 2020; Rombach et al. 2022) instead of GANs, given the superiority of DDPM to produce high-quality images in wild range of distribution. For deriving accurate representation of stylized faces, StyO leverages a disentanglement and recombination strategy. Specifically, StyO factorizes the style and content information of source and target images into different identifiers, then fuse them in a cross manner to form prompts, which plays as representations of stylized faces to drive DDPMs to generate reasonable images. In this way, StyO decomposes complex images into independent and specific attributes, and simplifies one-shot face stylization as combination of different attributes from input images, thus producing results better matching




Figure 1: Comparison of the proposed StyO and previous state-of-the-art work JoJoGAN (Chong and Forsyth 2022) for one-shot face stylization, which aims at applying the style of a single target image to the source one. We can see StyO obviously outperforms JoJoGAN for achieving suitable style from target image, *e.g.*, texture and geometry variations, while maintaining fine-grained content from source image, *e.g.*, head pose and hair style. Better see in color and 2x zoom.

face style of target image and content of source one.

In particular, StyO is implemented with the efficient Latent Diffusion Models (LDMs) (Rombach et al. 2022). It is composed of two core modules: 1) Identifier Disentanglement Learner (IDL). IDL aims to disentangle the style and content of an image into different identifiers. Here, IDL represents identifiers as text descriptors for fully leveraging the powerful capability of text-based image generation of LDMs. Then, IDL defines contrastive disentangled prompt template as text descriptions for input images with content and style identifiers, *e.g.*, “a drawing with [Source/Target Style Identifier][not Target/Source Style Identifier] style of [Source/Target Content Identifier][not Target/Source Content Identifier] of portrait” for source and target images, respectively. In addition, we introduce an auxiliary prompt template with only source and target style identifiers to describe an auxiliary image set with same style of source one. This helps style and content identifiers correctly to represent corresponding attributes of images, and also avoid the risk of relating style/content information with other words in prompts rather than corresponding identifiers. Given prompts defined above, StyO builds text-image pairs to fine-tune the pre-trained LDMs for injecting attributes of images into identifiers, thus achieving the disentanglement goal. 2) Fine-grained Content Controller (FCC). FCC aims to recombine style and content identifiers from IDL to generate stylized face images. Specifically, FCC constructs prompt template as “a drawing with [Target Style Identifier][not Source Style Identifier] style of [Source Content Identifier][not Target Content Identifier] of portrait” to describe stylized faces with style from target and content from source, which is used as condition to LDMs for deriving stylization results. However, we find that only with the above

reconstructed prompt will cause the loss of fine-grained details of source image, such as head pose, hair color, beard style, etc, due to the randomness feature of diffusion models yielding less control on diversity. To solve this issue, FCC presents a new manipulation mechanism on attention maps to improve controllability. This is inspired by the fact that attention maps from cross attention layers are semantically correlated with texts in prompts. Hence, FCC extracts attention maps for source content and uses them to replace that for stylized one, which effectively controls fine-grained details of results. In addition, FCC proposes to augment prompts by repeating identifiers. This simple augmentation strategy further successfully advances generation quality. By this design, StyO can produce faces with suitable target style while maintaining source content.

Extensive experiments show that the proposed StyO model can produce surprisingly good quality of facial images on various styles in the one-shot manner. It also outperforms state-of-the-art models quantitatively. Our contributions are summarized in four folds: First, to the best knowledge, we are the first to deploy diffusion models for tackling one-shot face stylization; Second, we present a novel StyO model with disentanglement and recombination strategy, which effectively solves core issues of prior works; Third, we present a classical way to encoding different attributes of images into multiple identifiers; Fourth, our StyO model sets new state-of-the-art for one-shot face stylization.

Related Work

Diffusion model recently breaks the long-time dominance of generative adversarial networks (GANs) (Goodfellow et al. 2014) on generative tasks. Since Sohl-Dickstein et al. (2015) first propose the denoising diffusion probabilistic models

(DDPM) (Ho, Jain, and Abbeel 2020; Nichol and Dhariwal 2021; Dhariwal and Nichol 2021) to form the image generation as a Markov Process, the community is inspired to investigate the generative model from non-Markov process (Song, Meng, and Ermon 2020), score matching (Song and Ermon 2020), and stochastic differential equation (Song et al. 2020). Besides, many sampling techniques are proposed to speedup (Song, Meng, and Ermon 2020; Liu et al. 2022; Karras et al. 2022) or condition (Dhariwal and Nichol 2021; Ho and Salimans 2022) the inference. Recent text-to-image diffusion models (Nichol et al. 2021; Ramesh et al. 2022; Rombach et al. 2022; Saharia et al. 2022) trained on extremely large-scale data have exhibited remarkable ability to generate diverse, creative images controlled by text prompts. (Gal et al. 2022a; Ruiz et al. 2022; Mokady et al. 2022; Kawar et al. 2022; Hertz et al. 2022; Zhang et al. 2024) shows that these models can edit images provided by users with a text prompt effectively, thereby enabling text-based stylization. However, compared with the fine-grained details provided by an artistic exemplar, the text prompt only provides a rough description of desirable style, thereby it is hard to obtain acceptable results with previous methods. To our best knowledge, StyO is the first to explore the exemplar-guided fine-style transfer in context of text-to-image diffusion model.

