Transforming the Latent Space of StyleGAN for Real Face Editing

Heyi Li[a,*], Jinlong Liu[b,**], Xinyu Zhang[c], Yunzhi Bai[b], Huayan Wang[b], Klaus Mueller[c]

[a]*Institute of Remote Sensing, China Academy of Space Technology, Beijing, China*
[b]*Y-tech, Kuaishou Technology, Beijing, China*
[c]*Department of Computer Science, Stony Brook University, Stony Brook, NY, USA*

Abstract

Despite recent advances in semantic manipulation using StyleGAN, semantic editing of real faces remains challenging. The gap between the W space and the $W+$ space demands an undesirable trade-off between reconstruction quality and editing quality. To solve this problem, we propose to expand the latent space by replacing fully-connected layers in the StyleGAN's mapping network with attention-based transformers. This simple and effective technique integrates the aforementioned two spaces and transforms them into one new latent space called $W++$. Our modified StyleGAN maintains the state-of-the-art generation quality of the original StyleGAN with moderately better diversity. But more importantly, the proposed $W++$ space achieves superior performance in both reconstruction quality and editing quality. Despite these significant advantages, our $W++$ space supports existing inversion algorithms and editing methods with only negligible modifications thanks to its structural similarity with the $W/W+$ space. Extensive experiments on the FFHQ dataset prove that our proposed $W++$ space is evidently more preferable than the previous $W/W+$ space for real face editing. The code is publicly available for research purposes at `https://github.com/AnonSubm2021/TransStyleGAN`

Keywords: Generative adversarial network, semantic editing, explainable

---

[*]Equal contribution
[**]Corresponding author
  Email address: `liujinlong@kuaishou.com` (Jinlong Liu)

artificial intelligence, disentanglement, Transformer

1. Introduction

The tremendous success of Generative Adversarial Networks (GANs) [1] has revolutionized the field of data-driven image generation and has sparked significant research attention. For human face synthesis in particular, the current state-of-the-art architecture StyleGAN [2, 3] generates high resolution (i.e. 1024×1024 pixel), photo-realistic images by first mapping the latent code to layer-wise style code and then feeding it into each convolution layer. Added via adaptive instance normalization (AdaIN) or its improved technique called weight demodulation, this style code S directly controls image features at various scales. The disentanglement of different attributes (coarse, medium, and fine) exhibited in the intermediate latent space is further explored by numerous follow-up works [4, 5, 6, 7, 8] to achieve semantically controllable human face generation. However, this manipulation capability is not directly applicable to real faces.

To mitigate this problem, an "invert then edit" methodology is adopted [4, 5, 9, 10, 11, 12, 13]. A real image is first projected into the latent space of StyleGAN. Then new latent codes are obtained by performing semantically meaningful edits on inverted latent codes. However, projected latent codes in the W space are not adequate for accurate regeneration of original images. The extended latent space $W+$ has been shown to be more powerful for inversion [14]. But editing latent codes in the $W+$ space is notoriously ill-posed because they fall out of the semantically meaningful manifold. In contrast, latent codes in the W space do not suffer from this problem and thus favor editing quality. This divide between the W space and the $W+$ space, unfortunately, mandates a compromise between reconstruction accuracy and manipulation naturalness.

In this paper, we present a straightforward but crucial enhancement to the StyleGAN architecture which expands the dimension of its latent space for better real face editing. We achieve this by replacing the fully-connected layers in
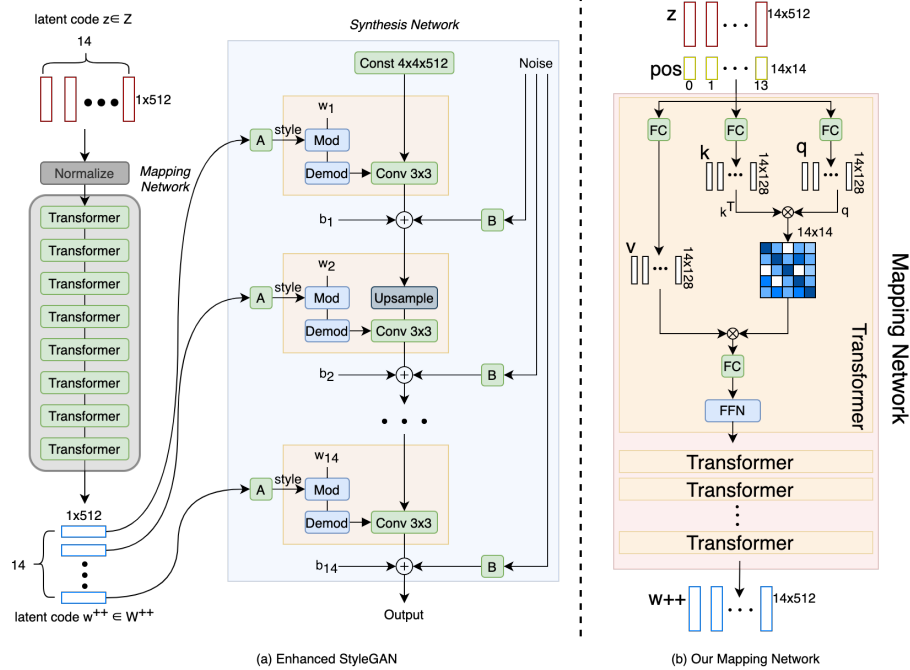
Figure 1: (a) **Model overview.** While the original StyleGAN [3] feeds the same style code to each convolution layer, we redesign the architecture of the StyleGAN mapping network using Transformer so that different style codes are provided to different layers. We illustrate the structure for resolution at 256×256. (b) An illustration of our proposed mapping network. We omit some standard layers that are not essential for our architecture to not clutter the exposition.

the mapping network M with attention-base transformer structures [15]. This modified mapping network M' maps a random input vector in the Z space to a new intermediate latent space named $W++$. The new $W++$ space has the same dimension as the $W+$ space, which is much larger than the W space and thus achieves a finer reconstruction quality. Unlike latent codes in the $W+$ space, latent codes in the proposed $W++$ space do not fall out of the semantically meaningful manifold after manipulation because they are directly utilized for image synthesis during training. As a result, our $W++$ space enjoys a better editing quality. In addition, our enhanced StyleGAN maintains the state-of-the-art performance of the original StyleGAN with a moderate improvement in

3

terms of generation diversity. Lastly, our proposed $W++$ space readily works with existing inversion algorithms and editing methods requiring only minor adjustments.

