

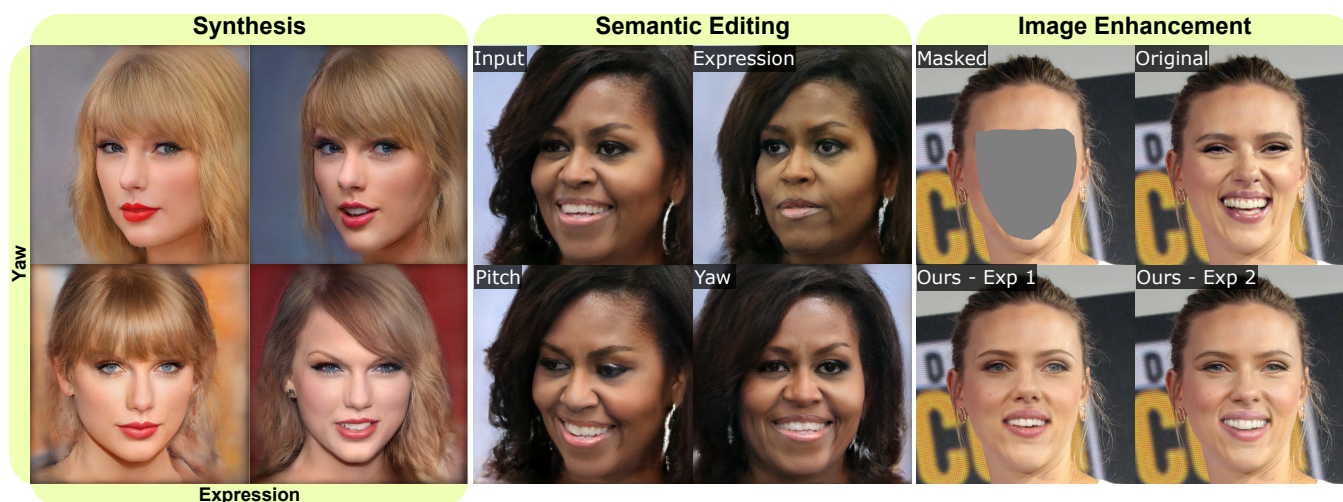
# MyStyle++: A Controllable Personalized Generative Prior

Libing Zeng  
libingzeng@tamu.edu  
Texas A&M University  
College Station, USA

Yi Xu  
yi.xu@oppo.com  
OPPO US Research Center  
Palo Alto, USA

Lele Chen  
lele.chen@sony.com  
Sony AI  
New York, USA

Nima Khademi Kalantari  
nimak@tamu.edu  
Texas A&M University  
College Station, USA



**Figure 1:** Our controllable personalized prior, trained on a collection of images of an individual, provides full control over a set of attributes, while generating results that accurately portray the facial features of that person. On the left, we demonstrate our method’s ability to synthesize images of Taylor Swift with user-defined expressions and yaw angles. In the middle, we show the editing results of our method for Michelle Obama. Our approach allows the user to directly generate an edited image with a set of desired attributes. Our method can also be used to enhance images with desired attributes, as shown on the right for image inpainting. Here, we show our inpainted results with two different expressions for an image of Scarlett Johansson. Middle and right block ©Gage Skidmore.

## ABSTRACT

In this paper, we propose an approach to obtain a personalized generative prior with explicit control over a set of attributes. We build upon MyStyle, a recently introduced method, that tunes the weights of a pre-trained StyleGAN face generator on a few images of an individual. This system allows synthesizing, editing, and enhancing images of the target individual with high fidelity to their facial features. However, MyStyle does not demonstrate precise control over the attributes of the generated images. We propose to address this problem through a novel optimization system that

organizes the latent space in addition to tuning the generator. Our key contribution is to formulate a loss that arranges the latent codes, corresponding to the input images, along a set of specific directions according to their attributes. We demonstrate that our approach, dubbed MyStyle++, is able to synthesize, edit, and enhance images of an individual with great control over the attributes, while preserving the unique facial characteristics of that individual.

## CCS CONCEPTS

• Computing methodologies → Image manipulation.

## KEYWORDS

Personalization, generative adversarial networks, controllability




This work is licensed under a Creative Commons Attribution International 4.0 License.

SA Conference Papers ’23, December 12–15, 2023, Sydney, NSW, Australia  
© 2023 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0315-7/23/12.  
<https://doi.org/10.1145/3610548.3618171>

## ACM Reference Format:

Libing Zeng, Lele Chen, Yi Xu, and Nima Khademi Kalantari. 2023. MyStyle++: A Controllable Personalized Generative Prior. In *SIGGRAPH Asia 2023 Conference Papers (SA Conference Papers ’23)*, December 12–15, 2023, Sydney,



**Figure 2: On the top, we show the editing results of MyStyle using expression direction from InterFaceGAN [Shen et al. 2020]. Since the original direction does not reside within the personalized subspace, editing with this direction produces results with altered identity (rightmost image). By performing the edit using the projected direction, the identity is better preserved, but the expression becomes entangled with the yaw angle. Our method preserves the identity and keeps the other attributes intact while removing the smile.**

NSW, Australia. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3610548.3618171>

## 1 INTRODUCTION

Ever since the introduction of generative adversarial networks (GAN) [Goodfellow et al. 2014], there has been a growing interest in unconditional image synthesis, which has led to a rapid improvement in resolution and quality of the images generated by GAN-based approaches. In particular, StyleGAN [Karras et al. 2021, 2019, 2020], one of the most popular image generators, produces high-resolution results that are indistinguishable from real images. Built on the success of StyleGAN, a large number of methods [Abdal et al. 2021; Gal et al. 2022; Härkönen et al. 2020; Patashnik et al. 2021; Shoshan et al. 2021; Tov et al. 2021; Wang et al. 2021a, 2022; Wu et al. 2021] use it as a prior for semantic face editing and other image enhancement tasks, such as inpainting and super-resolution. However, the major problem with these approaches is that they use a general prior, trained on a large number of diverse identities. Therefore, their edited or enhanced images may not preserve the identity and key facial features of the target person.

The recent approach by Nitzan et al. [2022], coined MyStyle, addresses this issue by personalizing the generative prior for an individual of interest. Specifically, given a few images of a person, MyStyle first projects these images into the latent space of a pre-trained StyleGAN to obtain a set of latent vectors, called anchors. It then tunes the generator by minimizing the error between the synthesized anchor images and their corresponding input images. Through this process, the generator becomes highly tuned to reconstruct the individual of interest with high fidelity in the specific regions in the latent space, covered by the anchors. MyStyle produces impressive results, preserving the identity and facial features of the target individual, for various tasks such as synthesis, semantic editing, and image enhancement.