Face stylization has been widely studied in the field of Neural Style Transfer (NST), image-to-image translation (I2IT), and GAN adaptation. NST methods match the feature distribution (Gatys, Ecker, and Bethge 2016; Li et al. 2017; Pan et al. 2019; Kolkin, Salavon, and Shakhnarovich 2019; Kalischek, Wegner, and Schindler 2021) to transfer artistic texture. These methods can achieve pleasing results when the artistic portraits have an intense color or texture pattern, but usually underperform in styles involving significant geometric distortion of facial features, such as anime and cartoons. I2IT methods (Isola et al. 2017; Zhu et al. 2017; Liu, Breuel, and Kautz 2017; Wang et al. 2018; Kim et al. 2019b) learn a mapping from a large-scale dataset to translate an input image from a source domain to a target domain. Men et al. (2022) design specific modules and losses based on facial knowledge (e.g., landmarks, expression) to facilitate the face stylization. These methods can achieve reasonable texture and geometric variation when the training dataset includes more than hundreds of artistic portraits. However, collecting such extensive images for a specific style is not non-trivial, thus limiting the usage of I2IT in practice. Recently, the works about domain adaptation of GANs (Ojha et al. 2021) provide a creative way to solve the task. Since the pre-trained StyleGAN (Karras, Laine, and Aila 2019) forms a generative prior of natural faces, an artistic portrait, which belongs to the proximal domain of natural face, can guide the adaptation of StyleGAN (Zhu et al. 2021; Gal et al. 2022b; Chong and Forsyth 2022; Zhang et al. 2022b) to generate artistic portraits. Thereafter, given a face photograph, a latent code can be found by inversion techniques (Xia et al. 2021) and fed into adapted StyleGAN to generate the stylized image. However, the strong facial prior of pre-trained StyleGAN limits the geometric variation, and compressing face images into latent code also drops the content partially.

Hence these methods perform not well when implementing artistic images of large geometric deformations and face photographs out of the facial prior. In contrast, StyO properly utilizes the potent generative and reconstructive abilities of LDM to alleviate these problems.

Methodology

Preliminaries

Text-to-image diffusion models (Rombach et al. 2022; Saharia et al. 2022) are a family of generative models to map a noise latent ϵ and a text prompt \mathcal{P} to an image x_0 . In this paper, we focus on the pre-trained latent diffusion model (LDM) (Rombach et al. 2022), which is widely used as a base model in recent research. LDM first encodes the image x^0 into latent space with a pre-trained encoder \mathcal{E} , i.e., $z^0 = \mathcal{E}(x^0)$, while x^0 can be restored by means of a pre-trained decoder \mathcal{D} . It then performs the destructive and reverse processes in the latent space. Specifically, the destructive process is known as a fixed-length Markov chain starting from z^0 to z^T with a Gaussian transition at each step. For each $t \in \{0, \dots, T\}$, the transition holds that

$$z^t = \sqrt{\alpha^t} z^0 + \sqrt{1 - \alpha^t} \epsilon, \quad \epsilon \sim N(0, I), \quad (1)$$

where $\{\alpha^t\}_{t=0}^T$ is a monotonic, strictly decreasing sequence, in which $\alpha^0 = 1$ and α^T vanishes as T increases. This ensures z^T roughly follows an isotropic Gaussian distribution. In contrast, the reverse process involves sequential Gaussian sampling given by

$$z^{t-1} = \mu(z^t, \epsilon_\theta, t) + \sigma(t)w, \quad w \sim N(0, I). \quad (2)$$

Here, μ and σ are pre-defined functions to compute parameters of Gaussian distribution. μ is parameterized by ϵ_θ that is a time-conditioned UNet (Ronneberger, Fischer, and Brox 2015) equipped with attention mechanism and trained to achieve the objective

$$\min_{\theta} \mathbb{E}_{z^0, \epsilon, t} \|\epsilon - \epsilon_\theta(z^t, t, \gamma)\|_2^2. \quad (3)$$

γ is the text embedding of \mathcal{P} mapped by a frozen model ϕ , e.g., the CLIP text encoder (Radford et al. 2021). The reverse process ensures the final sample z^0 follows the distribution of latents, and the generated image aligns with the prompt.

Framework of StyO

Given a source image x_{src} in natural face domain, and a target image x_{tgt} in artistic portrait domain, our objective is to synthesize an image with the source content as well as the target style. Building upon the state-of-the-art LDM (Rombach et al. 2022), we propose a novel approach named StyO, which consists identifier disentanglement learner (IDL) and fine-grained content controller (FCC) to tackle the task. In general, StyO learns disentangled identifiers representing content and style via IDL, and enables the retention of fine-grained content details through FCC when inferring the stylized image with the prompt comprising recombined identifiers. A detailed sketch of our proposed StyO is illustrated in Figure 2.