To demonstrate our $W++$ space's outstanding compatibility, we apply a few prevailing methods for inversion and editing, respectively. For the task of inversion, we choose the StyleGAN projector [3] to represent optimization-based methods and the pixel2style2pixel (pSp) [16] to illustrate learning-based methods. For the task of editing, we start with the InterfaceGAN [4, 5] for single attribution manipulation. Inspired by the idea of conditional GAN (cGAN) [17], we then propose a cGAN-based pipeline for attribute editing of real faces. Unlike traditional cGANs, our pipeline (see Figure 6) uses attributes extracted by pre-trained models as conditioning information. Another difference is that our discriminator learns to distinguish real and fake images without any conditioning information and thus is same as one in a naive GAN.

Our main contributions are summarized as follows:

- We propose to enhance the StyleGAN model by replacing fully-connected layers in its mapping network with attention-based transformers. This augmentation keeps the state-of-the-art generation quality and also moderately improves the generation diversity.

- Our proposed latent space $W++$ achieves superior performance in both reconstruction quality and editing quality.

- Our new $W++$ space offers excellent compatibility with existing algorithms for both inversion and editing. Only minor changes are needed.

2. Related Work

In the following, we briefly describe the existing literature on the subject of real face editing. The "invert then edit" methodology has become the de facto standard in this active research field. In the first inversion step, a given image is projected back into the latent space of StyleGAN and the inverted

4

latent code can faithfully reconstruct the input image through the generator. Then new latent codes are obtained by changing the old ones along semantically meaningful directions in the latent space.

2.1. GAN Inversion

The recent survey by Xia et al. [18] provides an exhaustive overview of GAN inversion algorithms. Existing methods can be classified into three main categories: optimization-based, learning-based, and a hybrid approach.

Optimization-based algorithms iteratively improve a latent code to minimize the error for a given image. Using an additive ramped-down noise, the original StyleGANv2 [3] proposes to embed images in the W space which enables better editing at the cost of worse reconstruction. To the contrary, Image2StyleGAN and Image2StyleGAN++ [14, 9] embed images into the extended $W+$ space which effectively optimizes a separate style for each scale. This approach sacrifices editing quality for reconstruction quality. To find a better balance, PIE [12] and StyleGAN2Encoder [19] adopt a two-stage encoding process which first embeds an image in the W space and then refines its initial latent code by optimizing in the $W+$ space.

Learning-based methods [20, 16, 21] aim to train an encoder network which maps an image to the latent space. Compared with optimization-based algorithms, learning-based methods have the advantage of low computation complexity but suffer from inferior reconstruction quality.

The hybrid type such as [10] combines the above two techniques where an encoder network is first used to obtain an approximate latent code and then this latent code is improved with optimization.

2.2. Latent Space Manipulation

Given its approximate linearity, the latent space of StyleGAN has been the primary target for semantic manipulation. Supervised methods find linear directions that correspond to changes in a given binary labeled attribute (such as

young *vs.* old) with the supervision of semantic annotations. StyleRig [22] utilizes a pre-trained 3DMM to find the mapping between rigging information and face manipulation. StyleFlow [11] modifies a set of predetermined attributes by learning a transformation between different vectors in the $W+$ space. InterfaceGAN [4, 5] trains linear support vector machines (SVMs) to classify latent codes based on semantic labels and uses the normal vector of each hyperplane as the latent direction of the selected attribute.

To facilitate attribute manipulation in an unsupervised manner, GANSpace [6] performs PCA on the sampled data to find primary directions in the latent space. Conversely, Collins *et al.* [7] discovers the connection between local semantics and components of latent codes using k-means clustering. Finally, SeFa [23] is a closed-form factorization method which computes interpretable directions without any kinds of training or optimization.

2.3. Other Space

Most recently, a few concurrent works [24, 25, 26] have been proposed to address this "reconstruction-editing" conundrum by exploring other space options.

Although having achieved a satisfactory trade-off, [24] fails to eliminate the fundamental conflict caused by performing reconstruction and editing in two different latent spaces. Their proposed P space is transformed from the W space by inverting the last Leaky ReLU layer in the StyleGAN mapping network, while the $P+$ space is extended from the P space by concatenation in a similar way as the $W+$ space being extended from the W space.

Instead of the latent space, [25, 26] investigate the style space which is spanned by all possible style vectors. However, as is pointed out in [26], the style space achieves worse manipulation naturalness and closer reconstruction when compared with the $W+$ space. This, unfortunately, even aggravates the problem that we are trying to solve.

Table 1: Quantitative comparison of our model and the original StyleGAN on generation quality with the FFHQ dataset. ↑ indicates that higher is better, and ↓ that lower is better.

Model	FID ↓	Precision ↑	Recall ↑	Density ↑	Coverage ↑
StyleGANv2	4.69	**0.670**	0.430	**1.278**	0.955
Ours	**4.67**	0.662	**0.454**	1.261	**0.961**
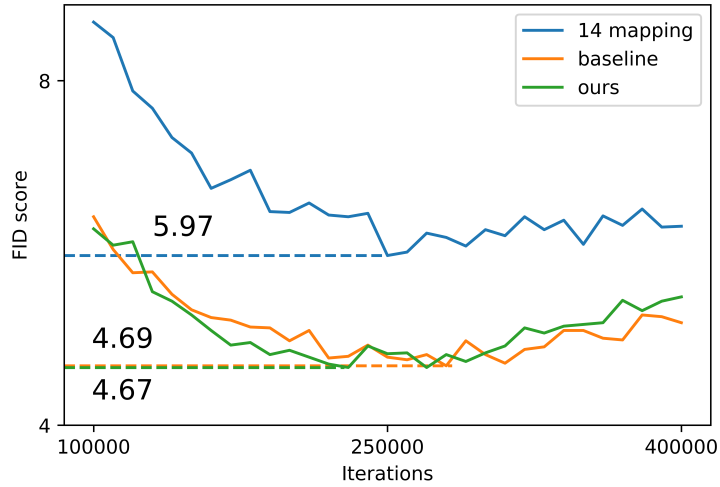
Figure 2: FID score curve for different StyleGAN models at 256×256 resolution during training. The StyleGAN with 14 independent mapping networks reaches its optimal score 5.97 at Iteration 250K. Our enhanced StyleGAN model achieves the best FID score 4.67 at Iteration 230K.

