

A Style-aware Discriminator for Controllable Image Translation

Kunhee Kim Sanghun Park Eunyeong Jeon Taehun Kim Daijin Kim
Pohang University of Science and Technology (POSTECH)

{kunkim, sanghunpark, eyjeon, taehoon1018, dkim}@postech.ac.kr

Abstract

Current image-to-image translations do not control the output domain beyond the classes used during training, nor do they interpolate between different domains well, leading to implausible results. This limitation largely arises because labels do not consider the semantic distance. To mitigate such problems, we propose a style-aware discriminator that acts as a critic as well as a style encoder to provide conditions. The style-aware discriminator learns a controllable style space using prototype-based self-supervised learning and simultaneously guides the generator. Experiments on multiple datasets verify that the proposed model outperforms current state-of-the-art image-to-image translation methods. In contrast with current methods, the proposed approach supports various applications, including style interpolation, content transplantation, and local image translation. The code is available at github.com/kunheek/style-aware-discriminator.

1. Introduction

Image-to-image (I2I) translation aims to manipulate the style of an existing image, where style refers to generic attributes that can be applied to any image in a dataset (e.g., texture or domain). Content generally refers to the remaining information, such as the pose and structure. This task has shown significant progress with generative adversarial network (GAN) [15] developments. Recent studies have expanded the functionality to multi-modal and multi-domains using domain-specific discriminators and latent injection [18], enabling the direct manipulation of existing images using domain labels or reference images [10, 11, 31, 32, 40].

However, despite promising functionality advances, there remains considerable room for development in terms of controllability. For example, users can only control the classes used for training. Although a reference image can be used to control output but this can often lead to erroneous results, particularly for misrecognition within the same class; and another common problem is inability to fine-tune the output. Since the label space does not con-

sider the semantic distance between classes, the learned style space cannot reflect these semantic distances, which leads to unrealistic images when controlling the results by manipulating the style code [32].

This study investigates I2I translation controllability, i.e., to be able to edit the result as desired using the style code, without being limited to the previously defined label space. The proposed model learns the style space using prototype-based self-supervised learning [6] with carefully chosen augmentations. Although the current domain-specific discriminators are not designed for an external continuous space, this is possible if the discriminator knows the style internally. Therefore, we propose a *Style-aware Discriminator*, combining a style encoder and a discriminator into a single module. Thus, the proposed model is somewhat lighter by reducing one module and achieves better performance because of to the better representation space of the discriminator. We used the style code sampled from prototypes during training to improve the controllability; and feature-level and pixel-level reconstructions to improve the consistency. Thus, the proposed model goes beyond image translation to support various applications, including style interpolation and content transplantation. Finally, we propose feedforward local image translation by exploiting spatial properties of the GAN feature space.

We evaluated the model on several challenging datasets: Animal Faces HQ (AFHQ) [11], CelebA-HQ [22], LSUN churches [41], Oxford-102 [34], and FlickrFaces-HQ (FFHQ) [24]. Extensive experiments confirm that the proposed method outperforms current state-of-the-art models in terms of both performance and efficiency without semantic annotations. The proposed model can also project an image into the latent space faster than baselines while achieving comparable reconstruction results.

The contributions from this study are summarized as follows: (i) We propose an integrated module for style encoding and adversarial losses for I2I translation, as well as a data augmentation strategy for the style space. The proposed method reduces the parameter count significantly and does not require semantic annotations. (ii) We achieve state-of-the-art results in *truly* unsupervised I2I translation

in terms of the Fréchet Inception Distance (FID) [17]. The proposed method shows similar or better performance compared with supervised methods. (iii) We extend image translation functionality to various applications, including style interpolation, content transplantation, and local image translation.

2. Related work

Multi-domain I2I translation StarGAN [10] enabled many-to-many translation using a given attribute label, but this and similar approaches have the disadvantage of being deterministic for a given input and domain. Subsequent studies suggested using reference images rather than labels [31, 40], enabling translation based on an image from unseen classes in the same domain. StarGAN v2 [11] introduced a noise-to-latent mapping network to synthesize diverse results for the same domain, but since all of these methods depend on labels defined for classification, representations for image manipulation cannot be learned. Therefore, we developed a new multi-domain I2I approach from two perspectives. The proposed method learns a style-specific representation suitable for image manipulation without relying on labels; and then provides more user-controllability while supporting various applications.

To overcome the problem of label dependency, Bahng *et al.* [4] clustered the feature space for pre-trained networks to create pseudo-labels and corresponding latent code. Similarly, TUNIT trained a guiding network using contrastive learning and clustering directly in target data [3]. These methods obtain pseudo-labels that can be substituted for class labels; however, the proposed approach models a continuous style space rather than discrete pseudo-labels. Thus the proposed model is significantly more efficient than previous approaches. CLUIT [27] recently proposed using contrastive learning through a discriminator, but used contrastive learning to replace the multi-task discriminator. Therefore, the style encoder exists independently, in contrast with the proposed model. Furthermore, CLUIT requires additional process (*e.g.*, clustering), to obtain a style code without a reference image.

