

# Unsupervised Disentanglement of Linear-Encoded Facial Semantics

Yutong Zheng, Yu-Kai Huang, Ran Tao, Zhiqiang Shen and Marios Savvides  
Carnegie Mellon University

{yutongzh, yukaih2, rant, zhiqians, marios}@andrew.cmu.edu



Figure 1: Facial manipulation result by sampling along latent representations in StyleGAN, learned by our unsupervised framework. **First column:** Original generations. **Column 2-4:** manipulation on univariate environmental semantics. **Column 5-9:** manipulation on localized facial semantics.

## Abstract

*We propose a method to disentangle linear-encoded facial semantics from StyleGAN without external supervision. The method derives from linear regression and sparse representation learning concepts to make the disentangled latent representations easily interpreted as well. We start by coupling StyleGAN with a stabilized 3D deformable facial reconstruction method to decompose single-view GAN generations into multiple semantics. Latent representations are then extracted to capture interpretable facial semantics. In this work, we make it possible to get rid of labels for disentangling meaningful facial semantics. Also, we demonstrate that the guided extrapolation along the disentangled representations can help with data augmentation, which sheds light on handling unbalanced data. Finally, we provide an analysis of our learned localized facial representations and illustrate that the semantic information is encoded, which surprisingly complies with human intuition. The overall unsupervised design brings more flexibility to representation learning in the wild.*

## 1. Introduction

In recent years, Generative Adversarial Networks (GANs) [12] have been a great success in synthesizing

photo-realistic images given a set of latent codes. Despite the rapid boost in image quality, the interpretability of the generation process has become another major area of research. In general, interpretability requires latent codes to encode disentangled semantic information of the image. Further, ideally, well-disentangled semantics are supposed to be factorized to practically interpretable components and each component should be linear-encoded in the latent space as representation [15, 9, 5, 19, 30, 11].

StyleGAN [17] proposes a new architecture by bringing an intermediate latent space to provide support for disentanglement property for face generation. Consequently, facial semantics are linear-encoded as latent representations. Based on StyleGAN, recent works show that sampling along the linear-encoded representation vector in latent space will change the associated facial semantics accordingly [32], which makes it possible to manipulate the face generations to meet a target requirement. However, in current frameworks, mapping a particular facial semantics to a latent representation vector relies on training offline classifiers with manually labeled datasets. Thus they require artificially defined semantics and provide the associated labels for all facial images. The disadvantages for training with labeled facial semantics include: first, they demand extra effort on human annotations for each new attributes proposed; second, each semantics is defined artificially, and the scope of semantics is limited to the linear

combination of such definitions; and third, by only training on each labeled semantics independently, we are unable to give any insights on the connections among different semantics.

In this work, we explore unsupervised methods to minimize the demand for human annotation. We propose a novel unsupervised framework to disentangle and manipulate facial semantics under the StyleGAN environment, while still maintain the interpretability for semantics (Fig. 1) as in labeled datasets.

- We motivate *decorrelation regularization* on StyleGAN to further enhance disentanglement for the latent representation.
- We introduce *mutual reconstruction* to stabilize training of an unsupervised 3D deformable face reconstruction method, such that it serves as an initial facial semantic extractor.
- For univariate semantics, e.g., yaw angle, we present a *linear regression method* to capture their perturbations from latent space. Given the manipulation vector, we further demonstrate the success of yaw manipulation on data augmentation to trivialize the unbalanced pose problem within the unsupervised paradigm.
- For pixel-level semantics, e.g., shape and texture, we propose a *localized representation learning algorithm* to capture sparse semantic perturbations from latent space. The associated analysis evinces the effectiveness of our method to provide interpretability without external supervision.

All methods proposed are purely based on a label-free training strategy. Only StyleGAN is trained with an in-the-wild face dataset. Therefore, we reduce a significant amount of human involvement in facial representation learning. Furthermore, with zero labels, our framework provides an unconstrained environment for the disentanglement algorithms to explore and shed light on how interpretable representations are learned in cutting-edge neural network models.

## 2. Related Works

### 2.1. Exploring latent space representations in GANs

GAN [12] is widely explored on applications with the photo-realistic images [13, 27, 28, 44]. The understanding of how GANs learn to construct the latent space with the semantics information defined in the real visual world draws attention recently [32, 34].