However, this technique does not demonstrate precise control over the attributes of the generated images. For example, to synthesize an image with a particular set of attributes, one should randomly sample the convex hull of the anchor points until a desired image is reached by chance. For image editing, MyStyle uses the editing directions provided by approaches, such as InterFaceGAN [Shen et al. 2020], to offer controllability over the attributes of the generated images. Since these editing directions are learned over the entire domain, they may not reside within the personalized subspace. As shown in Fig. 2 (top), by performing the edits using the original direction, the latent codes will quickly fall outside the personalized subspace, producing images with a different identity. To address this issue, MyStyle personalizes the editing direction by projecting it into the subspace. While the projected edit direction keeps the latent codes within the personalized subspace, it loses the ability to perform disentangled edits. As shown in Fig. 2 (middle), removing the expression also results in changing the yaw angle.

Our goal is to address these problems by providing full control over a set of pre-defined attributes of the generated images. To this end, we make a key observation that anchors corresponding to a single person are usually clustered together in a small region within the latent space. Therefore, we can organize the latent space within that region by rearranging the anchors. Since it is easier for a generator, like StyleGAN, to preserve the smoothness of the output variation over the space of the latent space, rearranging the anchors causes the space in between to be dragged with them, resulting in an organized latent space.


Armed with this observation, we propose a novel optimization system to personalize a generative prior by both tuning the generator and organizing the latent space through optimizing the anchors. Our key contribution is to formulate a loss function that arranges the anchors with specific attributes along a particular direction in the latent space. Specifically, we project the anchors into a set of principal axes and minimize the variance of the projection for all the anchors with the same attribute. By doing so, the generator becomes highly tuned to one individual, while the attributes can be controlled within a small hypercube in the latent space.

We demonstrate that our proposed method, called MyStyle++, allows synthesizing images with high fidelity to the characteristics of one individual, while providing full control over a set of pre-defined attributes. We also show that our method can better disentangle different attributes compared to MyStyle [Nitzan et al. 2022]. Moreover, we demonstrate that our system can produce images with a desired attribute during image enhancement.

## 2 RELATED WORK

### 2.1 Deep Generative Networks

Generative Adversarial Networks (GANs) consist of two main modules: a generator and a discriminator [Goodfellow et al. 2014]. The generator takes a noise vector as input and tries to capture the distribution of true examples. The generator focuses on producing an output that fools the discriminator, whose purpose is to classify whether the output is real or fake. GANs have been used extensively to synthesize images that are in line with the training data distribution [Brock et al. 2019; Karras et al. 2018, 2019; Zhu et al. 2017]. Among different variants, StyleGAN [Karras et al. 2021, 2019, 2020],



**Figure 3: Illustration of our data organization with two attributes  $M = 2$ , yaw and expression. We quantize the range of continuous attributes to obtain a set of discrete levels  $a_{m,p}$  across all attributes. The estimated attributes for each image are then assigned to their nearest discrete level.**

which is a carefully re-designed generator architecture, produces the results that are indistinguishable from real photographs, particularly for human faces. In our work, we use StyleGAN2 [Karras et al. 2020] as the base network and personalize it by tuning the generator and organizing the latent space.

## 2.2 Controllable GANs

StyleGAN generates photorealistic portrait images of faces, but it lacks control over semantic face parameters, such as face pose, expressions, and scene illumination. Recently, many StyleGAN variants [Abdal et al. 2021; Gal et al. 2022; Härkönen et al. 2020; Patashnik et al. 2021; Shoshan et al. 2021; Tov et al. 2021; Wang et al. 2021a, 2022; Wu et al. 2021] have been introduced to address this problem. For example, StyleFlow [Abdal et al. 2021] proposes flow models for non-linear exploration of a StyleGAN latent space. GANSpace [Härkönen et al. 2020] attempts to analyze the GAN space by identifying latent directions based on principal component analysis (PCA), applied either in latent space or feature space.

Most controllable portrait image generation methods [BR et al. 2021; Ji et al. 2022; Sun et al. 2022; Tewari et al. 2020a,b; Wang et al. 2021b; Zhou et al. 2019] either rely on 3D morphable face models (3DMMs) [Blanz and Vetter 1999] to achieve rig-like control over StyleGAN, or utilize another modality as guidance (e.g., facial landmark and audio) to control the generation. For instance, by building a bijective mapping between the StyleGAN latent code and the 3DMM parameter sets, StyleRig [Tewari et al. 2020b] achieves the controllable parametric nature of existing morphable face models and the high photorealism of generative face models. Ji et al. [2022] propose an approach to generate one-shot emotional talking faces controlled by an emotion source video and an audio clip.

To explicitly control the camera, several algorithms propose generative neural radiance fields [Chan et al. 2021; Deng et al. 2022] to produce 3D images. Others [Bergman et al. 2022; Simsar et al. 2022] extend these ideas by providing the ability to control other attributes of 3D GANs.

Unfortunately, all the approaches discussed in this section either struggle to retain crucial facial features (identity) after editing [Nitzan et al. 2022] or are unable to maintain explicit control over fully disentangled attributes.

## 2.3 Few-shot GANs and Personalization

Drawing inspiration from the human capability of picking up the essence of a novel object from a small number of examples and generalizing from there, many works [Chen et al. 2021; Liu et al. 2019; Nitzan et al. 2022; Ojha et al. 2021; Wang et al. 2019; Zakharov et al. 2019] seek to further improve the generation quality by adapting the pre-trained model to few-shot image samples. Zakharov et al. [2019] propose a framework that performs lengthy meta-learning on a large dataset of videos. After this training, this method is able to frame few- and one-shot learning of neural talking head models of previously unseen people as adversarial training problems with high capacity generators and discriminators. The appearance information of the unseen target person is learned by the adaptive instance normalization layers. More recently, MyStyle [Nitzan et al. 2022] tunes the weights of a pre-trained StyleGAN face generator to form a local, low-dimensional, personalized manifold in the latent space within a small reference set of portrait images of the target person. The synthesized images within the adapted personalized latent space have better identity-preserving ability compared with the original StyleGAN. However, MyStyle does not demonstrate precise control of the attributes of the generated images. We focus on addressing this issue by organizing the personalized subspace according to a set of pre-defined attributes.

## 3 ALGORITHM

Given a few images of an individual with a set of corresponding attributes, our goal is to obtain a personalized generative prior that allows us to synthesize images of that individual with high fidelity and *full control over the desired attributes*. Specifically, we use the pre-trained StyleGAN [Karras et al. 2019, 2020] face generator and adapt it to the target individual through a novel optimization system. During tuning, we organize the latent space by optimizing the anchors according to the attributes to be able to easily sample an image with a desired set of attributes. Additionally, we optimize the generator to ensure it can produce images that are faithful to the characteristics of the target individual. Below we discuss our approach in detail by first explaining our data pre-processing.