Learning-based image editing Recently, Karras *et al.* discovered that GANs naturally learn to disentangle pre-defined latent space [24, 25]. Several subsequent studies proposed image editing methods using StyleGANs [1, 2, 43]. However, these methods suffered from the long time it takes to find a corresponding latent. Recently, StyleMapGAN [26] and Swapping Autoencoder (SwapAE) [37] proposed directly encoding an image into the latent, enabling real-time and various image editing applications. Our study is different in that content and style can be manipulated separately because of disentangled latent space. SwapAE has a separate latent space called texture and structure, similar to the proposed method, but is challenging to operate without

a reference image. In addition, its texture-focused representation does not work well for tasks that require dramatic changes, such as the interspecies variation of animal faces (Fig. 3). On the other hand, since the proposed method learns the style space and the prototype, manipulating an image without a reference image is possible. Furthermore, because our method is designed for I2I translation, more challenging manipulations are possible.

Discriminator and self-supervised learning GANs [15] have always suggested that discriminators could be feature extractors, and many previous studies have demonstrated that GANs benefit from representation learning through a discriminator [9, 20, 21, 29, 30]. We also utilize self-supervised learning via the discriminator, but differ from previous approaches in that the primary purpose of the self-supervised learning is to function as an encoder, not just to improve the quality. Hence, our discriminator continues to work as an encoder after training; as opposed to most current GANs, which abandon discriminators after training.

3. Methods

Our aim was to build a flexible and manipulative style space. In particular, we considered the following objectives. (i) Visual similarity should be considered. For example, visually similar pairs, such as a wolf and a dog, should be placed in similar places in the style space. (ii) There should be a representative value, such as a discrete label, providing a good starting point when the user wants to fine-tune the result.

3.1. Framework

The framework overview is shown in Fig. 1, which shows the sampling strategies for the style code and the training procedures of the entire model.

Style-aware discriminator Given an image $\mathbf{x} \in \mathcal{X}$, discriminator D returns a vector as output. The discrimination head h_D determines whether \mathbf{x} is a real or fake image, and the style head outputs the latent code $\mathbf{z}_s = h_s(D(\mathbf{x}))$. We formulate the traditional discriminator $f_D(\mathbf{x})$ and the style encoder $f_s(\mathbf{x})$ as $h_D(D(\mathbf{x}))$ and $h_s(D(\mathbf{x}))$, respectively.

Prototypes We represent the style space using a set of L2-normalized vectors $\mathbf{C} \in \mathbb{R}^{K \times D}$ rather than predefined labels or pseudo-labels, where K and D denote the number of prototypes and style code dimension, respectively. We denote \mathbf{c}_k as an element of \mathbf{C} .

Generator The generator comprises an encoder and a decoder, similar to typical current I2I translation generators [11]. The encoder $G_{enc}(\mathbf{x})$ extracts style-agnostic content code $\mathbf{z}_c \in \mathbb{R}^{D \times W \times H}$ from input \mathbf{x} , and decoder $G_{dec}(\mathbf{z}_c, \mathbf{z}_s)$ synthesizes a new image reflecting content code and style code. Similar to Karras *et al.* [23], we use a 2-layer multi-layer perceptron (MLP) to transform the normalized style code \mathbf{z}_s into valid features. The generator uses

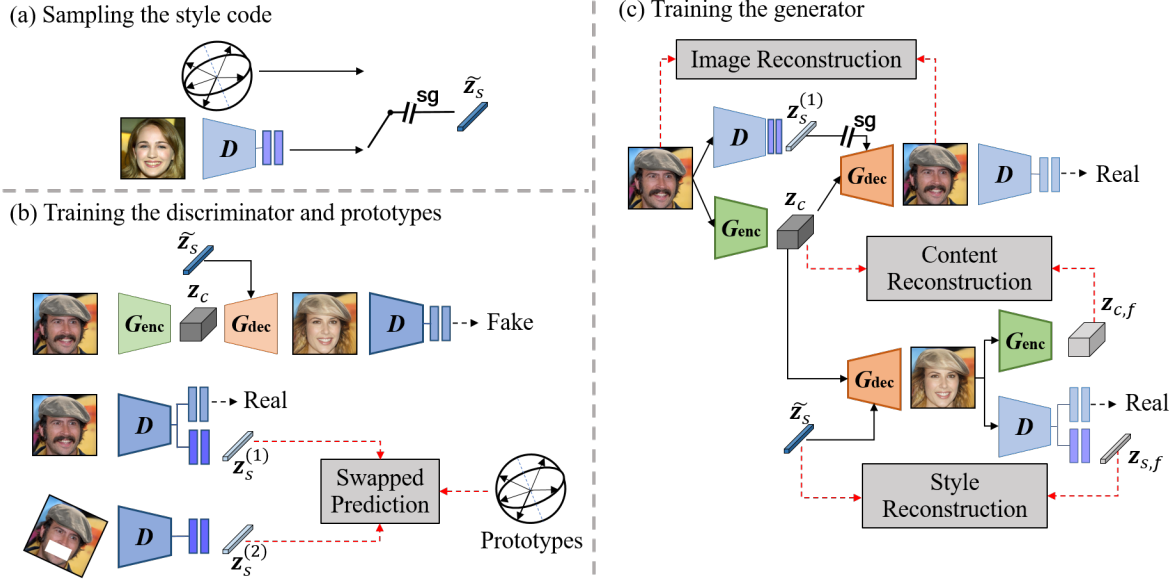


Figure 1. Framework overview. (a) The style code is sampled from the learned prototypes or dataset. (b) The discriminator not only learns to distinguish between real and fake images but also learns the style space via the swapped prediction loss. (c) The generator is enforced to utilize the style code via the style reconstruction and to preserve the input content via the content reconstruction.

weight modulation [25] or AdaIN [18] for latent injection.