Prior exploration of the latent space of GANs focuses on smoothly varying the output image from one synthesis to another with the interpolation in the latent space without

considering the underlying semantics [31, 21]. A simultaneous optimization strategy on the generator and latent code is developed to learn a better constructed latent space [3]. Other works with recurrent neural networks also explore the latent code space with the semantics information. One of which, [16] interprets the steer-ability from the perspective of camera motion and image color tone. Another [41] uses the hierarchical semantics to understand the deep generative representations for scene synthesis. There are also explorations using facial attributes for face synthesis, e.g. InterFaceGAN [32], which provided a detailed analysis of the semantics encoded in the latent space. They take into consideration both single and disentangled multi-semantics and, by a fixed pretrained GAN, it learns to explore semantics by relating the latent space with labeled semantic attributes on the synthesized images. DiscoFaceGAN [8] uses the pretrained 3DMM [2] parameters to provide guidance for interpretable face manipulation. Overall, The major commonsense for the learning of latent space semantic representation is that it is supervised and only in that way, the semantics are interpretable and controllable.

The editing of synthesized faces with GANs is another active research area. Face editing is generally conducted with semantic information such as facial attributes. With facial attributes act together, face images can be therefore constructed. Thus, face editing requires the GANs to have the ability to edit the disentangled information. In other words, the editing should change the target attribute but keep other information ideally unchanged. To achieve this effect, methods are mainly focused on three aspects: the comprehensive design of loss functions [6, 29, 36], the involvement of additional attribute features [35, 1, 22, 43] and the architecture designs [33, 10, 24, 45, 7]. However, these works either discard the latent code, resulting in the inability to continuously interpolate certain semantics, or fail to provide the same synthesis quality compared with state-of-the-art GANs [17, 18].

### 2.2. Unsupervised disentanglement method

Generally speaking, all the vanilla GANs are unsupervised representation learning methods because they only require training on unlabeled datasets. Variational auto-encoders (VAEs) [20], on the other hand, achieves a similar goal with self-reconstruction loss. Recent efforts have been made to make VAEs generate images as good as GANs [37, 38]. Unsupervised disentanglement methods both exists in GANs [6, 23] and VAEs [15, 19]. However, most of these unsupervised disentangled methods only work under simple datasets and are considered not disentangle realistic information [26]. According to [26], one should always introduce either supervision or inductive biases to the disentanglement method to achieve meaningful representations.

Examples of inductive bias rise from the symmetry of



Figure 2: The canonical albedo map of samples in CASIA WebFace after training for 60k iterations with batch size of 64.

natural objects and the 3D graphical information. Recent works successfully reconstruct the face images by carefully remodeling the graphics of camera principal [40]. This work makes it possible to decompose the facial images into environmental semantics and other facial semantics. However, a major drawback of this work is that it is unable to generate realistic faces and perform pixel-level face editing on it.

### 3. Method

The goal of our facial representation disentanglement method is to capture linear-encoded facial semantics. With a given collection of coarsely aligned faces, a Generative Adversarial Network (GAN) is trained to mimic the overall distribution of the data. In our case, to better learn linear-encoded facial semantics, we re-implement and train StyleGAN-2 [18] as it is well known for the ability to disentangle representations in latent space. Further, we investigate the latent space trained by StyleGAN and improve its capability to disentangle by adding a decorrelation regularization (Sec. 3.1). After training a StyleGAN model, we use the faces it generates as training data and trains a 3D deformable face reconstruction method modified from [40]. A mutual reconstruction strategy (Sec. 3.2) stabilizes the training significantly. Then, we keep a record of the latent code from StyleGAN and apply linear regression to disentangle the target semantics in the latent space (Sec. 3.3). Meanwhile, taking the reconstruction of the yaw angle as an example, we manipulate the latent representation as a data augmentation for training (Sec. 3.4). Finally, we describe the localized representation learning method to disentangle canonical semantics in Sec. 3.5.

#### 3.1. Decorrelating latent code in StyleGAN

In StyleGAN design, a latent code  $\mathbf{z} \in \mathcal{Z}^{d \times 1}$  is randomly generated from a Gaussian distribution. Then, a mapping network takes in  $\mathbf{z}$  and output a latent code  $\mathbf{w} \in \mathcal{W}^{d \times 1}$ . Space  $\mathcal{W}$  is proven to facilitate the learning of more disentangled representations. In our study, we find that we can further enhance the disentangled representation by

decorrelating latent codes in  $\mathcal{W}$ . Intuitively, a more decorrelated latent space enforces more independent dimensions to encode information, therefore encourages disentanglement in representations. In order to maximize the utilization of all dimensions in  $\mathcal{W}$ , we want to make all Pearson correlation coefficient  $\rho_{ij}$  to zero and variance of all dimensions  $\text{Var}[w_i]$  the same value, where  $i, j$  stand for the subscript of dimensions in  $\mathcal{W}$  space. Therefore, we introduce decorrelation regularization via a loss function:

$$\mathcal{L}_{decorr} = - \sum_{i \neq j} (\log(1 - |\rho_{ij}|)) + \sum_i \left( \text{Var}[w_i] - \sum_j (\text{Var}[w_j]) \right)^2. \quad (1)$$

Here,  $\rho_{ij}$  and  $\text{Var}[w_i]$  are all estimated by sampling  $\mathbf{w}$  from the mapping network  $\mathcal{F}(\mathbf{z})$ , given  $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ .

The overall objective for GAN with decorrelation regularization follows:

$$\min_{\mathcal{F}, \mathcal{G}} \max_{\mathcal{D}} \mathcal{L}_{GAN} + \mathcal{L}_{decorr},$$

where  $\mathcal{G}$  and  $\mathcal{D}$  stand for the generator and discriminator for GAN, respectively. Here the mapping network  $\mathcal{F}$  is the only one to update with the new loss,  $\mathcal{L}_{decorr}$ .

#### 3.2. Stabilized training for 3D face reconstruction

The unsupervised 3D deformable face reconstruction method [40] takes a roughly aligned face image and decomposes the faces into multiple semantics, i.e. view, lighting, albedo, and depth ( $\mathbf{y}_v, \mathbf{y}_l, \mathbf{y}_a$  and  $\mathbf{y}_d$ , respectively). During training, it uses these decomposed semantics to reconstruct the original input image  $I$  with the reconstruction loss:

$$\mathcal{L}_{recon} = \mathcal{C}(I)^\top |I - \hat{I}|, \quad \text{where } \hat{I} = \mathcal{R}(\mathbf{y}_v, \mathbf{y}_l, \mathbf{y}_a, \mathbf{y}_d),$$

where  $\mathcal{C}$  and  $\mathcal{R}$  stand for a confidence network and a 3D face renderer. We use this method to pre-decompose some external facial semantics, i.e. pose and lighting, from StyleGAN generations.

However, we find that the 3D face reconstruction algorithm struggles to estimate the pose of profile or near-profile faces. This finding coincides with the failure case analysis in the original paper [40]. To make things worse, if the dataset contains a decent number of profile and near-profile faces (e.g. CASIA WebFace), the 3D reconstruction fails to learn physically sounded semantics (Shown in Fig. 2(a)) and collapses into a sub-optimal state. That is, the algorithm tries to use extreme values to estimate the texture and shape of each face *independently*, which deviate far away from the actual texture and shape of the face.

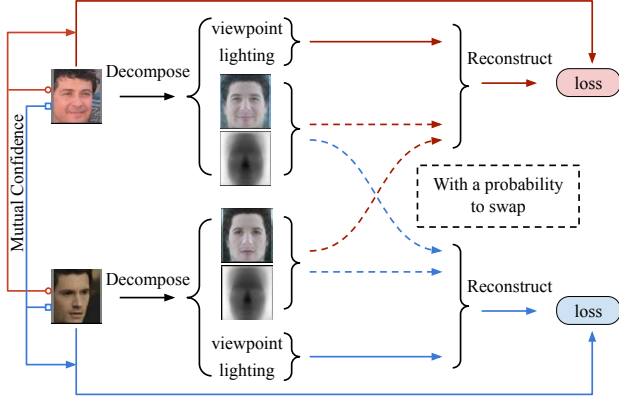


Figure 3: The implementation of mutual reconstruction during training.

To deal with the problem, we develop a mutual reconstruction strategy as illustrated in Fig. 3. To start with, we know that in general, human faces are very similar in shape and texture. Therefore, each face image should still be able to reconstruct itself comparatively if we swap its shape and texture with another face. Eventually, the strategy prevents the model from using extreme values to fit individual reconstruction, and the model learns to reconstruct faces with a minimum variance of the shape and texture among all samples, even for profile faces. Following this idea, during training, we swap the albedo and depth map between two images with a probability  $\epsilon$  to perform the reconstruction with the alternative loss:

$$\tilde{\mathcal{L}}_{recon} = \tilde{\mathcal{C}}(I, I')^T |I - \hat{I}'|,$$

where  $\hat{I}' = \mathcal{R}(\mathbf{y}_v, \mathbf{y}_l, \mathbf{y}'_a, \mathbf{y}'_d)$ ,

where  $\tilde{\mathcal{C}}$  is a mutual confidence network; and everything with the prime notation originates from another image  $I'$ . The overall loss to reconstruct each image becomes:

$$(1 - \epsilon)\mathcal{L}_{recon} + \epsilon\tilde{\mathcal{L}}_{recon}.$$

As a result, the shape and texture of faces with deviated facial semantics can be robustly estimated.

Moreover, since images are now reconstructed with two images, the confidence map in the original work should be yielded by these two images accordingly. We simply concatenate the two images channel-wise as input to the confidence network, where the top image provides environmental semantics and the bottom image provides texture and shape information.

### 3.3. Disentangle semantics with linear regression

With the 3D face reconstruction algorithm, face images generated by StyleGAN are decomposed to pose, lighting, depth, and albedo. Remember, the ultimate goal of disentangling semantics is to find a vector  $\mathbf{v} \in \mathcal{W}$  in StyleGAN, such that it only takes control of the target semantics.