### 3.1 Data Pre-processing

Given a set of  $N$  images of an individual, we first follow the pre-processing steps of MyStyle [Nitzan et al. 2022] to align, crop, and resize the images. We then estimate a set of  $M$  pre-defined attributes (e.g., yaw and expression) for each image. Certain attributes have a discrete domain, while others are continuous. We leave the discrete attributes unchanged, but quantize the range of the continuous ones to obtain  $a_{m,p}$ , where  $m$  refers to the attribute type  $m \in \{1, \dots, M\}$ , while  $p \in \{1, \dots, P(m)\}$  is the index of the attribute value. Note that the number of quantization levels  $P(m)$  could be different for each attribute  $m$ . The estimated attributes for each image are then snapped to the nearest quantized values. A simple example



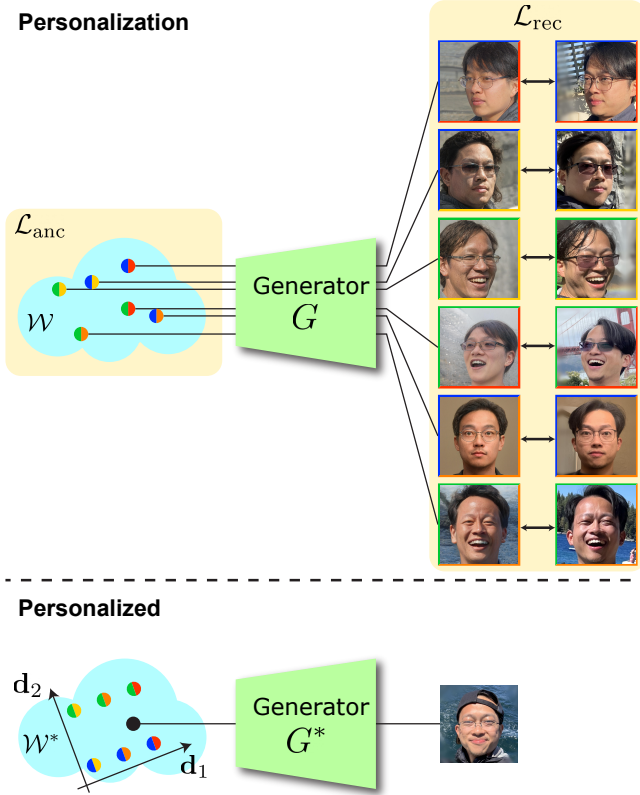


Figure 4: On the top, we show an overview of our controllable personalization approach. Given a set of input images of an individual with their corresponding attributes, we first encode them into  $\mathcal{W}$  space of StyleGAN to obtain a set of anchors, shown with circles. Note that the colors indicate the attributes of the images. In this case, different views are indicated with yellow, orange, and red, while the colors for different expressions are green and blue. We then minimize an objective function consisting of anchor  $\mathcal{L}_{anc}$  and reconstruction  $\mathcal{L}_{rec}$  losses to organize the latent space, by updating the anchors, while tuning the generator. After optimization, we obtain an organized latent space  $\mathcal{W}^*$  that can be easily sampled according to a set of attributes, and a tuned generator that can produce images that are faithful to the facial characteristics of the target individual.

illustrating this process is shown in Fig. 3. We provide more details on the attributes and our quantization strategy in Sec. 4.

### 3.2 Controllable Personalization

We begin by projecting the input images into the latent space of StyleGAN, using the pre-trained encoder by Richardson et al. [2021], to obtain a set of  $N$  latent codes  $\{\mathbf{w}_n\}_{n=1}^N$ . We follow MyStyle [Nitzan et al. 2022] terminology and call these latent codes, anchors. As discussed, in addition to tuning the generator to improve its fidelity to the target individual, we would like to organize the latent space to have full control over a set of attributes. An overview of our approach is shown in Fig. 4.

Our key observation is that we can organize the latent space by only rearranging the anchors. This is because the output of

StyleGAN changes smoothly with respect to the input, and thus as an anchor moves, its neighborhood will be dragged with it. Based on this observation, we formulate an anchor loss to rearrange the anchors based on their attributes.

Before explaining our anchor loss in detail, we discuss the properties of an ideal latent space: **1)** Each attribute should change along a known direction;  $\mathbf{d}_m$  for the  $m^{\text{th}}$  attribute. This is to ensure we can perform semantic editing and change a particular attribute by simply modifying a latent code along that attribute’s direction. **2)** All the latent codes that project to the same value along an attribute direction should have the same attribute. For example, all the latent codes that project to 0.5 along the yaw direction should correspond to images of front faces. This allows us to directly sample an image with a certain set of attributes by ensuring that the latent code projects to appropriate values along each attribute direction. **3)** The directions for different attributes should be orthogonal to guarantee that the attributes are fully disentangled and changing one will not result in modifying the other attributes.

We propose to codify the three properties into the following anchor loss:

$$\mathcal{L}_{anc} = \sum_{m=1}^M \mathcal{L}_d(\mathbf{d}_m), \text{ where } \mathcal{L}_d = \sum_{n=1}^N \|\mathbf{w}_n \cdot \mathbf{d}_m - c_{n,m}\|. \quad (1)$$

Here,  $\mathbf{w}_n \cdot \mathbf{d}_m$  computes the projection of the anchor for the  $n^{\text{th}}$  image onto the direction of  $m^{\text{th}}$  attribute through dot product. Moreover,  $c_{n,m}$  is the average of the projected anchors into direction  $\mathbf{d}_m$  for all the images with the same  $m^{\text{th}}$  attribute as the  $n^{\text{th}}$  image (subset denoted as  $\mathcal{N}_{n,m}$ ). Formally, we can write this as follows:

$$c_{n,m} = \frac{1}{|\mathcal{N}_{n,m}|} \sum_{k \in \mathcal{N}_{n,m}} \mathbf{w}_k \cdot \mathbf{d}_m, \quad (2)$$

where

$$\mathcal{N}_{n,m} = \{k \in \{1, \dots, N\} \mid k \neq n, f_a(\mathbf{I}_n)[m] = f_a(\mathbf{I}_k)[m]\}. \quad (3)$$

Here,  $f_a(\mathbf{I}_n)[m]$  returns the quantized  $m^{\text{th}}$  attribute of image  $\mathbf{I}_n$ . We note that  $c_{n,m}$  changes at every iteration of the optimization. By minimizing the loss in Eq. 1, we ensure that all the anchors with the same  $m^{\text{th}}$  attribute, project to the same point along  $m^{\text{th}}$  attribute direction  $\mathbf{d}_m$ , satisfying our second desired property. This loss also ensures that each attribute is changed along its specific direction, satisfying the first property. This can be seen visually in Fig. 3; for example, if all the images with a specific yaw (each column) project to the same point in the yaw direction, moving along this direction will change the yaw.

To satisfy the third property, we apply principle component analysis (PCA) to all the  $N$  anchors and use a subset of the principle components as our  $\mathbf{d}_m$ . We assign a specific principal component to each  $\mathbf{d}_m$  through the following objective:

$$\mathbf{d}_m = \arg \min_{\mathbf{v}_i \in \mathbf{V}} \mathcal{L}_d(\mathbf{v}_i) \quad (4)$$

where  $\mathbf{V}$  is the set of all the principle components and  $\mathcal{L}_d$  is defined in Eq. 1. The intuition behind this is that we would like to perform the least amount of rearrangement by ensuring that the latent space is already well aligned with respect to the selected directions. Note that we perform PCA at every iteration of training. Therefore, as we rearrange the anchor points in different iterations, the directions will be updated as well. We also note that although the objective in



Eq. 4 could potentially assign different principle components to a particular attribute direction  $\mathbf{d}_m$  in different iterations, we did not observe this phenomenon in our experiments.