3.2. Modeling the style space

Intuitively, an ideal style encoder would output the same code even though the input image was geometrically transformed. This idea is the fundamental concept underlying contrastive learning [6, 8, 16, 28, 35], which has been actively studied in recent years. We adopted self-supervised learning in our framework to learn the style space.

Data augmentation The goal of existing contrastive learning is to classify object instances other than the style in images. Chen *et al.* [8] proposed a specific augmentation pipeline (*e.g.*, random crop, color distortion) which has become the preferred approach. However, distorting the color does not serve our purpose since style is deeply related to color. Hence we use geometric transforms (*e.g.*, scale, rotation) to learn content invariant representation, and cutout [14] to learn styles such as gender and facial expressions for human faces. We also use random crop and resize following the work of [8].

SwAV We used the SwAV framework [6], online clustering based self-supervised learning, because it aligns with our goals in terms of updating prototypes and achieving better performance for small batch sizes. The basic concept is that encoded representations from both views (*i.e.*, augmented images) for the same image predict each other’s assignments \mathbf{q} . The objective for learning style space is expressed as:

$$\mathcal{L}_{swap} = l(\mathbf{q}^{(2)}, \mathbf{z}_s^{(1)}) + l(\mathbf{q}^{(1)}, \mathbf{z}_s^{(2)}), \quad (1)$$

where $l(\mathbf{q}, \mathbf{z}_s) = -\sum_k \mathbf{q}_k (\exp(\frac{\mathbf{z}_s \cdot \mathbf{c}_k}{\tau}) / \sum_{k'} \exp(\frac{\mathbf{z}_s \cdot \mathbf{c}_{k'}}{\tau}))$, τ is a temperature parameter, and \mathbf{q} is a code computed using the Sinkhorn algorithm [6, 13]. Note that swapped prediction loss can be replaced by other self-supervised learning objectives, such as InfoNCE [35], by sacrificing the advantages of the prototype.

3.3. Learning to synthesize

During training, we sample a target style code $\tilde{\mathbf{z}}_s$ from the prototype or dataset \mathcal{X} . When sampling from the prototype, we use perturbed prototypes or samples that are linearly interpolated between two prototypes (see Appendix A.3 for more details). Then, we apply a stop-gradient to prevent the style space from being affected by other objectives.

As shown in Fig. 1 (c), the generator G synthesizes a fake image $G(\mathbf{x}, \tilde{\mathbf{z}}_s)$. To enforce synthesized image be realistic, we adopted a non-saturating adversarial loss [15]:

$$\mathcal{L}_{adv} = \mathbb{E}_{\mathbf{x}} [\log(f_D(\mathbf{x}))] + \mathbb{E}_{\mathbf{x}, \tilde{\mathbf{z}}_s} [\log(1 - f_D(G(\mathbf{x}, \tilde{\mathbf{z}}_s)))] \quad (2)$$

We also employed R1 regularization [33] following previous works [3, 11, 26, 27, 37].

We adopted a *style reconstruction loss* to ensure the generator G utilize the style code:

$$\mathcal{L}_{style} = \mathbb{E}_{\mathbf{x}, \tilde{\mathbf{z}}_s} [\|\tilde{\mathbf{z}}_s - f_s(G(\mathbf{x}, \tilde{\mathbf{z}}_s))\|_2^2], \quad (3)$$

Previous multi-domain and multi-modal I2I translation methods [3, 11, 19] introduced similar objectives, the difference between the current and previous approaches is that we do not update a style encoder using this objective.

3.4. Disentanglement of style and content

An ideal image manipulation network should be able to separate an image into two mutually exclusive representations and synthesize them back into the original image without information loss [19]. Thus, the framework must satisfy the following:

$$\phi(\mathbf{x}, G(f_c(\mathbf{x}), f_s(\mathbf{x}))) = 0, \quad (4)$$

where $\phi(\cdot)$ is a distance measure in pixel space; and $f_c(\mathbf{x})$, $f_s(\mathbf{x})$ are encoding functions for content and style, respectively. To achieve this, we employ a *reconstruction loss*:

$$\mathcal{L}_{recon} = \mathbb{E}_{\mathbf{x}} [\phi(\mathbf{x}, G(\mathbf{x}, \text{sg}(f_s(\mathbf{x}))))], \quad (5)$$

where sg denotes a stop-gradient operation. This objective encourages G_{enc} to encode mutually exclusive features with the style code since $f_s(\mathbf{x})$ is not updated. Although any distance measure in pixel space can be used, we used learned perceptual image patch similarity (LPIPS) [42] since we empirically found this works better than Euclidean or Manhattan distance.