**Semantic gradient estimation:** Now consider a semantics  $y$  of a generated face image  $\mathcal{G}(\mathbf{w})$  that can be measured by a function  $f(\cdot)$ . The linear approximation of the gradient  $\nabla_y$  with respect to the latent code  $\mathbf{w}$  satisfies:

$$f(\mathcal{G}(\mathbf{w}_1)) \approx f(\mathcal{G}(\mathbf{w}_0)) + \nabla_y(\mathbf{w}_0)^T (\mathbf{w}_1 - \mathbf{w}_0).$$

Note that, generally, the gradient at location  $\mathbf{w}_0$ ,  $\nabla_y(\mathbf{w}_0)$ , is a function of latent code  $\mathbf{w}_0$ . However, with StyleGAN, it is observed that many semantics can be linear-encoded in the disentangled latent space  $\mathcal{W}$  [17, 32]. We take advantage of this merit from StyleGAN and assume all the semantics can be linear-encoded. In other words, the gradient is no longer independent of the input latent code  $\mathbf{w}_0$ . We get:

$$f(\mathcal{G}(\mathbf{w}_1)) \approx f(\mathcal{G}(\mathbf{w}_0)) + \nabla_y^T (\mathbf{w}_1 - \mathbf{w}_0),$$

or to be simplified as:

$$\Delta y \approx \nabla_y^T \Delta \mathbf{w},$$

where  $\Delta y = f(\mathcal{G}(\mathbf{w}_1)) - f(\mathcal{G}(\mathbf{w}_0))$  and  $\Delta \mathbf{w} = \mathbf{w}_1 - \mathbf{w}_0$ .

**Semantic linear regression:** Now it is obvious that in the ideal case, the target vector  $\mathbf{v} = \nabla_y$ . While in real world scenario, the gradient  $\nabla_y$  is hard to estimate directly because back-propagation only captures local gradient, making it less robust to noises. Therefore we propose a linear regression model to capture global linearity for gradient estimation. We randomly sample  $N$  pairs of  $(\mathbf{w}_1, \mathbf{w}_0)$ , generate images with StyleGAN and estimate its semantics. Finally, all samples of differences are concatenated, denoted as  $\Delta Y \in \mathbb{R}^{N \times 1}$  for semantics and  $\Delta \mathbf{W} \in \mathbb{R}^{N \times d}$  for latent codes. The objective is to minimize:

$$\min_{\mathbf{v}} \|\Delta Y - \Delta \mathbf{W} \mathbf{v}\|_2^2. \quad (2)$$

We have a closed-form solution when  $N > d$ :

$$\mathbf{v} = (\Delta \mathbf{W}^T \Delta \mathbf{W})^{-1} \Delta \mathbf{W}^T \Delta Y. \quad (3)$$

### 3.4. Image manipulation for data augmentation:

One useful application for guided image manipulation is to perform data augmentation. Data augmentation has proven to be efficient when dealing with unbalanced data during training. One related problem within our unsupervised framework is the inaccurate estimation of extreme yaw angle, which is also mentioned in [40]. This problem worsens when dealing with generations from CASIA StyleGAN since it contains a large amount but a small portion of profile faces (unbalanced yaw distribution).

In our experiment setting, we propose a data augmentation strategy based on self-supervised facial image manipulation. The end goal is to help the 3D face reconstruction network estimate the extreme yaw angle accurately. With



Table 1: Results on Decorrelation Regularization in StyleGAN

Decorrelated latent codes	CASIA (96 × 112)		FFHQ (256 <sup>2</sup> )	
	FID	PPL	FID	PPL
No	11.85	15.26	25.49	34.24
Yes	<b>9.96</b>	<b>12.30</b>	<b>15.57</b>	<b>20.23</b>

the linear regression method mentioned in Sec. 3.3, we can learn manipulation vectors  $\mathbf{v}$  for univariate semantics including the yaw angle, denoted as  $\mathbf{v}_{yaw}$ . Recall that by extrapolating along  $\mathbf{v}$  beyond its standard deviation, preferably, we can get samples with more extreme values for the associated semantics. Particularly, we can generate images with an extreme yaw angle to neutralize the unbalanced yaw distribution and train the 3D face reconstruction algorithm. Therefore, by seeing more profile faces deliberately, the system can better estimate extreme yaw angles.