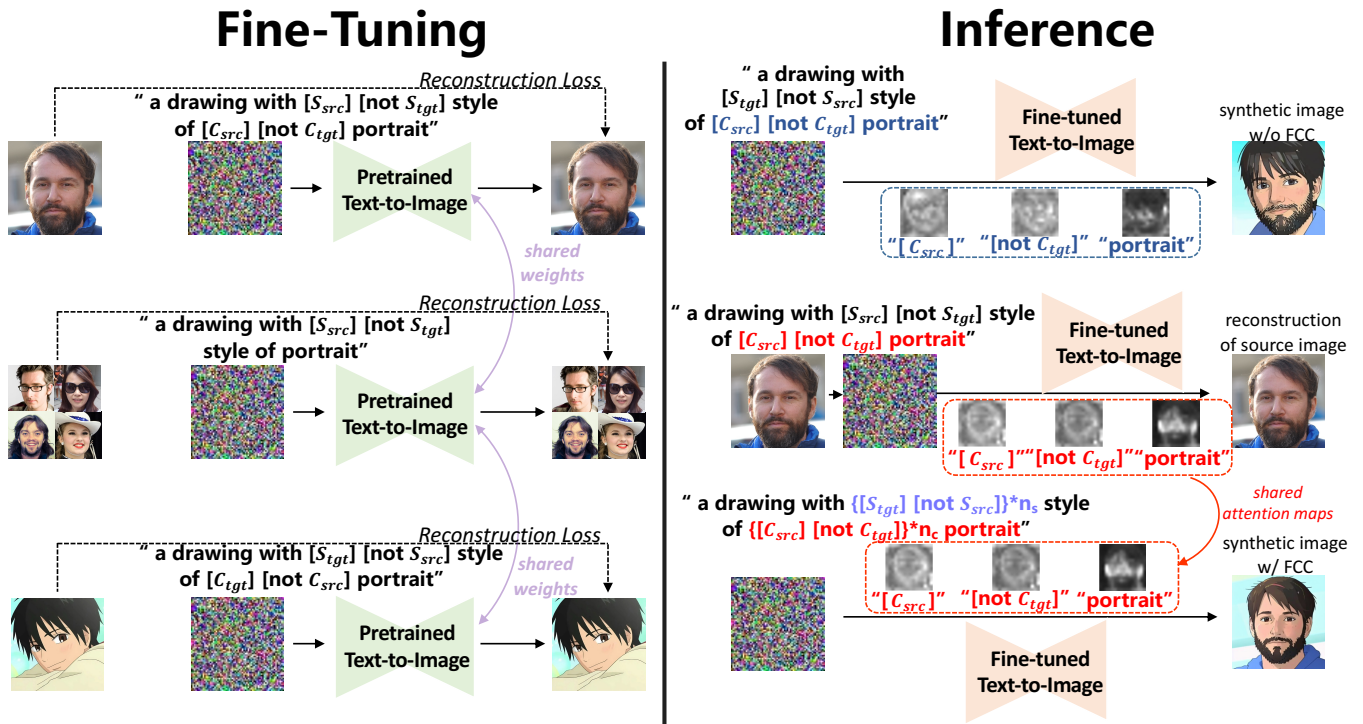


Figure 2: Pipeline of the proposed StyO. Left panel: The training phase. StyO constructs text-image pairs for source, target images and auxiliary image set with a contrastive disentangled prompt template. Then, StyO fine-tunes a pre-trained latent diffusion model in the *one-shot* manner. In this way, StyO injects attributes of images into different identifiers. Right panel: The inference phase. Given style and content identifiers, StyO recombines them to form text prompt to generating stylized face. To maintain fine-grained details, StyO extracts attention maps for source content and uses them to replace that for stylized one. This process yields faces with suitable target style while maintains good source content.



Figure 3: Qualitative comparison of our StyO with GOGA (Zhang et al. 2022b), Custom Diffusion (Kumari et al. 2023), and StyleID (Chung, Hyun, and Heo 2024). Better see in color and 2x zoom.



Figure 4: Ablation of Contrastive Disentangled Prompt Template. The arrows indicate the changes made to the current prompt template by sequentially plusing positive identifiers, negative identifiers, and auxiliary image set. Better see in color and 2x zoom.




Figure 5: Ablation of Fine-grained Content Controller. Better see in color and 2x zoom.

Identifier Disentanglement Learner In order to extract the content and style information from given images, we design a novel contrastive disentangled prompt template with positive-negative identifiers and information descriptors to construct the text label for each image. These image-text pairs will be used to fine-tune LDM with a triple reconstruction loss, such that the content and style are encoded into identifiers.