In order to learn content space through the reconstruction loss above, it is necessary to condition that the generator should not ignore input latents code. For example, the generator may ignore the content code and perform reconstruction with only style code. To prevent this, we enforce the generator to preserve input content code using a *content reconstruction loss*:

$$\mathcal{L}_{content} = \mathbb{E}_{\mathbf{z}_c, \tilde{\mathbf{z}}_c} \left[\frac{1}{WH} \sum_{i,j} \|\mathbf{z}_{c,i,j} - \tilde{\mathbf{z}}_{c,i,j}\|_2^2 \right], \quad (6)$$

where \mathbf{z}_c , $\tilde{\mathbf{z}}_c$ are $G_{enc}(\mathbf{x})$, $G_{enc}(G(\mathbf{z}_c, \tilde{\mathbf{z}}_s))$, respectively. This objective enforces patch-level similarity between inputs and outputs, similar to PatchNCE [36]. However, our proposed objective is simpler since we only compare the last layer features, and our objective does not contrast features between patches.

In practice, we found that there was no need to apply this loss every step, and hence we apply the objective every 16th step. We assume that this is because similar results can be obtained through a *reconstruction loss*.

Overall objectives Our final objective function for the discriminator is $\mathcal{L}_{StyleD} = \mathcal{L}_{adv} + \lambda_{swap} \mathcal{L}_{swap}$, and for the generator is $\mathcal{L}_G = \mathcal{L}_{adv} + \lambda_{sty} \mathcal{L}_{style} + \lambda_{rec} \mathcal{L}_{recon}$, where λ_{sty} , λ_{rec} are hyperparameters for each term, and we use for all $\lambda = 1.0$ except $\lambda_{rec} = 0.3$ for AdaIN-based models. We set K as 32 and 64 for AFHQ and CelebA-HQ, respectively. Please refer to Appendix A for more details.

3.5. Local image translation

One advantage of factored representations is having a higher degree of freedom when editing an image. The content of an image can easily be copied or moved by editing

in the content space [37]. To progress further, we propose a simple method of patch-level image translation. Kim *et al.* [26] proposed mixing spatial information in the latent space to enable local editing. Similarly, we mix spatial information in the feature space.

$$\mathbf{f}_o = \mathbf{m} \otimes \text{mod}(\mathbf{f}_i, \mathbf{z}_s^{(i)}) + (1 - \mathbf{m}) \otimes \text{mod}(\mathbf{f}_i, \mathbf{z}_s^{(j)}), \quad (7)$$

where \mathbf{f} and \mathbf{m} are feature map and mask, and mod is modulated convolution [25] or AdaIN. For patch-level image translation, we simply replace the entire modulated convolution layer [25] with above. To ensure content is maintained even when several styles are mixed, we mixed two styles with a random mask when calculating a content preserving loss.

4. Experiments

4.1. Experimental setup

We not only employed a StyleGAN2-based generator but also considered models using AdaIN to enable a fair comparison with I2I translation models that use AdaIN.

Datasets We trained the proposed and various comparator models on AFHQ, AFHQ v2 [11], CelebA-HQ [22], FFHQ [24], Oxford-102 [34], and LSUN churches [41]. Since high resolution models requires considerable training time, the proposed and comparison models were trained and evaluated at 256×256 resolution. For AFHQ and CelebA-HQ, we used the splits provided by Choi *et al.* [11].

Baselines Our primary goal is to synthesize an image with a reference image or a latent sampled from a learned space (*i.e.*, I2I translation). We compared the proposed approach with recent supervised [11, 32] and unsupervised [3, 27] methods. In contrast with most I2I translation methods, the proposed approach has further applications such as image editing. To compare real-time image editing capability, we compared our approach with Swapping Autoencoder (SwapAE) [37] and StyleMapGAN [26].

We used pre-trained networks provided by the authors whenever possible. Otherwise, we trained the models from scratch using the official implementation, except for CLUIT, where we employed our implementation because the authors have not yet published their code. We showed 1.6 and 5 M images to the AdaIN- and StyleGAN2-based models, respectively. For StyleMapGAN, we used pre-trained networks trained for 5 M images.

4.2. Main results

We quantitatively and qualitatively evaluated the proposed approach and the baselines on two datasets: AFHQ and CelebA-HQ.