To perform personalization, we minimize the combination of the anchor and reconstruction losses

$$\mathcal{L} = \mathcal{L}_{\text{anc}} + \mathcal{L}_{\text{rec}}, \quad (5)$$

where the reconstruction loss  $\mathcal{L}_{\text{rec}}$  minimizes the error between the synthesized  $G(\mathbf{w}_n)$  and the corresponding input images  $\mathbf{I}_n$ . We follow MyStyle and use a combination of LPIPS [Zhang et al. 2018] and L2 as our reconstruction loss. During optimization, both the latent codes corresponding to the anchors and the weights of the generator are updated. Note that in addition to adapting the generator to the input image set, the reconstruction loss plays a critical role in avoiding trivial solutions to the anchor loss, e.g., collapsing all the anchors to a single point.

Once the optimization is performed, we obtain an organized latent space  $\mathcal{W}^*$  and tuned generator  $G^*$ . All the attributes can be controlled within an  $M$ -dimensional hypercube in the organized latent space. The bounds of this hypercube can simply be found by projecting all the anchors into each axis of the hypercube  $\mathbf{d}_m$  and computing the minimum and maximum values. Note that all the other attributes, not being used during optimization, are encoded in the remaining PCA dimensions.

### 3.3 Controllable Synthesis, Edit, and Enhancement

We now describe how to use our personalized generative prior for various tasks.

*Synthesis:* Controlling the synthesized images can easily be done by ensuring that the sampled latent code projects to the desired location in the  $M$ -dimensional hypercube. However, special care must be taken to ensure the latent code does not fall outside of the personalized space. Following MyStyle, we define the convex hull of all the organized anchors  $\mathbf{w}_n^*$  as the personalized subspace within  $\mathcal{W}^*$ . This convex hull is represented through generalized barycentric coordinate as the weighted sum of the anchors, where the weights (coordinates)  $\alpha = \{\alpha_n\}_{n=1}^N$  sum up to 1 and are greater than  $-\beta$  ( $\beta$  is a positive value). The latter condition dilates the space by a small amount to ensure expressiveness.

We propose a simple strategy to perform controlled sampling in the dilated convex hull. Specifically, we first randomly sample  $\alpha$  to ensure the latent code is within the personalized subspace. We then project the sampled latent code into PCA and set the projected values along the attribute directions  $\mathbf{d}_m$  to the desired values. Note that, while it is possible for the modified latent codes to fall outside the dilated convex hull and require reprojection to the personalized space, we did not observe such cases in practice. This is mainly because our latent space is organized according to the attributes and our modifications are performed inside a hypercube which is part of the subspace.

*Semantic Editing:* Since our latent space is organized, the editing process for sampled images is straightforward. To edit an image, we project its latent code to PCA and perform the edit by changing

the coordinate in the hypercube. To edit a real image  $\mathbf{I}$ , we first project the image into the  $\alpha$  space through the following objective:

$$\alpha^* = \arg \min_{\alpha} \mathcal{L}_{\text{rec}}(G(\mathbf{W}^* \alpha), \mathbf{I}), \quad (6)$$

where  $\mathbf{W}^*$  is a matrix with organized anchors along its columns. Note that we follow MyStyle’s approach to ensure  $\alpha$  values satisfy the conditions of the dilated convex hull, i.e., they sum up to 1 and are greater than  $-\beta$ . Once we obtain the optimized latent code, following Roich et al. [2022], we further tune the generator to better match the input image. We then perform the semantic edits, by changing the latent code in PCA.

*Image Enhancement:* Given an input image  $I$  with a known degradation function  $Q$ , our goal is to enhance the image, while controlling the attributes of the reconstructed image. We propose to do this through the following objective:

$$\alpha^* = \arg \min_{\alpha} \mathcal{L}_{\text{rec}}(Q(G(\mathbf{W}^* \alpha)), \mathbf{I}) + \lambda \sum_{m=1}^M \|(\mathbf{W}^* \alpha) \cdot \mathbf{d}_m - a_m\|, \quad (7)$$

where  $\lambda$  controls the balance between the two terms and we set it to one in our implementation. Here, the first term ensures that the generated image, after applying the degradation function, is similar to the input image. The second term encourages the projection of the latent code  $\mathbf{W}^* \alpha$  onto the  $m^{\text{th}}$  attribute direction to be similar to the desired value  $a_m$ . Note that, we can perform enhancement by controlling a subset of the attributes, by only applying the second term to the attributes of interest. Similarly, for uncontrolled enhancement, we simply remove the second term.

## 4 RESULTS

We implement the proposed approach in PyTorch and adopt ADAM optimizer [Kingma and Ba 2015] with the default parameters. All the results are obtained after tuning a pre-trained StyleGAN2 [Karras et al. 2020] generator on FFHQ [Karras et al. 2019] dataset. We perform the tuning for 3000 epochs with a batch size of one and a learning rate of  $5e-3$  across all datasets.

We have tested our system on the following individuals: Barack Obama (93 images), Emma Watson (304 images), Joe Biden (142 images), Leonardo DiCaprio (217 images), Michelle Obama (138 images), Oprah Winfrey (106 images), Scarlett Johansson (179 images), and Taylor Swift (129 images). We consider the expression, as well as yaw and pitch angles as the attributes for all individuals. For Leonardo DiCaprio and Emma Watson, we include age in addition to the other three attributes. Throughout this section, we show our results on some of these individuals, but more results can be found in the supplementary materials.

We estimate the expression, yaw, and pitch by leveraging AWS Rekognition API [Amazon 2023], while we employ the DEX VGG network [Rothe et al. 2018] to estimate the age attribute. We quantize yaw and pitch angles by every 5 degrees and age by every 2 years during the data pre-processing stage, described in Sec. 3.1. For expression, we utilize a combination of the “Smile” and “MouthOpen” attributes of the AWS output, which indicates the presence of the attribute as true or false with a confidence level ranging from 50 to 100. We divide the confidence level by 20% and round it down to the nearest integer, resulting in three groups of presence and three groups of absence for each attribute. We then combine the lowest

groups of presence and absence (presence and absence with 50% to 60% confidence) into the same group, resulting in five quantization levels for both “Smile” and “MouthOpen”. The images with the same “Smile” and “MouthOpen” quantization levels are then grouped together.

We compare our approach against two versions of MyStyle, called MyStyle\_I and MyStyle\_P, where the editing directions are obtained from InterFaceGAN [Shen et al. 2020] and PCA (using Eq. 4), respectively. Note that in MyStyle\_P we do not organize the latent space and only tune the generator, i.e., minimize the reconstruction loss, but not the anchor loss. Although MyStyle does not demonstrate controllable synthesis, we use the approach discussed in Sec. 3.3 with the directions from InterFaceGAN and PCA to imbue MyStyle with this capability.

Here, we show a subset of our results, but more comparisons and evaluations can be found in our accompanying video and supplementary materials.