Contrastive Disentangled Prompt Template. As shown in Figure 2, we first label each image with a novel contrastive disentangled prompt template: “a drawing with $[S^+][S^-]$ style of $[C^+][C^-]$ portrait”, where “style” and “portrait” are the information descriptor, “ $[S^+]$ ”/“ $[S^-]$ ” is a positive/negative unique identifier linked to the style, and “ $[C^+]$ ”/“ $[C^-]$ ” is a positive/negative unique identifier linked to the content of subject. Specifically, the positive identifiers are expected to hold semantic meaning corresponding to the image. With regard to the source/target image, we use “ $[S_{src}]$ ”/“ $[S_{tgt}]$ ” and “ $[C_{src}]$ ”/“ $[C_{tgt}]$ ” for “ $[S^+]$ ” and “ $[C^+]$ ”, respectively. In contrast, the negative identifier is the nomenclature proposed to represent elements that are not presented in the image. We cooperate the negator, e.g., “not”, “without” or “except”, with above positive identifiers to construct negative identifiers. For instance, we set “ $[S^-]$ ” as “ $[not S_{tgt}]$ ”, and “ $[C^-]$ ” as “ $[not C_{tgt}]$ ” for the negative identifiers in the prompt of source image. Such

design of contrastive prompt highlights the significant visual disparity between the source and target images, and will facilitate the disentangled learning by the effective exploitation of the limited given data.

In addition to the task-relevant data x_{src} and x_{tgt} , we introduce a *no-cost* auxiliary image set x_{aux} to further enhance the disentanglement, in which about 200 natural faces are randomly sampled from FFHQ (Karras, Laine, and Aila 2019) and labeled with a same prompt “a drawing with $[S_{src}][not S_{tgt}]$ style of portrait”. The contrast between prompts of x_{aux} and x_{src} helps to encode source content into identifiers “ $[C^+]$ ”/“ $[C^-]$ ” instead of the descriptor “portrait”, and strengthen the style distinction between x_{src} and x_{tgt} . To sum up, we list prompts corresponding to different images or set as:

- Prompt of source image (\mathcal{P}_{src}): “a drawing with $[S_{src}][not S_{tgt}]$ style of $[C_{src}][not C_{tgt}]$ portrait”.
- Prompt of target image (\mathcal{P}_{tgt}): “a drawing with $[S_{tgt}][not S_{src}]$ style of $[C_{tgt}][not C_{src}]$ portrait”.
- Prompt of auxiliary image set (\mathcal{P}_{aux}): “a drawing with $[S_{src}][not S_{tgt}]$ style of portrait”.

Note that, different from (Ruiz et al. 2022) using rare words in T5-XXL tokenizer, we simply use firearms (e.g., “sks”, “aug” and “ak47”) as identifiers and also achieve phenomenal results.

Triple Reconstruction Loss. Fine-tuning of diffusion models is a powerful technique relevant to many use cases, e.g., image editing (Valevski et al. 2022; Saharia et al. 2022; Ruiz et al. 2022) and image-to-image translation (Wang et al. 2022). To make the identifiers work as intended, we opt for this simple approach with the text-image pair $\{(\mathcal{P}_i, x_i)\}_{i \in \{src, tgt, aux\}}$. We tune the text-to-image model with a triple reconstruction loss, thereby injecting different information to individual identifiers:

$$\min_{\theta} \mathbb{E}_{z_{src}, z_{tgt}, z_{aux}, \epsilon, t} \|\epsilon - \epsilon_{\theta}(z_{src}^t, t, \gamma_{src})\| + \|\epsilon - \epsilon_{\theta}(z_{tgt}^t, t, \gamma_{tgt})\| + \|\epsilon - \epsilon_{\theta}(z_{aux}^t, t, \gamma_{aux})\|, \quad (4)$$

where $z_{src} = \mathcal{E}(x_{src})$, $\gamma_{src} = \phi(\mathcal{P}_{src})$, z_{src}^t is the noisy latent according to Eq. (1), and same goes for other symbols. Note, we simultaneously fine-tune the text encoder ϕ to achieve better results. By implementing IDL, we can obtain a tuned model and disentangled identifiers biased to given images.

Fine-grained Content Controller The proposed IDL allows for the effective fusion of source identity and artistic style via the recombination of identifiers:

- Prompt of stylized portrait (\mathcal{P}_{sty}): “a drawing with $[S_{tgt}][not S_{src}]$ style of $[C_{src}][not C_{tgt}]$ portrait”

Nonetheless, as depicted in Figure 2, the synthesis that is highly correlated with the source face, exhibits noticeable diversity in terms of poses and some facial features from source image. Although this diversity can offer additional customization options, it would be desirable for the model to consistently produce artistic portraits that accurately capture the source face. To alleviate this problem, we introduce a Fine-grained Content Controller (FCC) module that further transfers delicate details of source image into syntheses. In