Latent-guided image synthesis We report Fréchet Inception Distance (FID) [17] and Kernel Inception Distance (KID) [5] to evaluate the latent-guided image synthesis

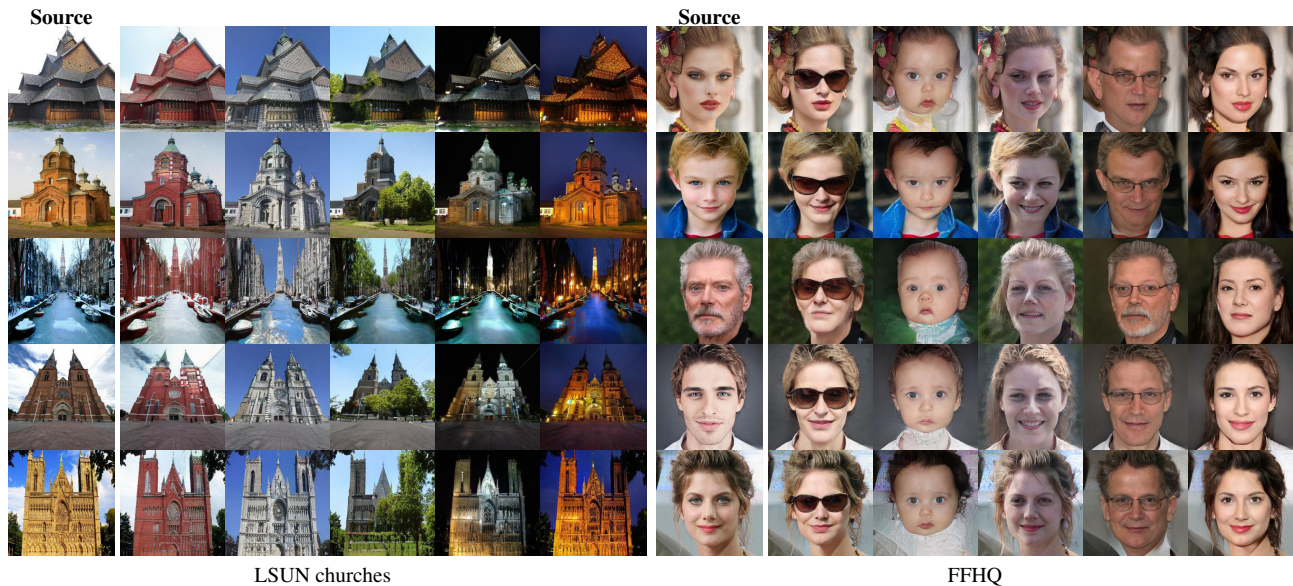


Figure 2. Prototype-guided synthesis. Our model discovers various style prototypes from the dataset in an unsupervised manner. The style prototype consists of a combination of various attributes including (left) time, weather, season, and texture; and (right) age, gender, and accessories. Each row shows the result of manipulating the leftmost image with learned prototypes.

quality, calculating FID and KID between 50,000 synthesized images and training samples. Parmer *et al.* [38] recently demonstrated that values of these metrics depend on the resizing method; therefore, we calculated FID and KID for all methods using Pillow-bicubic [12].

To synthesize images, we used a style code sampled using the strategy used in the training. To evaluate supervised methods [11, 32], we created a style code using randomly sampled domain and noise. We performed style mixing with randomly sampled latent with StyleMapGAN [26]. In Table 1, the proposed model showed better results than the existing unsupervised methods and comparable results to the supervised methods. Although the result of the proposed approach is slightly worse than StarGAN v2 in AFHQ, our approach allows users to choose one of several prototypes, whereas StarGAN v2 only allows users to choose from three classes. In Fig. 2, we show the prototype-guided synthesis results of our methods trained on unlabeled datasets. Note that we directly used prototypes obtained during the training without additional processing.

Reference-guided image synthesis Although FID/KID protocol can estimate the manipulated image quality, it provides good performance scores even if the generator ignores the given latent (*e.g.*, reconstruction). Therefore, we evaluated reference-guided image synthesis to evaluate whether the generator reflects the latent corresponding to each domain. Following [11], we synthesize images using a source-reference pair from each task (*e.g.*, cat→dog, male→female) and calculate FID and KID with a training set of a target domain. We report average values of all tasks (mFID and mKID).

As shown in the first two rows of Fig. 3, supervised approaches [11, 32] often misrecognized the style of reference images within the same classes. However, the proposed method successfully captures the styles of reference images. Furthermore, while other methods failed to preserve the details of the source image, the proposed method was the only method that preserved details such as pose and background. **User study** To investigate the human preferences, we conducted a survey using the Amazon MTurk platform. We randomly generated 100 source-reference pairs per dataset and asked the respondents to answer three questions: (Q1) Which one best reflects the style of the reference while preserving the content of the source? (Q2) Which one is the most realistic? (Q3) Which one would you use for manipulating an image? Each set was answered by 10 respondents. As shown in Table 2, the respondents obviously preferred our method in the AFHQ. In the CelebA-HQ, our model was not preferred over the supervised models (which use attribute labels and a pre-trained face alignment network); nevertheless, our model was still the most preferred among the unsupervised methods.

See Appendix B for additional results including experiments on AFHQ v2 and Oxford-102.