#### 4.1 Synthesis

We begin by comparing our controllable synthesis results for Oprah Winfrey, Barack Obama, Scarlett Johansson, and Leonardo DiCaprio against MyStyle\_I and MyStyle\_P. For each person, we show a set of results by fixing one attribute and randomly sampling the rest. As shown in Fig. 5, both MyStyle\_I and MyStyle\_P produce results with large variations in the attribute of interest, because the directions from InterFaceGAN [Shen et al. 2020] and PCA do not match the correct attribute directions in the personalized subspace. For example, on the top, a large smile is expected, whereas images generated by MyStyle\_I and MyStyle\_P exhibit a range of different expressions. While yaw is usually the dominant attribute in the latent space and relatively easy to control, MyStyle\_I and MyStyle\_P exhibit undesirable yaw variance for Barack Obama. Similarly, these baselines produce results with large pitch and age variations for Scarlett Johansson and Leonardo DiCaprio, respectively. In contrast, our approach produces results that are consistent in all four cases. Note that InterFaceGAN does not provide a direction corresponding to the pitch, and thus we only compare against MyStyle\_P for the case with fixed pitch.

We further numerically evaluate the ability of our method to control the attributes in comparison with MyStyle\_P and MyStyle\_I in Table 1. To accomplish this, we generate 100 images by fixing one attribute and randomly sampling the other ones. We then estimate the attributes of the generated images, using AWS Rekognition for expression, as well as the yaw and pitch angles, and DEX VGG [Rothe et al. 2018] for age, and compute the standard deviation of the estimated attribute for all the 100 images. For each attribute, we show the results for five normalized values (0.0, 0.25, 0.5, 0.75, 1.0). As seen, MyStyle\_P and MyStyle\_I generate inferior results as the PCA and InterFaceGAN attribute directions are not well-aligned with the correct attribute directions in the subspace. In contrast, our approach consistently demonstrates the smallest standard deviation across all attributes for both Scarlett Johansson and Leonardo DiCaprio.

A potential concern is whether our latent space organization could compromise the diversity and preservation of the identity of the results. To numerically evaluate this, we compute the ID metric,

**Table 1: We numerically compare our controlled synthesis results against MyStyle\_P and MyStyle\_I. We generate 100 images for each fixed attribute value and report the standard deviation of the estimated attribute of interest over the generated images. Note that the attribute values (e.g., 0.25) are in the normalized coordinate  $d_m$ . The best results are shown in bold. Note that we do not report any fixed pitch synthesis results for MyStyle\_I as InterFaceGAN [Shen et al. 2020] does not provide an edit direction for Pitch.**

		Scarlett Johansson				
		0.0	0.25	0.5	0.75	1.0
Exp	MyStyle_P	0.651	0.881	0.926	0.878	0.667
	MyStyle_I	0.242	0.771	0.788	0.391	0.001
	Ours	<b>0.025</b>	<b>0.182</b>	<b>0.470</b>	<b>0.230</b>	<b>0.000</b>
Yaw	MyStyle_P	7.815	7.331	4.418	5.563	6.568
	MyStyle_I	3.894	3.553	2.374	3.005	3.579
	Ours	<b>2.677</b>	<b>2.335</b>	<b>1.963</b>	<b>1.567</b>	<b>2.390</b>
Pitch	MyStyle_P	4.945	3.912	4.808	4.803	5.162
	MyStyle_I	-	-	-	-	-
	Ours	<b>3.717</b>	<b>3.030</b>	<b>2.501</b>	<b>2.710</b>	<b>2.670</b>
		Leonardo DiCaprio				
		0.0	0.25	0.5	0.75	1.0
Exp	MyStyle_P	0.217	0.242	0.092	0.174	0.123
	MyStyle_I	0.003	0.057	0.070	0.303	0.348
	Ours	<b>0.002</b>	<b>0.006</b>	<b>0.018</b>	<b>0.094</b>	<b>0.003</b>
Yaw	MyStyle_P	6.681	4.088	7.537	7.861	8.346
	MyStyle_I	4.067	2.648	3.263	2.943	3.106
	Ours	<b>2.879</b>	<b>1.462</b>	<b>1.554</b>	<b>1.640</b>	<b>2.444</b>
Pitch	MyStyle_P	3.258	1.680	3.006	5.029	4.596
	MyStyle_I	-	-	-	-	-
	Ours	<b>1.578</b>	<b>1.008</b>	<b>1.882</b>	<b>3.534</b>	<b>2.787</b>
Age	MyStyle_P	6.380	5.767	4.149	4.800	4.134
	MyStyle_I	3.234	3.682	4.166	4.933	3.725
	Ours	<b>1.918</b>	<b>2.577</b>	<b>2.096</b>	<b>2.478</b>	<b>1.906</b>

as proposed in MyStyle [Nitzan et al. 2022], on the results generated by both our approach and MyStyle for Scarlett Johansson and Leonardo DiCaprio. This metric measures the cosine similarity of the features extracted by a deep classifier between the generated image and the closest one from the training data. Besides measuring the ability to preserve the identity, we also compute the diversity of the synthesized images. We follow the protocol suggested by Ojha et al. [2021] to compute the intra-cluster diversity using the LPIPS score. Specifically, we generate 1000 images and assign them to one of the 10 training images, by using the lowest LPIPS distance. Then we compute the average pair-wise LPIPS distance within members of the same cluster and then average over the 10 clusters. As shown in Table 2, our method generates results that are comparable to MyStyle in terms of ID metric and diversity score, demonstrating that our latent space organization does not compromise the diversity and identity preservation of the results.

#### 4.2 Semantic Editing

We compare our semantic editing results against MyStyle\_P and MyStyle\_I in Fig. 6. Specifically, we modify the expression, yaw,

**Table 2: We compare our results against MyStyle in terms of the ID metric [Nitzan et al. 2022] and diversity score [Ojha et al. 2021]. Higher numbers are better. Our method produces similar results compared to MyStyle, which demonstrates that latent organization does not hurt the quality of our results. The best results are shown in bold.**

		Scarlett Johansson	Leonardo DiCaprio
ID $\uparrow$	MyStyle	0.760 $\pm$ 0.007	0.7828 $\pm$ 0.050
	Ours	<b>0.763<math>\pm</math>0.004</b>	<b>0.7856<math>\pm</math>0.070</b>
Diversity $\uparrow$	MyStyle	0.455 $\pm$ 0.081	0.417 $\pm$ 0.049
	Ours	<b>0.471<math>\pm</math>0.053</b>	<b>0.456<math>\pm</math>0.035</b>

**Table 3: We compare our editing results against MyStyle\_I and MyStyle\_P in terms of the mean standard deviation of the edited attribute to show editing consistency (marked with \*), and of fixed attributes to demonstrate attribute disentanglement. We additionally report the ID metric to evaluate identity preservation ability. Best results are shown in bold.**