We perform the data augmentation strategy alongside with the training of 3D face reconstruction. To be specific, we estimate  $\mathbf{v}$  and update with a historical moving average (momentum= 0.995) every 10 iterations. Then, augmentation is achieved by extrapolating a latent vector  $\mathbf{w}_i^{(s)}$  via:

$$\mathbf{w}_i^{(s)} = \mathbf{w}_i - \mathbf{w}_i^\top \mathbf{v} \mathbf{v} + s \cdot \sigma_{\mathbf{w}} \mathbf{v}, \quad (4)$$

where  $\mathbf{w}_i$  is a random sample drawn from  $\mathcal{F}(\mathbf{z})$ ,  $s$  is the scaling factor for the interpolation/extrapolation along the unit length manipulation vector  $\mathbf{v}$ . And  $\sigma_{\mathbf{w}}$  is the standard deviation for  $\mathbf{w}^\top \mathbf{v}$ . In this case  $\mathbf{v} = \mathbf{v}_{yaw}$ . Finally, the 3D face reconstruction method is trained with the augmented generations  $\mathcal{G}(\mathbf{w}_i^{(s)})$ .

### 3.5. Localized representation learning

In the case where  $f(\cdot)$  returns canonical outputs, i.e. depth and albedo maps, the outputs consist of pixels in spatial dimensions and the pixel values are highly correlated as the latent code changes. However, every pixel-level gradient estimation, i.e.  $\mathbf{v}$ , from Eqn. 3 is independently calculated and is thus extremely redundant. To deal with this problem, we reformulate our goal for canonical semantics; that is, to find the manipulation vectors  $\hat{\mathbf{v}}$  that capture interpretable *combinations* of pixel value variations. We start by defining a Jacobian matrix  $\mathbf{J}_{\mathbf{v}} \in \mathbb{R}^{S \times d}$ , which is the concatenation of all canonical pixel-level  $\mathbf{v}$ . Here  $S$  stands for the overall number of spatial and RGB dimensions of a given depth and albedo map.

One trivial definition of  $\hat{\mathbf{v}}$  is that it maximizes  $\|\mathbf{J}_{\mathbf{v}}^* \hat{\mathbf{v}}\|_2^2$ . However, we need to keep in mind that, ideally we expect a disentangled representation to manipulate *interpretable* facial semantics. That is to say, interpolation along  $\hat{\mathbf{v}}$  should result in significant but *localized* (i.e. sparse) change across the image domain, e.g. some  $\hat{\mathbf{v}}$  only control eye variations

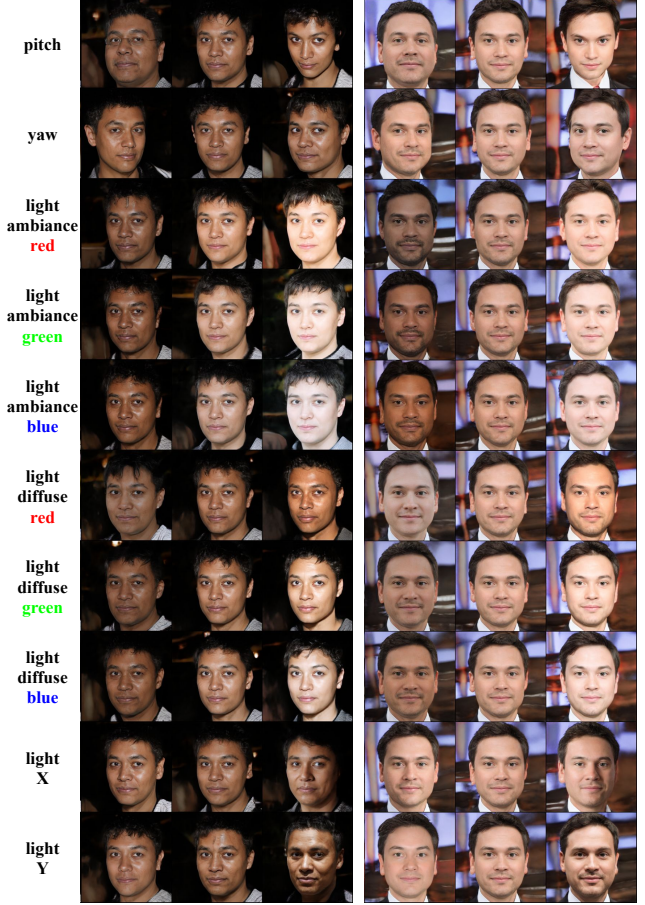


Figure 4: Two examples of the manipulation result on univariate semantics in StyleGAN trained with FFHQ.