3. Method

3.1. Motivation

The fundamental reason for having to sacrifice either reconstruction quality or editing quality for the other is the fact that these two operations are best suitable in two separate latent spaces. The original StyleGAN uses the same style vector for all varied scales during image generation, which essentially constrains the dimension of the intermediate latent space W to be 1×512. By traversing the W space, we can find semantically meaningful directions for editing. During inversion, on the other hand, such restriction is lifted which enlarges the dimen-

7

sion of the $W+$ space to be $k \times 512$ (k is the number of style codes). This extra flexibility allows for a more faithful restoration, but it also distorts all editing directions. To resolve this contradiction, we need to expand the input space Z to be of dimension $k \times 512$ and in turn expand the intermediate latent space W. But how do we achieve that? (For simplicity and conformity, we fix the image resolution to 256×256 and thus set the value of k to 14 for the following analysis.)

One straightforward solution is to expand the dimension of fully-connected layers in the mapping network from 512 to 7168 (which is equal to 14×512). Although the dimension of the intermediate latent space is successively increased in this way, the computation cost also grows quadratically. The estimated FLOPs of a fully-connected layer is 512×512. And the estimated FLOPs of an expanded fully-connected layer becomes 7168×7168, which is 196 times of an original one. And it grows to 324 times at 1024×1024 resolution as k reaches 18.

Another possible solution is to use an individual mapping network for each layer, which avoids the explosion of computational complexity occurred above. This means 14 independent mapping networks each creating a different style vector. Since each new mapping network is structurally identical to the original one, this approach limits the computational growth to linear. However, as is shown in Figure 2, the generation quality indicated by the FID metric deteriorates significantly. The main reason for this degradation is the absence of any correlation among style vectors. In the original StyleGAN, consistency of global features across scales is ensured by feeding the same style vector to all layers in the synthesis network. However, using fully independent mapping networks to compute style vectors completely prohibits such a correlation.

3.2. W++ Space

An ideal solution to this problem should therefore satisfy the following two conditions: (i) limited growth of computational cost; (ii) some degree of correlation among style vectors. The attention-based transformer structure [15] fits

8

these two conditions naturally.

A detailed illustration of our mapping network is provided in Figure 1(b). We follow the original design of the transformer with a multi-head self-attention and a simple position-wise fully connected feed-forward network. A latent code z is randomly sampled from the input latent space Z. This Z space is almost identical to its counterpart in the original StyleGAN except that its dimension is enlarged to 14×512. Then a fixed 1D position embedding is added to retain the layer index of each style vector [27]. The query (q), the key (k), and the value (v) all derive from this same latent code z concatenated with position embedding through separate linear transformations. According to the ablation study by Karras et al [2] on the depth of the mapping network, 8 is the optimal choice. Therefore, we implement our mapping network by stacking 8 transformers. Outputs of this network form our new latent space, **W++**.

Computation cost. In our case, the overall computational complexity of the transformer is dominated by linear transformations instead of the similarity calculations. Each linear transformation is bounded by $O(kd^2)$ where the dimension d is 512 and the number of style codes k is 14. We also add a compression ratio c as a trick to reduce the total computation, which lowers the complexity to $O(k(\frac{d}{c})^2)$; in practice, we set the value of c to 4. In total, our $W++$ space introduces an approximately five-fold increase in computational cost for the mapping network. However, such expansion has limited impact on the whole model given the fact that most of the StyleGAN's computation burden falls upon the synthesis network and the discriminator.

Style correlation. The cosine similarity between two linear transformations of the same input, the query and the key, is calculated by the self-attention module. This similarity matrix is then multiplied with the value which is another linear transformation of the input. Therefore, the latent code in our $W++$ space is essentially a weighted sum of the input latent code in the Z space along the k ($= 14$) dimension. All coefficients are learned during the training stage. Hence, the mapping network of the original StyleGAN is a special case of ours where

Table 2: Quantitative comparison on image reconstruction with StyleGAN projector [3] in different latent spaces. Reported results are average values over 400 images. ↓ means lower number is better.

Metrics	inversion in W space	inversion in $W+$ space	inversion in $W++$ space
avg Perceptual loss ↓	0.2079 ± 0.0389	0.0855 ± 0.0177	0.0937 ± 0.0179
avg MSE loss ↓	0.0776 ± 0.0512	0.0217 ± 0.0179	0.0255 ± 0.0183

correlation between any two elements in the k dimension is rigidly constrained to 1. Using 14 independent mapping networks is at the other end of the spectrum because the correlation is then equivalent to 0. Depending on the input values, our correlation coefficients fall within $[0, 1]$ after training and vary for different pairs.

As is shown in Figure 1(a), the output latent code is divided into 14 different style codes along the k dimensions. Each style code with a dimension of 512 is then fed to a different layer in the synthesis network at different scale. No changes are made to the original styleGAN architecture to accommodate the proposed $W++$ space.

4. Experiments

We evaluate our proposed $W++$ space using the FlickrFaces-HQ (FFHQ) dataset [2] from three aspects: generation quality, reconstruction quality, and editing quality. All experiments are performed using PyTorch at 256×256 resolution with 8 NVIDIA Tesla V100 GPUs.

4.1. Image Synthesis

We first evaluate our proposed $W++$ space by comparing it with the original StyleGANv2 [1]. We train both models from scratch using exactly the same

---

[1] Our work builds upon the PyTorch implementation of StyleGANv2 by rosinality, which is publicly available at `https://github.com/rosinality/stylegan2-pytorch`
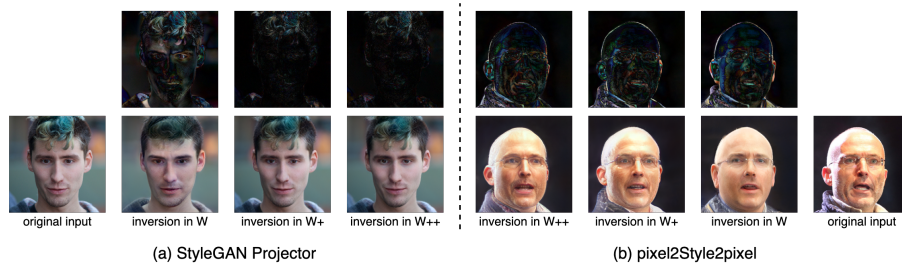
Figure 3: (a) Qualitative comparison on image reconstruction with StyleGAN projector [3] in different latent spaces. (b) Qualitative comparison on image reconstruction with pSp encoder [16] in different latent spaces. Top row displays differences between inverted images and original images.

hyperparameters. The generation quality is measured by Frechet inception distances (FID) [28] and the diversity is reflected by Recall [29] and Coverage [30] [2]. The results are provided in Table 1.