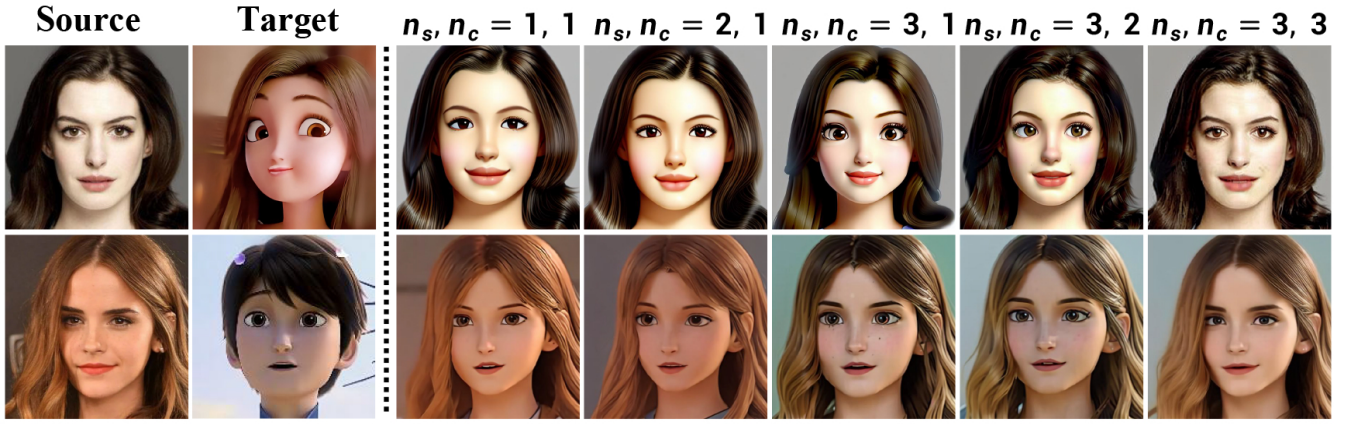


Figure 6: Ablation of hyperparameter in augmented text prompt. Better see in color and 2x zoom.

technique, FCC devises the cross attention control and augments text prompt to ensure content attention maps of the stylized portrait correspond closely to those of the source face, as depicted in the right panel of Figure 2.

Cross Attention Control. LDM utilizes a cross attention mechanism to modify the latent features based on the conditional text features. Here we slightly abuse some notations: given the text embedding $\gamma \in \mathbb{R}^{s \times d_1}$, latent features $z \in \mathbb{R}^{h \times w \times d_2}$, where s denotes the number of tokens, h and w are the height and width of latent feature, and d_1 and d_2 are the dimensions of features, respectively. The cross attention first computes query $Q = zW^q$, key $K = \gamma W^k$, and value $V = \gamma W^v$. Then latent features are updated by

$$M(z, \gamma) = \text{softmax} \left(\frac{QK^T}{\sqrt{d}} \right), \quad (5)$$

$$\text{Attn}(z, \gamma) = M(z, \gamma)V, \quad (6)$$

where d is the dimension of key and query feature.

Hertz et al. (2022) find that the attention map M in Eq. (5) has a strong influence on the spatial layout of syntheses. In the face domain, we are inspired to control the pose and facial details of stylized portraits by constraining the attention mask of syntheses to that of source image. Specifically, considering the feed-forward process of LDM prompted by \mathcal{P}_{src} and \mathcal{P}_{sty} , their latent features and text features are (z_{sty}, z_{src}) and $(\gamma_{sty}, \gamma_{src})$, respectively. We modify the cross attention to a condition form:

$$\text{Attn}(z_{sty}, \gamma_{sty}; z_{src}, \gamma_{src}) = M_{ctr} V_{sty}, \quad (7)$$

where $M_{ctr}^i = M_{src}^i$ if $i \in \text{Content Index}$ else M_{sty}^i .

M_{sty} and M_{src} are the attention maps computed by the corresponding text feature and latent feature, and the superscript i denotes the i -th column of matrix, and the Content Index includes the indexes of “[C_{src}]”, “[$not C_{tgt}$]” and “portrait” in \mathcal{P}_{src} . Note that the tokens of \mathcal{P}_{src} and \mathcal{P}_{sty} have a same length, thus FCC indeed swaps the attention maps of those content identifiers to pull the content of the stylized portrait towards source image.

Augmented Text Prompt. The text prompt is a critical factor in controlling the content of syntheses for LDM, and

improving its quality usually enhances the semantic details of syntheses directly. Leveraging the disentangled identifiers obtained from IDL, we introduce an augmented text prompt that offers greater customization and control over stylized portraits. Given a identifier “[C]” and $n \in \mathbb{Z}^+$, we define the augmented identifier “[C]” $\ast n$ as repeating the identifier n times. Augmented text prompts \mathcal{P}_{src}^{aug} and \mathcal{P}_{sty}^{aug} comprises augmented identifiers:

- Augmented prompt of source face (\mathcal{P}_{src}^{aug}): “a drawing with ($[S_{src}][not S_{tgt}])\ast n_s$ style of ($[C_{src}][not C_{tgt}])\ast n_c$ portrait”.

- Augmented prompt of stylized portrait (\mathcal{P}_{sty}^{aug}): “a drawing with ($[S_{tgt}][not S_{src}])\ast n_s$ style of ($[C_{src}][not C_{tgt}])\ast n_c$ portrait”.