while some only control mouth variations, etc.. However,  $\|\mathbf{J}_{\mathbf{v}}^* \hat{\mathbf{v}}\|_2^2$  captures the *global* pixel value perturbation. Thus, we propose the **localized canonical representation learning** by solving:

$$\begin{aligned} \min_{\mathbf{U}, \hat{\mathbf{V}}} \quad & \left\| \mathbf{J}_{\mathbf{v}}^* - \mathbf{U} \hat{\mathbf{V}}^\top \right\|_F^2 + \alpha \|\mathbf{U}\|_1 + \beta \sum_{i \neq j} (\hat{\mathbf{v}}_i^\top \hat{\mathbf{v}}_j)^2 \\ \text{s.t.} \quad & \|\hat{\mathbf{v}}_p\|_2 = 1, \end{aligned} \quad (5)$$

where  $p, i, j \in \{1, \dots, P\}$  and  $P$  is the number of components to learn. Each column in  $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_P] \in \mathbb{R}^{S \times P}$  is a sparse component of the canonical albedo and depth perturbation, and  $\hat{\mathbf{V}} = [\hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_P] \in \mathbb{R}^{d \times P}$  consists of the associated latent representation in  $\mathcal{W}$  space.  $\alpha$  and  $\beta$  are tuning parameters to control the trade-off among the reconstruction accuracy of  $\mathbf{J}_{\mathbf{v}}^*$ , sparsity of perturbation in semantics and orthogonality of associated latent representations.

## 4. Experiments

The datasets used for training StyleGAN is CASIA Web-Face [42] or Flickr-Face-HQ (FFHQ). We choose the two datasets since they represent very different types of face im-

Table 2: Correlation coefficients between learned latent representations and ground truth facial semantics.

latent representations	facial semantics	
	pitch	yaw
pitch	<b>0.7647</b>	0.0146
yaw	0.0390	<b>0.9458</b>
ambiance lighting red	0.0195	0.1001
ambiance lighting green	0.0414	0.0913
ambiance lighting blue	0.0134	0.1175
diffuse light red	0.0528	0.0402
diffuse light green	0.1051	0.1243
diffuse light blue	0.0823	0.1250
light direction X	0.0338	0.2316
light direction Y	0.2166	0.0057

ages, CASIA WebFace is at low resolution and high pose variation, while FFHQ is at high resolution, high environmental variations but mostly frontal faces. We demonstrate the effectiveness of our unsupervised disentanglement method in both cases. For CASIA WebFace, we roughly align and crop the faces to  $112 \times 96$  such that the faces are all in the center of the image with minimum roll angle variations. The images are resized to  $128 \times 128$  before being fed into StyleGAN for training. For FFHQ, we resize the original images to  $256 \times 256$  for training StyleGAN. All StyleGAN generations are resized to  $64 \times 64$  to decompose into semantics via the 3D reconstruction algorithm.

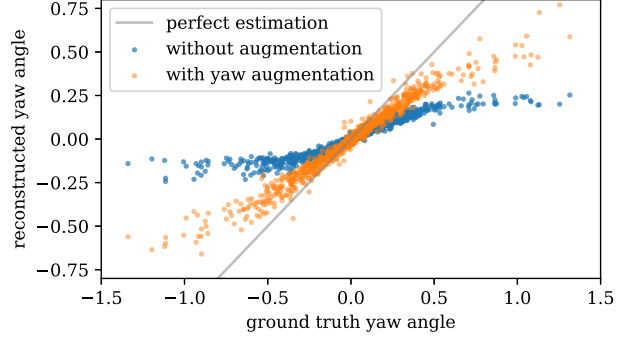
#### 4.1. Effectiveness of decorrelation regularization

We follow the implementation and keep all parameters the same as in StyleGAN-2 [18] and train until the model sees a total of 25 million images. The decorrelation regularization applies its update purely on the mapping network of StyleGAN. We test the performance of the model in terms of FID [14] and PPL [17] metrics.


The results shown in Table 1 indicate that by decorrelating the latent codes in  $\mathcal{W}$ , the generated images become more realistic, meanwhile, the latent representations are more disentangled.

#### 4.2. Stabilized 3D face reconstruction

CASIA WebFace [42] is a more challenging dataset for the original 3D face reconstruction algorithm [40] since it contains faces with extreme pose variations. We illustrate the difficulty by training the algorithm with its default parameter setting. Note that the default setting successfully trained with the FFHQ dataset in our experiments and




(a) The correlation of estimated yaw angle vs the ground truth yaw angle.



(b) GAN generations.



(c) Reconstruction without augmented training.



(d) Reconstruction with augmented yaw training.

Figure 5: The reconstruction result for extreme yaw generations from CASIA StyleGAN.

CelebA [25] dataset as reported in their original work.

We stabilize the training process by applying the mutual reconstruction strategy and set the swap probability  $\epsilon = 0.5$  for CASIA WebFace. All faces are decomposed into reasonable semantics with a stabilized estimation of depth and albedo (examples are shown in Fig. 2). We also set the swap probability  $\epsilon = 0.1$  for FFHQ training for a more smoothed estimation of the depth and albedo.