Our modified StyleGAN achieves comparable generation quality than the original StyleGANv2 [3]. Furthermore, there is a slight shift from precision to recall as our model achieves higher recall and lower precision. But as is pointed out in [3], this is generally desirable because recall can be traded into precision via truncation, whereas the opposite is not true. A similar phenomenon is observed in the density and coverage pair, where density measures fidelity and coverage quantifies diversity.

In summary, our new StyleGAN maintains the state-of-the-art generation quality of the original StyleGAN. Also, our model gains slightly in terms of diversity thanks to its larger latent space Z.

---

[2]The code is publicaly available at https://github.com/clovaai/generative-evaluation-prdc

[3]The best FID score announced using rosinality's Pytorch implementation at 256×256 resolution is 4.5. While the best FID score we have achieved after multiple runs is 4.69 at the same resolution. However, this performance gap does not affect our findings because the training codes are identical.

4.2. Real Face Inversion

In this experiment, we compare inversion in the proposed $W++$ space with both W and $W+$ space. We choose two well-established inversion algorithms, the projector from [3] as an example for the optimization-based methods, and the pSp [16] as an instance for the learning-based methods. Apart from its outstanding reconstruction quality, we also demonstrate that our $W++$ space readily supports existing works.

4.2.1. Optimization-Based

Although initially proposed for inversion in the W space, the StyleGANv2's projector has been extended to support inversion in the $W+$ space as well. Since the $W++$ space shares the same dimension as the $W+$ space, this projector easily works in our new space. Quantitative evaluations between inversion in the W space, the $W+$ space, and our $W++$ space are provided in Table 2. Figure 3(a) displays qualitative comparisons for visual inspection.

For quantitative examination, we choose two metrics, the mean square error (MSE) loss and the perceptual loss defined as the LPIPS distance [31] between target images and restored images. We perform optimization on a total of 400 images and report the average value for both metrics in all three spaces. Our proposed $W++$ space outperforms the W space by a large margin in both metrics.

For qualitative examination, the W space clearly performs the worst in reconstruction quality. Due to its low dimension, it simply fails to faithfully restore the identity of the person in the input image. The $W+$ and our $W++$ space, on the other hand, do not suffer from this drawback. Inverted images in these two spaces are almost indistinguishable because they have the same dimensions. These findings are also confirmed by the element-wise difference map shown in the top row of Figure 3(a).

Therefore, our proposed $W++$ space clearly outperforms the W space in reconstruction quality while it achieves comparable results with the $W+$ space.

Table 3: Quantitative comparison on image reconstruction with pSp [16] in different latent spaces. Reported results are average values over 400 images. ↓ means lower number is better.

Metrics	inversion in W space	inversion in $W+$ space	inversion in $W++$ space
avg Perceptual loss ↓	0.4158 ± 0.0590	0.3557 ± 0.0517	0.3641 ± 0.0530
avg MSE loss ↓	0.0661 ± 0.0298	0.0398 ± 0.0204	0.0412 ± 0.0204

4.2.2. Learning-Based

Our proposed $W++$ space also supports existing learning-based inversion algorithms such as pSp [16] [4] effortlessly. The only modification necessary is to replace the previous StyleGAN model with our enhanced one. We follow the default hyperparameters for training.

The quantitative and qualitative results are shown in Table 3 and Figure 3(b) respectively. We adopt the same numerical metrics as above. In accordance with findings in previous works, the reconstruction quality of the learning-based pSp is worse than the reconstruction quality of the optimization-based StyleGAN projector. Again, our proposed $W++$ space accomplishes equivalent superiority with the $W+$ space in reconstruction quality.

4.3. Real Face Editing

After projecting real images into the latent space, we perform manipulation in the obtained latent codes for semantic editing. Again, we choose two methods to prove the advantage of our proposed $W++$ space in editing quality. One is the well-known InterfaceGAN [4, 5] and the other is a conditional GAN-based editing pipeline which uses pre-trained models as attribute extractors.

4.3.1. InterfaceGAN

Following the procedure in [4, 5], we train a linear support vector machine (SVM) [5] to find the editing direction which manipulates a real image.

---

[4]The code is publicly available at `https://github.com/eladrich/pixel2style2pixel`.

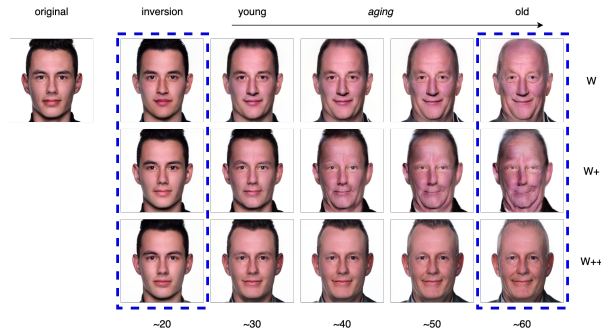[5]The code is publicly available at `https://github.com/genforce/interfacegan`.

Figure 4: Qualitative comparison of age transformation using InterfaceGAN [4, 5] in different latent spaces. The age label for each image is created using the pre-trained DEX [32] model. Columns identified by blue boxes are on display in Figure 6. "$\sim N$" denotes "approximately N years old". Our results (the bottom row) achieve considerably stronger robustness for long-distance manipulation.

100000 random latent codes are first generated and then a pre-trained predictor network, such as DEX [32] [6], is selected to create attribute labels for their corresponding images. Afterwards, these data pairs are used to train the linear SVM. For this demonstration, we use the StyleGANv2's projector to obtain latent codes for real images.

Since the step size plays a crucial role in latent space manipulation, we first conduct an empirical study to align both the visual effect and the classification score of edited images in different spaces. Figure 4 and Figure 5 display the intermediate frames of the entire age transformation and gender transitioning respectively. The first column shows the inverted images and subsequent columns contain images edited to different extents. As the editing distance between the latent code and the classifying boundary keeps increasing, the editing quality of both the W and the $W+$ space starts to deteriorate, especially for editing in the $W+$ space. For gender manipulation in particular (see Figure 5), changes in the gender direction also transforms the person's race from white to Asian. In contrast, our proposed $W++$ space does not suffer from

---

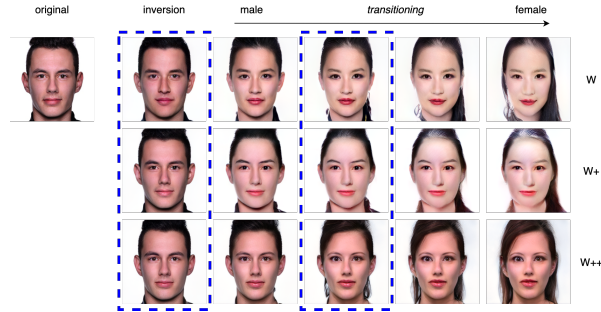[6]The code is publicly available at https://github.com/siriusdemon/pytorch-DEX.