Instead of \mathcal{P}_{src} and \mathcal{P}_{sty} , we use prompt \mathcal{P}_{src}^{aug} to reconstruct the source face and \mathcal{P}_{sty}^{aug} to generate the stylized portrait with the aid of cross attention control. It is important to note that two hyper-parameters, n_s and n_c , are introduced to control the degree of stylization and identity preservation. Compared with the single identifier as used during training, a balance of n_s and n_c will result in better effects. Actually, it is not complicated to determine the optimal values for the two hyper-parameters. Through experiments, we find setting n_s and n_c to small values (≤ 3) is adequate to produce satisfactory results.

Inference with FCC. On this basis, the process of generation with FCC can be summarized as follows: We carry out destructive process Eq. (1) to obtain the noisy latent z_{src}^T of source image. Note that compared with random Gaussian noise, z_{src}^T provides a better prior to restore x_{src} . Then we perform the reverse diffusion starting from z_{src}^T and prompted by \mathcal{P}_{src}^{aug} . Meanwhile, their attention maps M_{sty} in all cross attention layers are recorded as Eq. (5). Once again, we perform the reverse diffusion starting from z_{src}^T and prompted by \mathcal{P}_{sty}^{aug} to synthesize the stylized portrait, while it implements the conditional cross attention as Eq. (7) to fetch fine-grained details from source image.

Model	Identity \uparrow	Geometry \uparrow	Texture \uparrow
GOGA (Zhang et al. 2022b)	0.11	0.14	0.23
Custom Diffusion (Kumari et al. 2023)	0.09	0.27	0.11
StyleID (Chung, Hyun, and Heo 2024)	0.38	0.12	0.28
Ours	0.42	0.47	0.38

Table 1: Quantitative comparison of our StyO and state-of-the-arts based on the user study. Please see details in text.

Experiments

Comparison With SOTA Methods

Baselines. We compare StyO with two recent representative state-of-the-arts of GOGA (Zhang et al. 2022b), Custom Diffusion (Kumari et al. 2023) and StyleID (Chung, Hyun, and Heo 2024) to show the advantage of StyO. All these methods are implemented by their official codes. We refrain from comparing with InstantID (Wang et al. 2024) due to its lack of capability in extracting style from target images. The artistic portraits are taken from AAHQ dataset (Liu et al. 2021b), while the natural faces employ photos of celebrities.

Qualitative Comparison. Figure 3 shows the qualitative results. The top row images are the source natural faces and the leftmost images are the target artistic portraits. From the results, we can draw the following conclusions: 1) Our method not only achieves facial identity consistency between the source and output images, but also highly preserves local details such as hair and eye color, which is a significant improvement over other methods. In contrast, other methods often fail to generate images that are faithful to the identity information of source images and transfers unexpected colors in some cases. 2) Our method captures the desired geometric variations in the target image, such as big eyes and round face. Instead of using GAN based generative model, we use potential text-image diffusion model to successfully capture the geometric variation in reference image and inject it to unique identifier. Obviously, other methods can not associate big eyes and round faces as consistent geometric variations and smoothly introduce both of these into the source image. (3) Our method produces syntheses that flexibly acquire vivid local stroke characteristics and overall appearance, resulting in high-quality images with sharp and realistic high-frequency details. In contrast, the results produced by other methods tend to be fuzzy.

Quantitative Comparison. As portrait stylization is often regarded as a subjective task, we resort to user studies to evaluate which method generates results that are most favored by humans. We conducted three user studies on the results in terms of the identity, geometric and texture. In the first study, participants are asked to select the stylized images best preserve the identity of source image. In the second study, participants are asked to point out the best stylized images with reasonable exaggerated geometric deformation. In the third study, we encourage participants to touch the image with the finest texture and less distorted artifacts. We receive 150 answers on 50 source-target pairs in total for each study. As shown in Table 1, over 0.38% of our results are selected as the best in both three metrics, which proves a significant ad-

vantage in stylization.

Ablation Study

Effect of Contrastive Disentangled Prompt Template. Our proposed contrastive disentangled prompt template acts as a powerful tool to disentangle and extract style and content information of given images. In contrast, a prompt template that only contains positive identifiers, *e.g.*, the prompt used in Dreambooth (Ruiz et al. 2022), faces two challenging issues: 1) Over-fitting to source or target image. 2) Failed to decouple style and content information. To better characterize the claims, we fine-tune LDM with three different configurations, positive identifiers, positive-negative identifiers, and our positive-negative identifiers plus the auxiliary image set. As shown in Figure 4, the results generated by positive identifier overfit to source image and often generates “haloing” artifacts around faces. Adding negative identifiers tends to alleviate this problem and generate acceptable images. Nevertheless, it may fail to capture the local stroke characteristics and generate sharp details, and there still exists entanglement between style and content. In comparison, the full prompt template achieves the best performance in terms of style effect and content preservation.