4.3. Controllable image translation

Real image projection To edit an image in the latent space, we first need to project the image into the latent space. What matters here is how quickly and accurately the image can be reconstructed. We measured the runtime and LPIPS [42] between the input and reconstructed images. As shown in Table 3, our model can embed an image into the

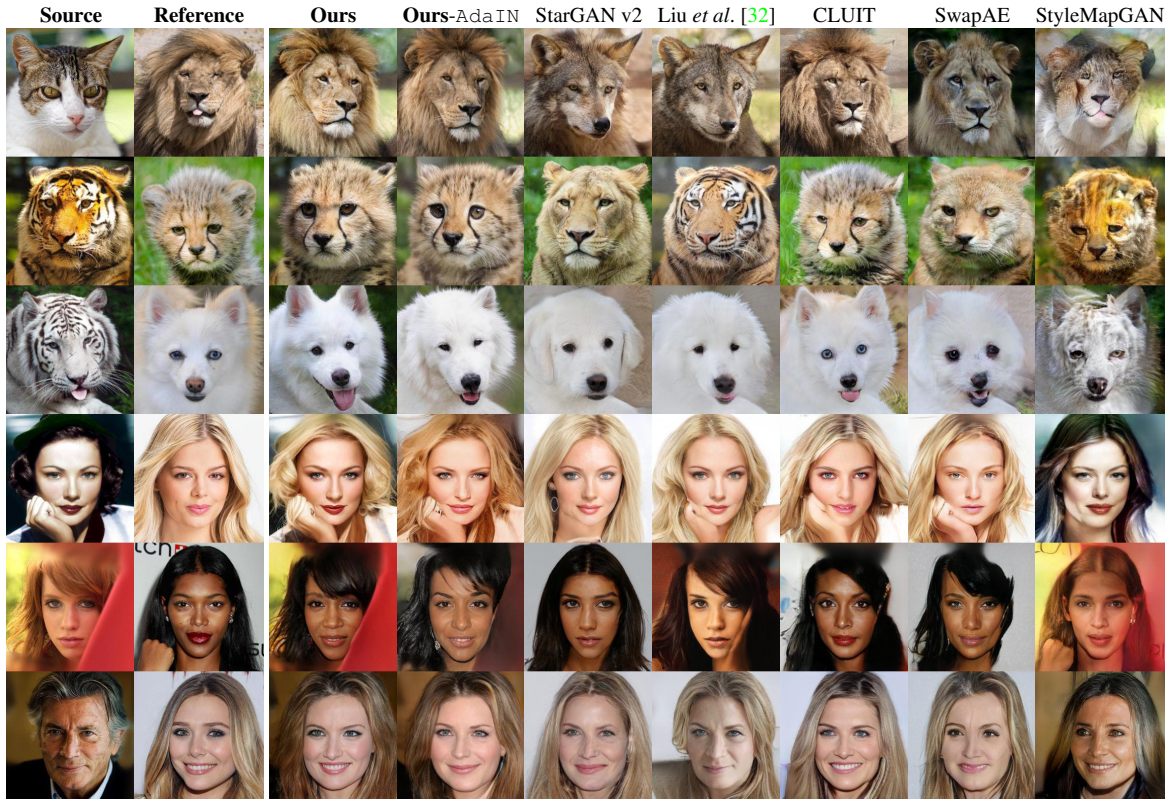


Figure 3. Qualitative comparison of reference-guided image synthesis on AFHQ (top three rows) and CelebA-HQ (bottom three rows).

Method	Param. (M)	Latent-guided synthesis				Reference-guided synthesis			
		AFHQ		CelebA-HQ		AFHQ		CelebA-HQ	
		FID↓	KID↓	FID↓	KID↓	mFID↓	mKID↓	mFID↓	mKID↓
Ours	56.51	10.0	2.1	6.8	2.8	10.6	2.1	12.6	4.9
Ours-AdaIN	57.75	12.5	2.5	10.9	4.9	14.7	5.4	17.6	8.6
*StarGAN v2 [11]	87.67	9.8	2.3	13.9	8.0	20.0	9.8	28.3	17.3
*Liu <i>et al.</i> [32]	87.67	26.0	7.0	17.8	11.0	51.7	28.6	26.7	16.8
TUNIT [3]	107.70	116.1	99.7	128.0	122.0	223.0	187.7	173.7	193.7
CLUIT [27]	80.54			N/A		22.6	10.5	28.9	18.1
SwapAE [37]	109.03			N/A		61.2	28.8	25.4	17.8
*StyleMapGAN [26]	126.23	32.8	18.7	24.3	15.2	64.3	51.3	28.8	25.1

Table 1. Quantitative comparison on image synthesis. We report FID and $KID \times 10^3$. An asterisk (*) denotes that we used the pre-trained networks provided by authors. **Bold** indicates the best result and ***bold+italicize*** indicates the best result among the *unsupervised* methods.


latent space faster and more accurately than other real-time image editing methods.

Style interpolation With the proposed method, it is possible to control only the style of the image as desired. In Fig. 4 (a), we first projected images into content and style space, then interpolated style code with randomly selected prototypes. The results show that the proposed approach is suitable for controlling the results of synthesized images.


Content transplantation Although we did not specifically target content transplantation, the proposed method

supports this application. We achieved this by copying the content code from another content code. After manipulating the content code, we synthesized the image using a style code of the source image. As shown in Fig. 4 (b), our model shows qualitatively similar results to the StyleMapGAN, which specifically targeting the local editing. Since our model separated the content and style, it is also possible to transplant only the content (*i.e.*, a big smile) without changing the style (*i.e.*, a beard) (bottom).