		Leonardo DiCaprio				
		Exp*	Yaw	Pitch	Age	ID $\uparrow$
Exp	MyStyle_P	0.794	4.866	1.965	3.584	0.743 $\pm$ 0.103
	MyStyle_I	0.527	1.538	1.227	3.179	0.731 $\pm$ 0.108
	Ours	<b>0.268</b>	<b>1.204</b>	<b>0.999</b>	<b>2.201</b>	<b>0.752<math>\pm</math>0.107</b>
		Yaw*	Exp	Pitch	Age	ID $\uparrow$
Yaw	MyStyle_P	4.069	0.213	2.570	2.893	0.717 $\pm$ 0.108
	MyStyle_I	3.925	0.111	2.426	2.700	0.716 $\pm$ 0.117
	Ours	<b>2.097</b>	<b>0.075</b>	<b>2.108</b>	<b>2.212</b>	<b>0.728<math>\pm</math>0.115</b>
		Pitch*	Exp	Yaw	Age	ID $\uparrow$
Pitch	MyStyle_P	5.463	0.281	3.030	3.720	0.717 $\pm$ 0.121
	MyStyle_I	-	-	-	-	-
	Ours	<b>3.591</b>	<b>0.071</b>	<b>1.786</b>	<b>3.023</b>	<b>0.726<math>\pm</math>0.114</b>
		Age*	Exp	Yaw	Pitch	ID $\uparrow$
Age	MyStyle_P	5.113	0.230	2.808	2.824	0.734 $\pm$ 0.118
	MyStyle_I	7.152	0.134	1.294	2.095	0.723 $\pm$ 0.120
	Ours	<b>3.473</b>	<b>0.087</b>	<b>0.467</b>	<b>1.217</b>	<b>0.739<math>\pm</math>0.113</b>

pitch, and age of Scarlett Johansson, Michelle Obama, Joe Biden, and Leonardo DiCaprio, respectively. MyStyle\_P has difficulties editing Scarlett Johansson’s expression and predominantly changes the yaw. While MyStyle\_I is better able to edit the expression, it slightly changes the yaw (see the supplementary video) and produces a neutral face with altered identity (the leftmost image). Moreover, both MyStyle\_P and MyStyle\_I change the expression when editing Michelle Obama’s yaw angle. For Joe Biden, MyStyle\_P struggles to properly edit the pitch angle as the PCA direction is not well-aligned with the pitch attribute direction in the subspace. Finally, when editing the age of Leonardo DiCaprio, both MyStyle\_P and MyStyle\_I exhibit noticeable changes to the expression and pitch, respectively. Additionally, both approaches struggle to preserve the identity of the edited images in extreme cases (rightmost for MyStyle\_P and leftmost for MyStyle\_I). In contrast to these techniques, our method only changes the attribute of interest when producing edited results and is able to better preserve the identity. Again, we note that we do not show pitch editing for MyStyle\_I as

InterFaceGAN does not provide a direction corresponding to the pitch attribute.

Next, we compare our method against the other techniques for editing real images of Barack Obama, Michelle Obama, Emma Watson, and Leonardo DiCaprio, in Fig. 7. MyStyle\_P changes the yaw and MyStyle\_I has difficulties preserving the identity of Barack Obama when adding a smile. For Michelle Obama, both MyStyle\_P and MyStyle\_I change the expression when editing the yaw angle. For Emma Watson, MyStyle\_P is unable to edit the pitch. Finally, MyStyle\_P changes the yaw angle when editing Leonardo DiCaprio’s age, while MyStyle\_I has difficulties maintaining the identity. In contrast to these methods, our approach disentangles the attributes more effectively and is better at preserving the identities in all four cases.

We note that the reason behind MyStyle’s occasional failure to preserve the identity is that the edited latent codes, in some cases, fall outside the personalized subspace. While the loss of identity can be resolved by projecting the edited latent codes back to the convex hull, using MyStyle’s suggested strategy, this process produces results with undesirable attributes. This is shown in Fig. 8 where the objective is to add a smile to Barack Obama and produce a teenage Leonardo DiCaprio. MyStyle\_I produces results with altered identities as evident both visually and numerically through the ID metric. The identity is improved by projecting the edited latent codes to the subspace (third column), but this process reduces the smile (top) and increases the age (bottom).

We further numerically compare our real image editing results against MyStyle\_P and MyStyle\_I on Leonardo DiCaprio in Table 3. Specifically, we evaluate the editing consistency by computing the mean standard deviation of the edited attribute, while we measure the attribute disentanglement by calculating the mean standard deviation of the unedited attributes. The standard deviations are computed over 21 edits and they are averaged over 21 images. We additionally evaluate the ability of different methods to preserve the identity using the ID metric. As seen, our method consistently outperforms MyStyle\_P and MyStyle\_I across all metrics.

### 4.3 Image Enhancement

As discussed in Sec. 3.3, since our method provides precise control over the attributes, it can be used to perform controllable image enhancement. This is shown in Figs. 9 and 10 for image inpainting and super-resolution, respectively. As seen our method can produce inpainted and super-resolved images with the desired expressions.

### 4.4 Analysis

We begin by evaluating the effect of number of images on the quality Yaw-editing for Scarlett Johansson in Table 3. As seen, our system works reasonably well with 92 images, but the editing quality substantially decreases when using only 46 images. Furthermore, we show the information encoded in an uncontrolled PCA component by editing the sampled images along the uncontrolled axis. As shown in Fig. 11, this particular PCA component slightly darkens the color of Scarlett Johansson’s hair (top), while changing the same component raises Michelle Obama’s eyebrows (bottom). It is, however, expected for the same component to correspond to different attributes for different individuals as the latent space is



**Table 4: Evaluating the effect of the number images on the quality of Yaw-editing for Scarlett Johansson in a manner akin to Table 3.**

	Yaw*	Exp	Pitch	ID ↑
185 Images	2.349	0.774	1.669	0.707±0.069
92 Images	2.643	0.791	1.921	0.681±0.085
46 Images	4.840	0.951	2.989	0.624±0.093

personalized for each person through our optimization. Finally, we demonstrate our method’s extrapolation ability by showing the editing results beyond the convex hull in Fig. 12. As shown, our approach can produce reasonable extrapolation results.

## 5 LIMITATIONS AND FUTURE WORK

Our approach is able to produce high-quality results with great control over a set of attributes. However, it has a few limitations. First, the number of images required for personalization increases significantly with the number of desired attributes. This is because we rely on the propagation of the anchors to the neighboring regions. If the anchors in certain regions are sparse, those areas are not going to be personalized appropriately. However, this is not unique to our approach and MyStyle suffers from the same drawback. For example, if MyStyle is personalized with images of a young subject, it cannot produce images of the subject at an old age with high fidelity. Second, while our approach provides great control over the attributes, our reconstructions for attributes like view are not physically accurate. In the future, it would be interesting to incorporate the image formation process into our system to improve accuracy.

We note that although our approach has the potential to be applied to cases beyond MyStyle, such as organizing the entire latent space of StyleGAN, one significant challenge arises: organizing the entire latent space necessitates a large number of anchor images, resulting in time-consuming and difficult optimization. Furthermore, special attention must be given to prevent anchors with different identities from being placed closely together after optimization; this is not an issue when handling a single individual.