Figure 5: Qualitative comparison of gender transitioning results using InterfaceGAN [4, 5] in different latent spaces. Columns identified by blue boxes are on display in Figure 6. Our results (the bottom row) achieve considerably stronger robustness for long-distance manipulation.

the "distance effect" problem and shows considerably stronger robustness for long-distance manipulation.

The blue boxes indicate the columns we select for each attribute and the final results are shown in Figure 6. The results in the $W+$ space (see the middle row) clearly exhibits the worst editing quality in both cases. Although enjoying a better editing quality, the W space suffers from the "distance effect" problem [4, 5] where only near-boundary manipulation works well. When the latent code goes further from the boundary, the quality of edited images deteriorates dramatically as manipulating one attribute starts to affect others. In comparison, our proposed $W++$ space (see the bottom row) achieves an exceedingly more satisfying editing quality. The improvement is extremely obvious in the third column where images are modified to transit gender. In this case, the race of the target person changes from white to Asian when moving the latent code along the gender direction in the W space. Edited images in our $W++$ space faithfully preserve personal identity even for long-distance manipulations.

Two more examples for age transformation are provided in Figure 10 and Figure 11. Figure 12 and Figure 13 exhibit two extra examples for gender transitioning.
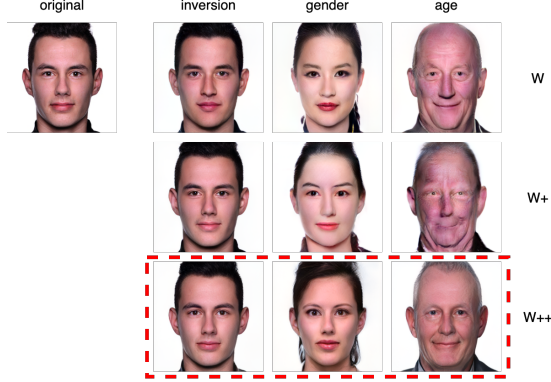
Figure 6: Manipulating real faces with respect to the attribute age in different latent spaces. Given a real image to edit, we first invert it back to the latent space using StyleGAN projector [3] and then manipulate the latent code with InterFaceGAN [4, 5]. Our results (highlighted by the red box) achieve considerably stronger robustness for long-distance manipulation.

4.3.2. cGAN-based Pipeline

To further demonstrate the superiority of our proposed $W++$ space, we utilize a conditional GAN-based pipeline for attribute manipulation of real faces. As is shown in Figure 7, our pipeline adds conditioning information to the StyleGAN generation process. During training, the pre-trained attribute extractor first extracts attribute information from target images. Then this attribute code is concatenated with the latent code in our $W++$ space. Different style codes are concatenated with exactly the same attribute code and are fed to the synthesis network to generate outputs. Afterwards, the same attribute extractor extracts attribute information from generated images. An $L1$ attribute loss between the two attribute codes is minimized as a form of supervision. Unlike traditional cGAN, the discriminator in our case does not take any conditioning information as input. It is only trained to distinguish realistic-looking images from fake ones and thus is identical to the discriminator in StyleGANv2. Also, only the StyleGAN part (the mapping network, the synthesis network, and the discriminator) is trainable since the attribute extractor is fixed.

For editing real faces, the corresponding latent code for a given input is obtained through inversion. Different attribute classifiers are selected to extract
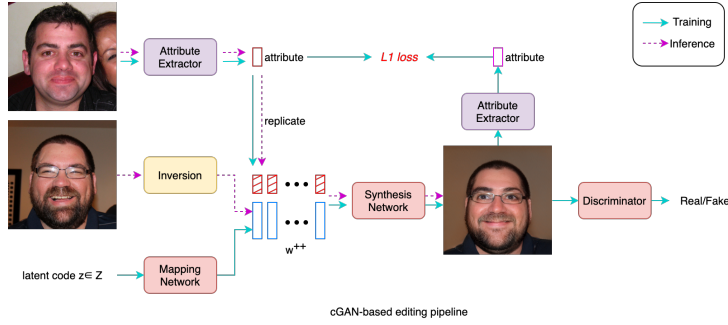
Figure 7: Our conditional GAN-based framework for editing real images. Here we cast one person's facial expression(*eyes open and mouth closed*) to another person.

different attributes (such as age, gender, smile, eyes open) from target images. Concatenated with various attribute codes, the synthesis network consequently generates images of the same person but with intended attributes.

In this experiment, we use the StyleGANv2's projector to invert real images into latent codes. The default hyperparameters are used to train projectors in all three spaces (W, $W+$, and $W++$) for a fair comparison. We utilize a pre-trained feature extractor which captures various facial movements essential for depicting a person's smile. The attribute code extracted is a $1D$ vector of 51 facial coefficients. For editing the smile expression in particular, we extract the attribute code from a single source image and apply it to all input images. The comparison results are shown in Figure 8.

Here we show two examples, one female and one male. Inverted images in the W space show the lowest quality, which in turn impairs the quality of edited images and makes the W space ill-suited for real image editing. The $W+$ space achieves excellent reconstruction quality but the worst editing quality. The smile expression after editing is indistinguishable. Our proposed $W++$ space, on the other hand, attains outstanding quality in both reconstruction and editing. The smile expression in the edited images are both apparent and natural. Additionally, our cGAN-based editing pipeline enables exceptional attribute disentanglement. As is shown in the last row of Figure 8, all irrelevant features remain unaltered. Additional results can be found in Figure 14 and
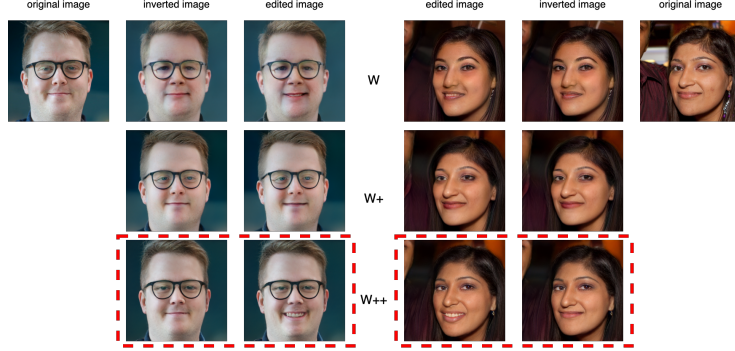
17

Figure 8: Manipulating real faces with respect to the attribute smile in different latent spaces. Given a real image to edit, we first invert it back to the latent space using StyleGAN projector [3] and then manipulate the latent code with our proposed cGAN-based editing pipeline. Our results (highlighted by the red box) exhibit the most apparent and natural smile expression.