Local image translation Fig. 4 (c) shows the results of



(a) Reconstruction and style interpolation results on FFHQ, and AFHQ. The first two source images are from CelebA-HQ.



(b) Content transplantation comparison on CelebA-HQ



(c) Local image translation results on CelebA-HQ and AFHQ.

Figure 4. Examples of various applications. The proposed method is capable of manipulating the style and content of an image in real-time.

Method	AFHQ (%)			CelebA-HQ (%)		
	Q1	Q2	Q3	Q1	Q2	Q3
Ours	24.5	22.4	25.0	25.0	19.3	23.2
CLUIT	18.7	19.7	18.1	14.2	15.0	15.5
TUNIT	21.9	18.0	17.9	11.8	10.7	9.2
★Liu <i>et al.</i>	19.3	19.2	20.1	26.7	28.2	26.3
★StarGAN v2	15.6	20.6	18.8	22.4	26.8	25.7

Table 2. User study. Q1: content and style. Q2: realism. Q3: preference. A star (★) denotes models trained with extra information.

Method	Runtime (sec)	AFHQ		CelebA-HQ	
		MSE	LPIPS	MSE	LPIPS
Ours	0.029	0.012	0.269	0.007	0.202
SwapAE	0.037	0.009	0.303	0.005	0.241
StyleMapGAN	0.092	0.039	0.316	0.026	0.255

Table 3. Quantitative comparison for real image projection. We used a single NVIDIA Xp GPU to measure the runtime.

Method	Param. (M)	k-NN \uparrow	mFID \downarrow
Ours-AdaIN	57.75	99.1	14.7
separated	76.77	80.6	159.7

Table 4. Quantitative comparison using the AFHQ dataset.

local image translation. The first two rows are the result of using vertically split masks. The red box in the bottom row indicates the mask for reference 1. The proposed method can synthesize the content using multiple styles.

4.4. Analysis

Effect of the style-aware discriminator We trained the model with a separated discriminator and style encoder to analyze the effect of integrating the discriminator and style encoder. The difference is that we used the hard-assigned prototypes as pseudo-label for the multi-task discriminator. To evaluate the alignment between learned style representation and domain labels, we measured the k-NN accuracy used for self-supervised learning [7, 16]. In Table 4, *separated* achieved significantly lower k-NN accuracy,




Figure 5. Similarity search results on the AFHQ and CelebA-HQ datasets. We projected the query and test set into the style space and performed a nearest neighbor search. We plot here the five most similar images in style space.




Figure 6. Comparison of the results for various augmentations.

and failed to reflect the style of the target images (high mFID). See Appendix C.1 for a further discussion.

Effect of data augmentation We employed random resized crop, rotation, and scale for augmentation, along with random erasing for facial datasets (e.g., CelebA-HQ, FFHQ). Among them, we analyzed the effect of color distortion and cutout, which are major differences compared with other methods [3, 27]. As shown in Fig. 5, different augmentation choice leads to different style space. This result further leads to incorrect or unwanted synthesis results (Fig. 6). For example, when the color distortion is used, the style space ignores the color. On the other hand, if the cutout is not applied in the human face domain, learned style space failed to capture the attribute information such as gender.

Speed and memory Table 5 shows the trainable parameter counts and the training time of each method. The proposed approach is more efficient and faster than conventional I2I translation methods because it requires one less module for training and has fewer modules than SwapAE, which uses two discriminators. Nevertheless, the proposed method achieved comparable or better performance, which shows the efficiency of our method.

5. Discussion and limitation

In this study, we proposed a *style-aware discriminator*, which learns a style space in a self-supervised manner and guides the generator. Here, we discuss reasons why the proposed approach can be successfully trained. First, rep-

Method	Parameters (M)			sec/iter
	G	D	E	
Ours	36.8	19.7		0.383
Ours-AdaIN	38.1	19.7		0.351
StarGAN v2 [11]	43.5	20.9	20.9	0.678
TUNIT [3]	27.4	71.0	9.3	0.667
CLUIT [27]	34.4	25.2	20.9	1.016
SwapAE [36]	25.1	53.4	30.6	0.692
StyleMapGAN [26]	79.7	28.9	17.6	1.475

Table 5. Efficiency of proposed method. We measured the training speed (s/iter) with minibatch size 2 on a single TITAN Xp GPU.

resentation learning using human-defined labels cannot be a representation for style space. In contrast, the proposed method learns latent space specifically designed for style. Second, in the existing I2I translation, both the generator and the style encoder are updated together by the signal from the discriminator. In this case, the separation between content and style is ambiguous. Conversely, the proposed model can have a separate content space with the style encoder being updated completely separately from the generator, which results in better disentanglement. Finally, a style-aware discriminator can provide a better signal to the generator since it has a better understanding of the style space.