Effect of Fine-grained Content Controller. To verify the effectiveness of the proposed FCC module, we perform the evaluation by inferring artistic portrait without it. As shown in Figure 5, StyO without FCC will generate diverse artistic portraits that maintains source identity, due to the content identifiers also store the content information of source image. This characteristic allows users to have more choices, such as different poses, different hairstyles and eye colors. With the introduction of FCC module, it is obvious that the diversity of syntheses is significantly reduced and the identity fidelity to source image is also improved.

Hyper-parameters in Augmented Text Prompt. The ablation in Figure 6 analyzes the visual results by varying the number of style identifiers n_s and the number of content identifiers n_c in FCC. Here, we find that with the increase of n_s , the style of syntheses gradually change from photorealistic to animated. Similarly, a larger n_c will force the output image to “overfit” to the source image. Actually, there is a trade-off between content faithfulness and style similarity related to hyper-parameter n_s and n_c . Note that in practice, we can experiment with various hyper-parameters to synthesize multiple artistic portraits, allowing the user to choose among different levels of artistic effects.

Conclusions

In this paper, we present a novel StyO model to effectively solve the one-shot face stylization. StyO exploits a disentanglement and recombination strategy. For this purpose, StyO introduces two new modules, *i.e.*, Identifier Disentanglement Learner (IDL) and Fine-grained Content Controller (FCC). IDL aims to disentangle style and content attributes of images into different identifiers and FCC recombines style and content identifiers to form prompts describing the stylization faces. Extensive results show that our StyO model can generate high-quality stylization results in one-shot manner, clearly outperforming existing methods.

Acknowledgements

This paper is supported is supported by National Key R&D Program of China (2021YFA1000403), National Natural Science Foundation of China (Nos. U23B2012, 11991022), and Fundamental Research Funds for the Central Universities (E3E41904X2).

References

- Akleman, E. 1997. Making caricatures with morphing. *CGIT*.
- Cao, K.; Liao, J.; and Yuan, L. 2018. Carigans: Unpaired photo-to-caricature translation. *arXiv preprint arXiv:1811.00222*.
- Chong, M. J.; and Forsyth, D. 2022. Jojogan: One shot face stylization. In *ECCV*, 128–152. Springer.
- Chung, J.; Hyun, S.; and Heo, J.-P. 2024. Style injection in diffusion: A training-free approach for adapting large-scale diffusion models for style transfer. In *CVPR*, 8795–8805.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. In *NeurIPS*, 8780–8794.
- Gal, R.; Alaluf, Y.; Atzmon, Y.; Patashnik, O.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2022a. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*.
- Gal, R.; Patashnik, O.; Maron, H.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2022b. StyleGAN-NADA: CLIP-guided domain adaptation of image generators. *ACM Transactions on Graphics*, 41(4): 1–13.
- Gatys, L. A.; Ecker, A. S.; and Bethge, M. 2016. Image style transfer using convolutional neural networks. In *CVPR*, 2414–2423.
- Gong, J.; Hold-Geoffroy, Y.; and Lu, J. 2020. Autotoon: Automatic geometric warping for face cartoon generation. In *WACV*, 360–369.
- Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A. C.; and Bengio, Y. 2014. Generative Adversarial Nets. In *NeurIPS*.
- Han, F.; Ye, S.; He, M.; Chai, M.; and Liao, J. 2021. Exemplar-Based 3D Portrait Stylization. *IEEE Transactions on Visualization and Computer Graphics*, 29: 1371–1383.
- Hertz, A.; Mokady, R.; Tenenbaum, J.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2022. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. In *NeurIPS*, 6840–6851.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *CVPR*, 1125–1134.
- Kalischek, N.; Wegner, J. D.; and Schindler, K. 2021. In the light of feature distributions: moment matching for Neural Style Transfer. In *CVPR*.
- Karras, T.; Aittala, M.; Aila, T.; and Laine, S. 2022. Elucidating the design space of diffusion-based generative models. *arXiv preprint arXiv:2206.00364*.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *CVPR*, 4401–4410.
- Kawar, B.; Zada, S.; Lang, O.; Tov, O.; Chang, H.; Dekel, T.; Mosseri, I.; and Irani, M. 2022. Imagic: Text-based real image editing with diffusion models. *arXiv preprint arXiv:2210.09276*.
- Kim, J.; Kim, M.; Kang, H.; and Lee, K. 2019a. U-gat-it: Unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation. *arXiv preprint arXiv:1907.10830*.
- Kim, J.; Kim, M.; Kang, H.; and Lee, K. 2019b. U-gat-it: Unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation. *arXiv preprint arXiv:1907.10830*.
- Kolkin, N. I.; Salavon, J.; and Shakhnarovich, G. 2019. Style Transfer by Relaxed Optimal Transport and Self-Similarity. In *CVPR*.
- Kumari, N.; Zhang, B.; Zhang, R.; Shechtman, E.; and Zhu, J.-Y. 2023. Multi-concept customization of text-to-image diffusion. In *CVPR*, 1931–1941.
- Li, Y.; Wang, N.; Liu, J.; and Hou, X. 2017. Demystifying neural style transfer. *arXiv preprint arXiv:1701.01036*.
- Liu, L.; Ren, Y.; Lin, Z.; and Zhao, Z. 2022. Pseudo numerical methods for diffusion models on manifolds. *arXiv preprint arXiv:2202.09778*.
- Liu, M.; Li, Q.; Qin, Z.; Zhang, G.; Wan, P.; and Zheng, W. 2021a. Blendgan: Implicitly gan blending for arbitrary stylized face generation. In *NeurIPS*, 29710–29722.
- Liu, M.; Li, Q.; Qin, Z.; Zhang, G.; Wan, P.; and Zheng, W. 2021b. BlendGAN: Implicitly GAN Blending for Arbitrary Stylized Face Generation. *ArXiv*, abs/2110.11728.
- Liu, M.-Y.; Breuel, T.; and Kautz, J. 2017. Unsupervised image-to-image translation networks. In *NeurIPS*.
- Men, Y.; Yao, Y.; Cui, M.; Lian, Z.; and Xie, X. 2022. DCT-net: domain-calibrated translation for portrait stylization. *ACM Transactions on Graphics*, 41(4): 1–9.
- Mokady, R.; Hertz, A.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2022. Null-text Inversion for Editing Real Images using Guided Diffusion Models. *arXiv preprint arXiv:2211.09794*.
- Nichol, A.; Dhariwal, P.; Ramesh, A.; Shyam, P.; Mishkin, P.; McGrew, B.; Sutskever, I.; and Chen, M. 2021. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *ICML*, 8162–8171.
- Ojha, U.; Li, Y.; Lu, J.; Efros, A. A.; Lee, Y. J.; Shechtman, E.; and Zhang, R. 2021. Few-shot image generation via cross-domain correspondence. In *CVPR*, 10743–10752.
- Pan, L.; Zhao, L.; Xu, D.; and Lu, D. 2019. Optimal Transport of Deep Feature for Image Style Transfer. In *MM*.

- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *CVPR*, 10684–10695.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 234–241. Springer.
- Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2022. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. *arXiv preprint arXiv:2208.12242*.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E.; Ghasemipour, S. K. S.; Ayan, B. K.; Mahdavi, S. S.; Lopes, R. G.; et al. 2022. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. *arXiv preprint arXiv:2205.11487*.
- Shi, Y.; Deb, D.; and Jain, A. K. 2019. Warpgan: Automatic caricature generation. In *CVPR*, 10762–10771.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2256–2265.
- Song, G.; Luo, L.; Liu, J.; Ma, W.-C.; Lai, C.-P.; Zheng, C.; and Cham, T.-J. 2021. AgileGAN: stylizing portraits by inversion-consistent transfer learning. *ACM Trans. Graph.*, 40: 117:1–117:13.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Song, Y.; and Ermon, S. 2020. Improved techniques for training score-based generative models. In *NeurIPS*, 12438–12448.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.
- Valevski, D.; Kalman, M.; Matias, Y.; and Leviathan, Y. 2022. Unitune: Text-driven image editing by fine tuning an image generation model on a single image. *arXiv preprint arXiv:2210.09477*.
- Wang, Q.; Bai, X.; Wang, H.; Qin, Z.; and Chen, A. 2024. Instantid: Zero-shot identity-preserving generation in seconds. *arXiv preprint arXiv:2401.07519*.
- Wang, T.; Zhang, T.; Zhang, B.; Ouyang, H.; Chen, D.; Chen, Q.; and Wen, F. 2022. Pretraining is all you need for image-to-image translation. *arXiv preprint arXiv:2205.12952*.
- Wang, T.-C.; Liu, M.-Y.; Zhu, J.-Y.; Tao, A.; Kautz, J.; and Catanzaro, B. 2018. High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs. In *CVPR*.
- Xia, W.; Zhang, Y.; Yang, Y.; Xue, J.-H.; Zhou, B.; and Yang, M.-H. 2021. GAN Inversion: A Survey. *arXiv preprint arXiv:2101.05278*.
- Zhang, Y.; Yao, M.; Wei, Y.; Ji, Z.; Bai, J.; and Zuo, W. 2022a. Towards Diverse and Faithful One-shot Adaption of Generative Adversarial Networks. *ArXiv*, abs/2207.08736.
- Zhang, Z.; Li, B.; Nie, X.; Han, C.; Guo, T.; and Liu, L. 2024. Towards consistent video editing with text-to-image diffusion models. In *NeurIPS*.
- Zhang, Z.; Liu, Y.; Han, C.; Guo, T.; Yao, T.; and Mei, T. 2022b. Generalized One-shot Domain Adaption of Generative Adversarial Networks. *arXiv preprint arXiv:2209.03665*.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2223–2232.
- Zhu, P.; Abdal, R.; Femiani, J.; and Wonka, P. 2021. Mind the gap: Domain gap control for single shot domain adaptation for generative adversarial networks. *arXiv preprint arXiv:2110.08398*.