Figure 15.

4.4. User Study

In addition, we conduct a user study to evaluate our proposed $W++$ space quantitatively. Participants are randomly assigned inverted images and edited images in either the W space or the $W+$ space, paired with corresponding images in the $W++$ space. Then they are asked to select the one of the highest quality according to their own judgements. In this way, we are able to determine which space generates more authentic images to the human eye.

4.4.1. Details of the User Study

The user study has 4 tasks, namely *inversion*, *aging*, *smiling* and *transition*. In each task, the original image is shown to the left along with two generated images to the right. One of the two images is in the proposed $W++$ space, while the other one is randomly picked either in the W space or the $W+$ space. The sequence of these two images shown is shuffled so that the image in the $W++$ space could either be in the second column or the third column. More importantly, all tasks in our study follow the double-blind procedure to guard

18

against any potential bias.

Inversion: In this task, the goal is to reconstruct an image from a latent code. 3 sets of images are randomly picked out of a total number of 12. The participant is then asked to select one image per set which looks more identical to the original image among the two.

Aging: In this task, the goal is to edit a portrait so that the person looks like at 60 years old. 3 randomly selected sets of images are shown. And the participant is asked to select the one more realistic per set.

Smiling: In this task, the goal is to edit a portrait so that the person puts on a smiling face. 2 sets of images are randomly picked out of a total number of 6. The participant is then asked to select the more genuine one per set among the two images shown.

Transition: In this task, the goal is to edit an image so that the person is transited to the opposite gender. 2 randomly selected sets of images are shown. And the participant is asked to select the one more authentic per set.

We recruit a total number of 41 users with diverse backgrounds. 21 of them are self-identified as male, while the rest is self-identified as female. There are 12 participants who are younger than 25, 19 between age 25 and 30, 7 between age 30 and 35, and 3 above the age of 35. 11 people has some experience in the field of computer vision, while the others do not.

4.4.2. Result Analysis

We designate the W space and the $W+$ space as 2 independent control groups and compare the result of the $W++$ space with each of them respectively. As is shown in Figure 9, our proposed $W++$ space is overwhelmingly more preferable among participants than either the W space or the $W+$ space in most tasks. For example, 91.7 percent of users conclude that the reconstruction quality of our $W++$ space is better than the W space(according to Figure 9(a)) and 96.3 percent of participants believe that the editing quality for aging in the proposed $W++$ space is better than the $W+$ space(according to Figure 9(b)). Figure 9(b) also shows that 55.6 percent of human subjects prefer the inverted
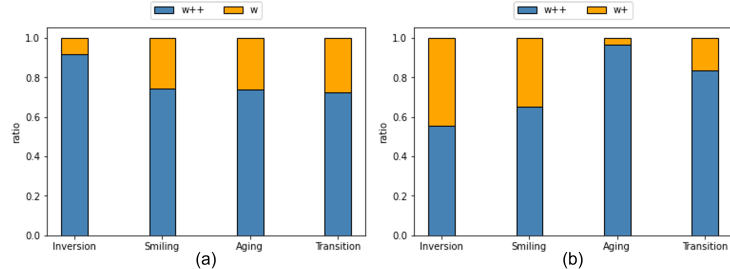
Figure 9: Stacked bar charts of the preference ratio (a) the $W++$ space *vs.* the W space (b) the $W++$ space *vs.* the $W+$ space

images in our $W++$ space to those in the $W+$ space, which resonates with the fact that these two spaces achieve comparable perceptual loss and MSE loss values.

To evaluate the statistical significance, we supplement our user study with a two-way ANOVA test and the result is shown in Table 4. The p values for the models in both tables are less than the threshold of 0.001 which indicates that the difference of $W++$ and the other model is highly statistically significant rather than caused by random factors. Therefore, our proposed $W++$ space is superior to both the W space and the $W+$ space with highly statistical significance. Furthermore, the p values for the tasks in both tables are more than the threshold of 0.05. This implies that our results are not attributable to the choice of tasks. As a result, the superiority of our proposed $W++$ space is universal across all tasks.

5. Conclusion

In this work, we propose to upgrade the StyleGAN architecture by replacing its mapping network with 8 attention-based transformers. This modification transforms its original latent space to a new latent space called $W++$. Our StyleGAN model retains the state-of-the-art generation quality and moderately improves generation diversity. However, unlike the previous W or $W+$ spaces, our proposed $W++$ space achieves superior performance in both reconstruction

Table 4: Two-way ANOVA analysis for (a) the $W++$ space *vs.* the W space (b) the $W++$ space *vs.* the W space to analyze the statistical significance.

Table (a)	SS	df	MS	F	P
Models	53.33	1	53.33	127.59	**< 0.0001**
Tasks	2.24	3	0.75	1.79	0.1494
Models*Tasks	6.68	3	2.23	5.33	0.0014
Error	109.51	262	0.42	-	-
Total	171.76	269	-	-	-
Table (b)	SS	df	MS	F	P
Models	36.11	1	36.11	99.77	**< 0.0001**
Tasks	1.76	3	0.59	1.62	0.1851
Models*Tasks	15.35	3	5.12	14.14	< 0.0001
Error	93.38	258	0.36	-	-
Total	146.6	265	-	-	-

quality and editing quality. Additionally, it supports existing inversion algorithms and editing methods with only minor adjustments needed. Experiments using FFHQ dataset clearly demonstrate the merits of our method.