#### 4.3. Results on univariate semantics manipulation

In our experiment setting, there are 14 univariate semantics: facial pose (pitch, yaw, roll), face translation (X, Y, Z-axis), ambiance lighting (RGB), diffuse lighting (RGB and direction in X, Y-axis). We apply our linear regression method to get all manipulation vectors for these semantics. Then, the interpolation along each learned manipulation vector is performed with Eqn. 4.

The generation result for StyleGAN trained with FFHQ is shown in Fig. 4.

**Testing Setup:** To test the behavior of our manipulation vectors, we examine the correlation coefficient between projection length  $\mathbf{w}_i^\top \mathbf{v}$  and the ground truth value of the associated semantics. We showcase our result on StyleGAN trained with FFHQ in terms of pitch, yaw angle since the ground truth can be estimated with existing open-source




Figure 6: The clustering result with Ward variance minimization algorithm [39]. The Y-axis stand for the Ward’s distance metric calculated from the absolute value of cosine distance. **Green** cluster: consists of all perturbations in the background. **Red** cluster: contains most facial perturbations inside the face area. **Cyan** cluster: contains mostly facial perturbations around the face region.

3D landmarker [4]. We ignore roll and all translation semantics because they are almost removed during data pre-processing, and as a result, StyleGAN is not sensitive to these semantics. For lighting semantics, on the other hand, we are currently unable to find any method to accurately estimate their ground truth, thus it is impossible to test their performance quantitatively. Nevertheless, we showcase the visual results of all univariate semantics of both FFHQ and CASIA StyleGAN in the supplementary materials.

**Testing Results:** In Table 2, all coefficient values greater than 0.5 are labeled as bold. It showcases that the pitch and yaw angles are highly correlated with their associated latent representation. Meanwhile, a perturbation in semantically irrelevant latent representations is less likely to affect pitch and yaw. One unexpected phenomenon is the light direction in the X and Y-axis seems to be slightly correlated to yaw and pitch angle, respectively. This is because the view-point aligns with the lighting within the same 3D direction in the real world, resulting in a more entangled generation in GANs. The performance drop under extreme lighting conditions is also mentioned in the failure analysis in [40].

**Results on self-supervised data augmentation:** We apply yaw angle augmentation on generation from CASIA StyleGAN to improve 3D facial reconstruction, because the generations contain profile or near-profile faces that are hard to reconstruct. The augmentation scale  $s$  is a random uniform distributed value with range  $[-10, 10]$  for extreme extrapolation. We stick to the same optimization settings and compare the profile face matching performance. As shown in Fig 5, the 3D facial reconstruction algorithm predicts the yaw angle more accurately when trained with yaw augmented generations. We also showcase the reconstruction results of randomly selected profile or near profile examples.

Due to the success of data augmentation, canonical semantics are ensured to have been better estimated for CA-

SIA StyleGAN generations as well. Thus, for the following canonical semantics, we only report results on the data augmented model for CASIA StyleGAN.

#### 4.4. Results on canonical semantics manipulation

For our FFHQ experiments, we set  $\alpha = 1$ ,  $\beta = 1$  and  $P = 200$  in objective 5 and optimize with a Adam optimizer with learning rate of 0.0001. While for CASIA WebFace experiments, we set  $\alpha = 5$ ,  $\beta = 10$  for best performance. Note that the range for selecting  $\alpha$  and  $\beta$  is not very strict. We showcase some results by changing the  $\alpha$ ,  $\beta$  values in the supplementary materials. The training converges after 500k iterations (25 mins in single RTX 2080 GPU).

**Optimization Result:** We discard any  $\hat{\mathbf{v}}_p$  if its associated localized facial semantics has  $\|\mathbf{u}_p\| < 0.01$ , i.e. these  $\hat{\mathbf{v}}_p$  does not encode linear semantics in albedo and depth map. Finally, 43 components remain valid for FFHQ StyleGAN. We illustrate their albedo maps as the X-axis in Fig. 6. Among these components, 28 of them are associated with non-background localized facial semantics. In Fig. 7, we demonstrate some example results when performing interpolation/extrapolation along specific localized facial representations  $\hat{\mathbf{v}}_p$ . The sampling scale  $s$  in Eqn. 4 ranges from -4 to 4. Note that samples with extreme scale are typically rare for StyleGAN to generate under normal circumstances [32]. Therefore it is likely for the representation to be less disentangled when extrapolating at a high scale.

**Interpretability of Manipulation:** Due to the sparse constraint applied to the localized facial semantics, the resulting  $\mathbf{u}_p$  only captures perturbation of limited areas. The sparsity significantly enhances the interpretability of the disentangled facial semantics. For example, in Fig. 7, the semantics is fairly easy to be interpreted.