## 6 CONCLUSION

We have presented an approach to obtain a controllable personalized generative prior from a set of images of an individual. Our system allows for reconstructing images of the individual that faithfully preserve the key facial features of the person, while providing full control over a set of pre-defined attributes. In addition to tuning a pre-trained generator, we organize its latent space such that different attributes change along certain known directions. To do this, we formulate a loss that rearranges the latent codes, corresponding to the input images, according to the attributes. We show that our method better disentangles the attributes than MyStyle, while providing full control over the attributes.

## ACKNOWLEDGMENTS

We express our gratitude to the anonymous reviewers for their insightful comments and suggestions. Additionally, we would like to thank Keqiang Yan and Yongqing Liang for the valuable discussions.

## REFERENCES

- Rameen Abdal, Peihao Zhu, Niloy J Mitra, and Peter Wonka. 2021. Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows. *ACM Transactions on Graphics (TOG)* 40, 3 (2021), 1–21.
- Amazon. 2023. AWS Rekognition. <https://aws.amazon.com/rekognition/>.
- Alexander W. Bergman, Petr Kellnhofer, Wang Yifan, Eric R. Chan, David B. Lindell, and Gordon Wetzstein. 2022. Generative Neural Articulated Radiance Fields. In *NeurIPS*.
- Volker Blanz and Thomas Vetter. 1999. A morphable model for the synthesis of 3D faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*. 187–194.
- Mallikarjun BR, Ayush Tewari, Abdallah Dib, Tim Weyrich, Bernd Bickel, Hans-Peter Seidel, Hanspeter Pfister, Wojciech Matusik, Louis Chevallier, Mohamed Elgharib, et al. 2021. Photoapp: Photorealistic appearance editing of head portraits. *ACM Transactions on Graphics* 40, 4 (2021), 1–16.
- Andrew Brock, Jeff Donahue, and Karen Simonyan. 2019. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=B1xsqj09Fm>
- Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. 2021. Efficient Geometry-aware 3D Generative Adversarial Networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Lele Chen, Chen Cao, Fernando De la Torre, Jason Saragih, Chenliang Xu, and Yaser Sheikh. 2021. High-fidelity face tracking for AR/VR via deep lighting adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 13059–13069.
- Yu Deng, Jiaolong Yang, Jianfeng Xiang, and Xin Tong. 2022. GRAM: Generative Radiance Manifolds for 3D-Aware Image Generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Rinon Gal, Or Patashnik, Haggai Maron, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. 2022. Stylegan-nada: Clip-guided domain adaptation of image generators. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–13.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger (Eds.), Vol. 27. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf>
- Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. 2020. Ganspace: Discovering interpretable gan controls. *Advances in Neural Information Processing Systems* 33 (2020), 9841–9850.
- Xinya Ji, Hang Zhou, Kaisiyuan Wang, Qianyi Wu, Wayne Wu, Feng Xu, and Xun Cao. 2022. EAMM: One-Shot Emotional Talking Face via Audio-Based Emotion-Aware Motion Model. In *ACM SIGGRAPH 2022 Conference Proceedings* (Vancouver, BC, Canada) (SIGGRAPH '22). Association for Computing Machinery, New York, NY, USA, Article 61, 10 pages. <https://doi.org/10.1145/3528233.3530745>
- Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=Hk99ZCeAb>
- Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2021. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems* 34 (2021), 852–863.
- Tero Karras, Samuli Laine, and Timo Aila. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 4401–4410.
- Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 8110–8119.
- Diederik Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*.
- Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, Jaakko Lehtinen, and Jan Kautz. 2019. Few-shot unsupervised image-to-image translation. In *Proceedings of the IEEE/CVF international conference on computer vision*. 10551–10560.
- Yotam Nitzan, Kfir Aberman, Qiurui He, Orly Liba, Michal Yarom, Yossi Gandelsman, Inbar Mosseri, Yael Pritch, and Daniel Cohen-Or. 2022. MyStyle: A Personalized Generative Prior. *arXiv preprint arXiv:2203.17272* (2022).
- Utkarsh Ojha, Yijun Li, Jingwan Lu, Alexei A Efros, Yong Jae Lee, Eli Shechtman, and Richard Zhang. 2021. Few-shot image generation via cross-domain correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10743–10752.
- Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. 2021. Styleclip: Text-driven manipulation of stylegan imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2085–2094.
- Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. 2021. Encoding in Style: a StyleGAN Encoder for Image-to-Image Translation. In *IEEE/CVF Conference on Computer Vision and Pattern*

- Recognition (CVPR)*.
- Daniel Roich, Ron Mokady, Amit H. Bermano, and Daniel Cohen-Or. 2022. Pivotal Tuning for Latent-Based Editing of Real Images. *ACM Trans. Graph.* 42, 1, Article 6 (aug 2022), 13 pages. <https://doi.org/10.1145/3544777>
- Rasmus Rothe, Radu Timofte, and Luc Van Gool. 2018. Deep Expectation of Real and Apparent Age from a Single Image Without Facial Landmarks. *International Journal of Computer Vision* 126 (2018), 144–157.
- Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. 2020. InterFaceGAN: Interpreting the Disentangled Face Representation Learned by GANs. *TPAMI* (2020).
- Alon Shoshan, Nadav Bhonker, Igor Kviatkovsky, and Gerard Medioni. 2021. Gan-control: Explicitly controllable gans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 14083–14093.
- Enis Simsar, Alessio Tonioni, Evin Pinar Örneke, and Federico Tombari. 2022. LatentSwap3D: Semantic Edits on 3D Image GANs. arXiv. <https://doi.org/10.48550/ARXIV.2212.01381>
- Yasheng Sun, Hang Zhou, Kaisiyuan Wang, Qianyi Wu, Zhibin Hong, Jingtuo Liu, Errui Ding, Jingdong Wang, Ziwei Liu, and Koike Hideki. 2022. Masked Lip-Sync Prediction by Audio-Visual Contextual Exploitation in Transformers. In *SIGGRAPH Asia 2022 Conference Papers*. 1–9.
- Ayush Tewari, Mohamed Elgharib, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhöfer, and Christian Theobalt. 2020a. Pie: Portrait image embedding for semantic control. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–14.
- Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhofer, and Christian Theobalt. 2020b. Stylerig: Rigging stylegan for 3d control over portrait images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6142–6151.
- Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. 2021. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–14.
- Sheng-Yu Wang, David Bau, and Jun-Yan Zhu. 2021a. Sketch your own gan. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 14050–14060.
- Tengfei Wang, Yong Zhang, Yanbo Fan, Jue Wang, and Qifeng Chen. 2022. High-fidelity gan inversion for image attribute editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11379–11388.
- Ting-Chun Wang, Ming-Yu Liu, Andrew Tao, Guilin Liu, Jan Kautz, and Bryan Catanzaro. 2019. Few-shot Video-to-Video Synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. 2021b. One-shot free-view neural talking-head synthesis for video conferencing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10039–10049.
- Zongze Wu, Dani Lischinski, and Eli Shechtman. 2021. Stylespace analysis: Disentangled controls for stylegan image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12863–12872.
- Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. 2019. Few-shot adversarial learning of realistic neural talking head models. In *Proceedings of the IEEE/CVF international conference on computer vision*. 9459–9468.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hang Zhou, Yu Liu, Ziwei Liu, Ping Luo, and Xiaogang Wang. 2019. Talking face generation by adversarially disentangled audio-visual representation. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 9299–9306.
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*. 2223–2232.