Our work has some limitations that we leave to future work. Although our proposed $W++$ space has achieved excellent reconstruction quality by expanding the original latent space, the inverted image still looks a little different from the input real image. This dissimilarity adversely impacts the editing quality of real images. In the future, we would like to close this gap.
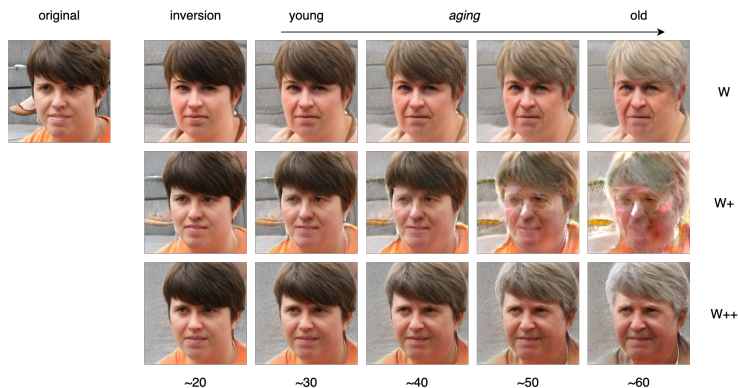
Acknowledgments

Figure 10: Qualitative comparison of age transformation using InterfaceGAN [4, 5] in different latent spaces. The age label for each image is created using the pre-trained DEX [32] model. "$\sim N$" denotes "approximately N years old". Our results (the bottom row) achieve considerably stronger robustness for long-distance manipulation.

References

[1] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, in: Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, K. Q. Weinberger (Eds.), Advances in Neural Information Processing Systems, Vol. 27, Curran Associates, Inc., 2014.

[2] T. Karras, S. Laine, T. Aila, A style-based generator architecture for generative adversarial networks, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 4401–4410.

[3] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, T. Aila, Analyzing and improving the image quality of stylegan, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 8110–8119.

[4] Y. Shen, J. Gu, X. Tang, B. Zhou, Interpreting the latent space of gans for semantic face editing, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 9243–9252.
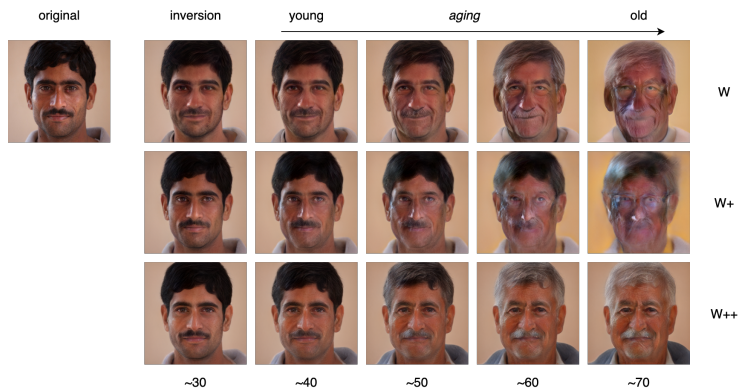
Figure 11: Qualitative comparison of age transformation using InterfaceGAN [4, 5] in different latent spaces. The age label for each image is created using the pre-trained DEX [32] model. "∼ N" denotes "approximately N years old". Our results (the bottom row) achieve considerably stronger robustness for long-distance manipulation.

[5] Y. Shen, C. Yang, X. Tang, B. Zhou, Interfacegan: Interpreting the disentangled face representation learned by gans, IEEE Transactions on Pattern Analysis and Machine Intelligence.

[6] E. Härkönen, A. Hertzmann, J. Lehtinen, S. Paris, Ganspace: Discovering interpretable gan controls, in: H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, H. Lin (Eds.), Advances in Neural Information Processing Systems, Vol. 33, Curran Associates, Inc., 2020, pp. 9841–9850.

[7] E. Collins, R. Bala, B. Price, S. Susstrunk, Editing in style: Uncovering the local semantics of gans, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 5771–5780.

[8] A. Shoshan, N. Bhonker, I. Kviatkovsky, G. Medioni, Gan-control: Explicitly controllable gans, arXiv preprint arXiv:2101.02477.

[9] R. Abdal, Y. Qin, P. Wonka, Image2stylegan++: How to edit the embedded images?, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 8296–8305.

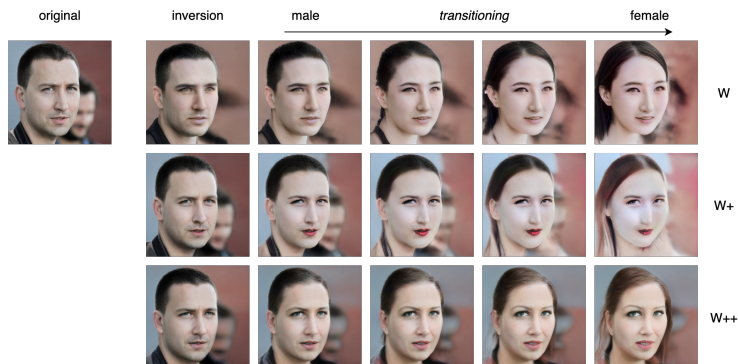[10] J. Zhu, Y. Shen, D. Zhao, B. Zhou, In-domain gan inversion for real image

Figure 12: Qualitative comparison of gender transitioning results using InterfaceGAN [4, 5] in different latent spaces. Our results (the bottom row) achieve considerably stronger robustness for long-distance manipulation.

editing, in: European Conference on Computer Vision, Springer, 2020, pp. 592–608.

[11] R. Abdal, P. Zhu, N. Mitra, P. Wonka, Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows, arXiv e-prints (2020) arXiv–2008.

[12] A. Tewari, M. Elgharib, F. Bernard, H.-P. Seidel, P. Pérez, M. Zollhöfer, C. Theobalt, Pie: Portrait image embedding for semantic control, ACM Transactions on Graphics (TOG) 39 (6) (2020) 1–14.

[13] X. Hou, X. Zhang, L. Shen, Z. Lai, J. Wan, Guidedstyle: Attribute knowledge guided style manipulation for semantic face editing, arXiv preprint arXiv:2012.11856.

[14] R. Abdal, Y. Qin, P. Wonka, Image2stylegan: How to embed images into the stylegan latent space?, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 4432–4441.