Yet still, the proposed method cannot preserve the face identity of the source image, unlike [11, 32]. One can therefore consider using a pre-trained network for identity or landmark following previous works [11, 39]. However, preserving the identity may increase risks of misuse or abuse. Therefore, we did not force the proposed method to preserve the facial identity of a source image. Though, preserving the facial identity without using additional information (e.g., face landmark or id) will be a valuable future work.

Acknowledgements This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.B0101-15-0266, Development of High Performance Visual BigData Discovery Platform for Large-Scale Realtime Data Analysis) and (No.2017-0-00897, Development of Object Detection and Recognition for Intelligent Vehicles)

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *ICCV*, 2019. 2
- [2] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan++: How to edit the embedded images? In *CVPR*, 2020. 2
- [3] Kyungjune Baek, Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Hyunjeong Shim. Rethinking the truly unsupervised image-to-image translation. In *ICCV*, 2021. 2, 3, 4, 6, 8
- [4] Hyojin Bahng, Sunghyo Chung, Seungjoo Yoo, and Jaegul Choo. Exploring unlabeled faces for novel attribute discovery. In *CVPR*, 2020. 2
- [5] Mikołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. In *ICLR*, 2018. 4
- [6] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020. 1, 3
- [7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 7
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 3
- [9] Ting Chen, Xiaohua Zhai, Marvin Ritter, Mario Lucic, and Neil Houlsby. Self-supervised gans via auxiliary rotation loss. In *CVPR*, 2019. 2
- [10] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *CVPR*, 2018. 1, 2
- [11] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *CVPR*, 2020. 1, 2, 3, 4, 5, 6, 8
- [12] Alex Clark. Pillow (pil fork) documentation, 2015. 5
- [13] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *NeurIPS*, 2013. 3
- [14] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 3
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. In *NeurIPS*, 2014. 1, 2, 3
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 3, 7
- [17] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017. 2, 4
- [18] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, 2017. 1, 3
- [19] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, 2018. 3, 4
- [20] Eunyeong Jeon, Kunhee Kim, and Daijin Kim. Fa-gan: Feature-aware gan for text to image synthesis. In *ICIP*, 2021. 2
- [21] Jongheon Jeong and Jinwoo Shin. Training gans with stronger augmentations via contrastive discriminator. In *ICLR*, 2020. 2
- [22] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *ICLR*, 2018. 1, 4
- [23] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *NeurIPS*, 2020. 2
- [24] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 1, 2, 4
- [25] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, 2020. 2, 3, 4
- [26] Hyunsu Kim, Yunjey Choi, Junho Kim, Sungjoo Yoo, and Youngjung Uh. Exploiting spatial dimensions of latent in gan for real-time image editing. In *CVPR*, 2021. 2, 3, 4, 5, 6, 8
- [27] Hanbit Lee, Jinseok Seol, and Sang goo Lee. Contrastive learning for unsupervised image-to-image translation. *arXiv preprint arXiv:2105.03117*, 2021. 2, 3, 4, 6, 8
- [28] Junnan Li, Pan Zhou, Caiming Xiong, and Steven Hoi. Prototypical contrastive learning of unsupervised representations. In *ICLR*, 2020. 3
- [29] Ru Li, Shuaicheng Liu, Guangfu Wang, Guanghui Liu, and Bing Zeng. Jigsawgan: Self-supervised learning for solving jigsaw puzzles with generative adversarial networks. *arXiv preprint arXiv:2101.07555*, 2021. 2
- [30] Bingchen Liu, Yizhe Zhu, Kunpeng Song, and Ahmed Elgammal. Towards faster and stabilized gan training for high-fidelity few-shot image synthesis. In *ICLR*, 2020. 2
- [31] Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, Jaakko Lehtinen, and Jan Kautz. Few-shot unsupervised image-to-image translation. In *ICCV*, 2019. 1, 2
- [32] Yahui Liu, Enver Sangineto, Yajing Chen, Linchao Bao, Haoxian Zhang, Nicu Sebe, Bruno Lepri, Wei Wang, and Marco De Nadai. Smoothing the disentangled latent style space for unsupervised image-to-image translation. In *CVPR*, 2021. 1, 4, 5, 6, 8
- [33] Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. Which training methods for gans do actually converge? In *ICML*, 2018. 3
- [34] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics and Image Processing*, 2008. 1, 4
- [35] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 3

- [36] Taesung Park, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *ECCV*, 2020. 4, 8
- [37] Taesung Park, Jun-Yan Zhu, Oliver Wang, Jingwan Lu, Eli Shechtman, Alexei A. Efros, and Richard Zhang. Swapping autoencoder for deep image manipulation. In *NeurIPS*, 2020. 2, 3, 4, 6
- [38] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On buggy resizing libraries and surprising subtleties in fid calculation. *arXiv preprint arXiv:2104.11222*, 2021. 5
- [39] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *ICCV*, 2021. 8
- [40] Kuniaki Saito, Kate Saenko, and Ming-Yu Liu. Coco-funit: Few-shot unsupervised image translation with a content conditioned style encoder. In *ECCV*, 2020. 1, 2
- [41] 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. 1, 4
- [42] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 4, 5
- [43] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *ECCV*, 2020. 2