Figure 5: We present a comparison of our synthesis results with those of MyStyle\_I and MyStyle\_P for four individuals: Oprah Winfrey, Barack Obama, Scarlett Johansson, and Leonardo DiCaprio. Our method consistently produces attribute-controlled image synthesis, whereas MyStyle\_I and MyStyle\_P exhibit significant inaccuracies in attribute controllability.

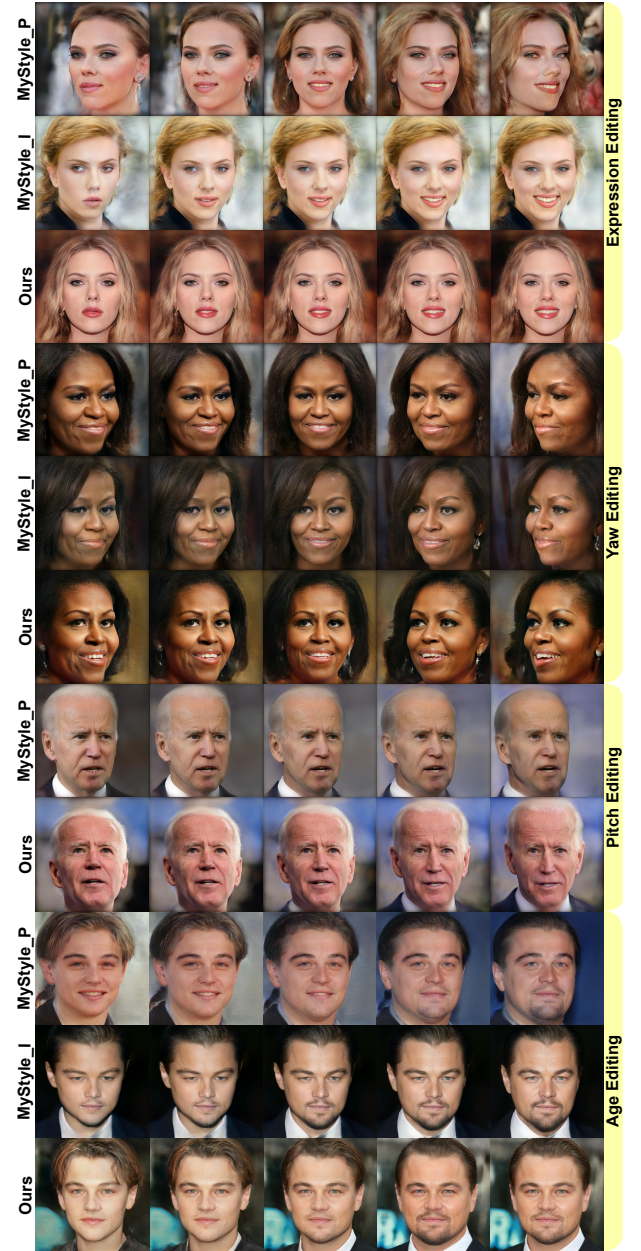


Figure 6: We show our results for editing expression, yaw, pitch, and age for sampled images against MyStyle\_P and MyStyle\_I on Scarlett Johansson, Michelle Obama, Joe Biden, and Leonardo DiCaprio. Our method maintains the identity and consistency of other attributes while making modifications to the expression, yaw, pitch, or age of the image.



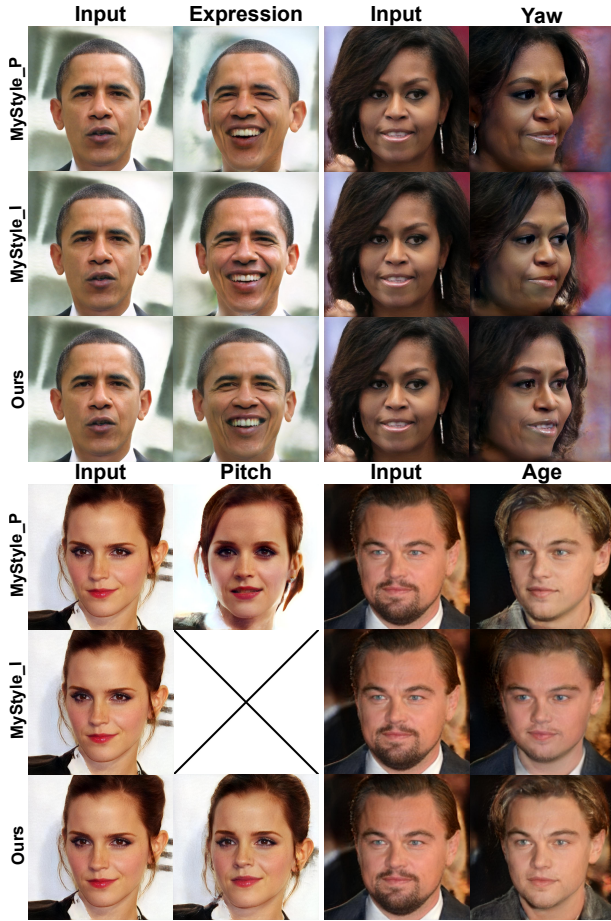


Figure 7: We show comparisons for semantic editing of real images. Input is the projected images into the latent space of MyStyle and our generators. Our method disentangles the attributes better than MyStyle\_P and MyStyle\_I, producing results that preserve the unchanged attributes and identity. Note that InterFaceGAN [Shen et al. 2020] (MyStyle\_I) does not provide an edit direction for pitch. From top-left to bottom-right ©Steve Jurevson, ©Gage Skidmore, ©David Shankbone, ©PictureCapital.



Figure 8: MyStyle\_I produces edited results with altered identity (see identity scores - bottom right). The identity can be improved by projecting the latent code to subspace (third column), but the projected images do not have the desired attributes. Top to bottom ©Steve Jurevson, ©PictureCapital.



Figure 9: We demonstrate our expression controllability on the image inpainting task. For each case, we present two expressions (“Exp 1” and “Exp 2”) in our results. Our approach generates inpainting outputs that smoothly blend with the unmasked regions while adhering to the desired expression. Top to bottom ©Elen Nivrae, ©David Shankbone.



Figure 10: We show our results for the super-resolution task with two different expressions (“Exp 1” and “Exp 2”). Credit: ©Gage Skidmore.



Figure 11: We show the information encoded in an uncontrolled PCA component by performing editing the sampled images of two individuals along this direction. As seen, hair color and eyebrow location are encoded in this particular uncontrolled PCA component for Scarlett Johansson and Michelle Obama, respectively.



Figure 12: We show our extrapolation results by editing Leonardo Dicaprio’s age beyond the convex hull.