[15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, I. Polosukhin, Attention is all you need, in: I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett
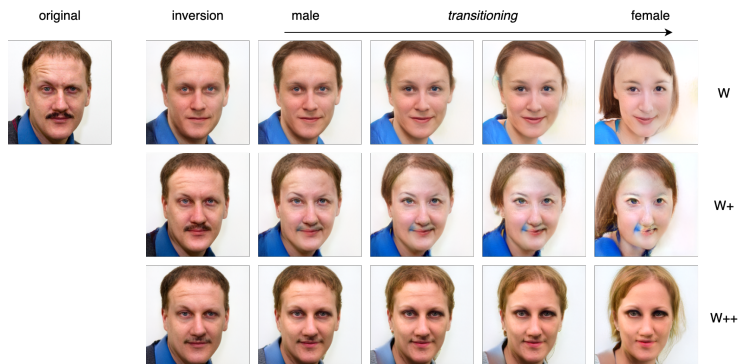
Figure 13: Qualitative comparison of gender transitioning results using InterfaceGAN [4, 5] in different latent spaces. Our results (the bottom row) achieve considerably stronger robustness for long-distance manipulation.

(Eds.), Advances in Neural Information Processing Systems, Vol. 30, Curran Associates, Inc., 2017.

[16] E. Richardson, Y. Alaluf, O. Patashnik, Y. Nitzan, Y. Azar, S. Shapiro, D. Cohen-Or, Encoding in style: a stylegan encoder for image-to-image translation, arXiv preprint arXiv:2008.00951.

[17] M. Mirza, S. Osindero, Conditional generative adversarial nets, arXiv preprint arXiv:1411.1784.

[18] W. Xia, Y. Zhang, Y. Yang, J.-H. Xue, B. Zhou, M.-H. Yang, Gan inversion: A survey, arXiv preprint arXiv:2101.05278.

[19] robertluxemburg, Git repository: stylegan2encoder https://github.com/robertluxemburg/stylegan2encoder.

[20] S. Guan, Y. Tai, B. Ni, F. Zhu, F. Huang, X. Yang, Collaborative learning for faster stylegan embedding, arXiv preprint arXiv:2007.01758.

[21] O. Tov, Y. Alaluf, Y. Nitzan, O. Patashnik, D. Cohen-Or, Designing an encoder for stylegan image manipulation, arXiv preprint arXiv:2102.02766.
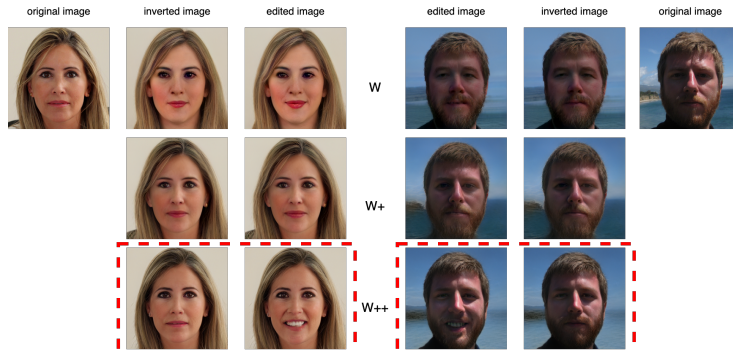
Figure 14: Manipulating real faces with respect to the attribute smile in different latent spaces. Given a real image to edit, we first invert it back to the latent space using StyleGAN projector [3] and then manipulate the latent code with our proposed cGAN-based editing pipeline. Our results (highlighted by the red box) exhibit the most apparent and natural smile expression.

[22] A. Tewari, M. Elgharib, G. Bharaj, F. Bernard, H.-P. Seidel, P. Pérez, M. Zollhofer, C. Theobalt, Stylerig: Rigging stylegan for 3d control over portrait images, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 6142–6151.

[23] Y. Shen, B. Zhou, Closed-form factorization of latent semantics in gans, arXiv preprint arXiv:2007.06600.

[24] P. Zhu, R. Abdal, Y. Qin, P. Wonka, Improved stylegan embedding: Where are the good latents?, arXiv preprint arXiv:2012.09036.

[25] Y. Liu, Q. Li, Z. Sun, T. Tan, Style intervention: How to achieve spatial disentanglement with style-based generators?, arXiv preprint arXiv:2011.09699.

[26] Z. Wu, D. Lischinski, E. Shechtman, Stylespace analysis: Disentangled controls for stylegan image generation, arXiv preprint arXiv:2011.12799.

[27] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit,
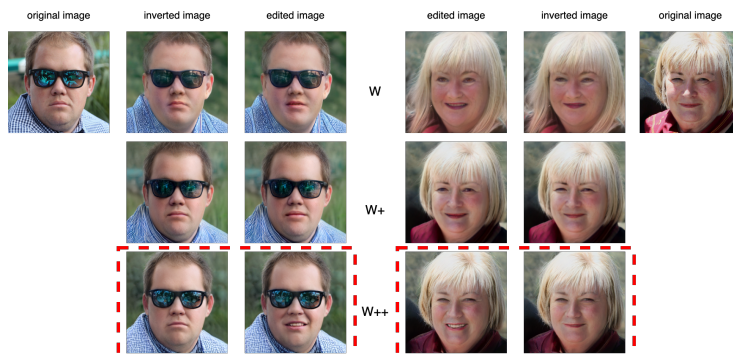
Figure 15: Manipulating real faces with respect to the attribute smile in different latent spaces. Given a real image to edit, we first invert it back to the latent space using StyleGAN projector [3] and then manipulate the latent code with our proposed cGAN-based editing pipeline. Our results (highlighted by the red box) exhibit the most apparent and natural smile expression.

N. Houlsby, An image is worth 16x16 words: Transformers for image recognition at scale, in: International Conference on Learning Representations, 2021.

[28] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, S. Hochreiter, Gans trained by a two time-scale update rule converge to a local nash equilibrium, in: I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (Eds.), Advances in Neural Information Processing Systems, Vol. 30, Curran Associates, Inc., 2017.

[29] T. Kynkäänniemi, T. Karras, S. Laine, J. Lehtinen, T. Aila, Improved precision and recall metric for assessing generative models, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (Eds.), Advances in Neural Information Processing Systems, Vol. 32, Curran Associates, Inc., 2019.

[30] M. F. Naeem, S. J. Oh, Y. Uh, Y. Choi, J. Yoo, Reliable fidelity and diversity metrics for generative models, in: International Conference on Machine Learning, PMLR, 2020, pp. 7176–7185.

[31] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, O. Wang, The unreasonable effectiveness of deep features as a perceptual metric, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 586–595.

[32] R. Rothe, R. Timofte, L. V. Gool, Dex: Deep expectation of apparent age from a single image, in: IEEE International Conference on Computer Vision Workshops (ICCVW), 2015.