**Disentanglement Analysis:** We further check whether these image-level perturbations are semantically encoded in the latent space. The latent representations for semantically






Figure 7: Sampling along the localized representations. A possible interpretation: **first column**: hooked nose, height of nasal bridge, nasolabial folds, eye glasses, height of eyebrow, mustache/beard, smiling; **second column**: bags under eyes, darkness of hair, brightness of forehead, depth of orbital rim, opened mouth, long hair, bangs on the forehead.

independent perturbations are supposed to be more orthogonal with each other; and vice versa. To showcase the result of this test, we decide to use a linkage clustering algorithm on the latent representations  $\hat{v}_p$  and apply Ward’s variance minimization [39]. The absolute value of the pairwise cosine distance between latent representations is used as the distance metric for the clustering. The result is shown in Fig. 6, where three major clusters are detected. Amazingly, the formation of the clusters is according to the semantics of the facial image, i.e. background areas, facial areas, and areas around the face. Moreover, we find that although the background cluster is tightly distributed, the representations in the other two clusters are fairly independent, as we see the representations in non-background clusters split in a relatively earlier stage in the dendrogram. In general, the localized canonical representations are surprisingly easy to interpret and are complying with human intuitions.

We also test the semantic correlation of the localized facial semantics and CelebA [25] facial attributes. We showcase the similarity by concatenating CelebA SVM weight vectors with  $\hat{v}_p$  and following the same process to generate

Fig. 6. The result shows that some semantics are highly correlated with facial attributes, e.g. glasses, smiling, big nose, big lips, etc. Full result is shown in the supplementary materials.

Besides, CASIA StyleGAN follows a similar trend, where we report in the supplementary materials.

## 5. Conclusion

We have presented an unsupervised learning framework for disentangling linear-encoded facial semantics from StyleGAN. The system can robustly decompose facial semantics from any single view GAN generations and disentangle facial semantics that is easily interpretable. We also illustrate a potential direction to get the facial manipulation frameworks to work, i.e. performing guided data augmentation to counteract the problem brought by unbalanced data. Extensive analysis suggests that the manipulation of localized facial semantics is easily interpretable and intuitive.



## References

- [1] Jianmin Bao, Dong Chen, Fang Wen, Houqiang Li, and Gang Hua. Towards open-set identity preserving face synthesis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6713–6722, 2018. 2
- [2] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194, 1999. 2
- [3] Piotr Bojanowski, Armand Joulin, David Lopez-Paz, and Arthur Szlam. Optimizing the latent space of generative networks. *arXiv preprint arXiv:1707.05776*, 2017. 2
- [4] Adrian Bulat and Georgios Tzimiropoulos. How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks). In *International Conference on Computer Vision*, 2017. 7
- [5] Ricky TQ Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. In *Advances in Neural Information Processing Systems*, pages 2610–2620, 2018. 1
- [6] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, pages 2172–2180, 2016. 2
- [7] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8789–8797, 2018. 2
- [8] Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, and Xin Tong. Disentangled and controllable face image generation via 3d imitative-contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5154–5163, 2020. 2
- [9] Guillaume Desjardins, Aaron Courville, and Yoshua Bengio. Disentangling factors of variation via generative entangling. *arXiv preprint arXiv:1210.5474*, 2012. 1
- [10] Chris Donahue, Zachary C Lipton, Akshay Balsubramani, and Julian McAuley. Semantically decomposing the latent spaces of generative adversarial networks. *arXiv preprint arXiv:1705.07904*, 2017. 2
- [11] Cian Eastwood and Christopher KI Williams. A framework for the quantitative evaluation of disentangled representations. In *International Conference on Learning Representations*, 2018. 1
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 1, 2
- [13] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in neural information processing systems*, pages 5767–5777, 2017. 2
- [14] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems*, pages 6626–6637, 2017. 6
- [15] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. 2016. 1, 2
- [16] Ali Jahanian, Lucy Chai, and Phillip Isola. On the “steerability” of generative adversarial networks. *arXiv preprint arXiv:1907.07171*, 2019. 2
- [17] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 1, 2, 4, 6
- [18] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. 2, 3, 6
- [19] Hyunjik Kim and Andriy Mnih. Disentangling by factorising. *arXiv preprint arXiv:1802.05983*, 2018. 1, 2
- [20] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2
- [21] Samuli Laine. Feature-based metrics for exploring the latent space of generative models. 2018. 2
- [22] Guillaume Lample, Neil Zeghidour, Nicolas Usunier, Antoine Bordes, Ludovic Denoyer, and Marc’Aurelio Ranzato. Fader networks: Manipulating images by sliding attributes. In *Advances in neural information processing systems*, pages 5967–5976, 2017. 2
- [23] Zinan Lin, Kiran Koshy Thekumparampil, Giulia Fanti, and Sewoong Oh. Infogan-cr: Disentangling generative adversarial networks with contrastive regularizers. *arXiv preprint arXiv:1906.06034*, 2019. 2
- [24] Ming Liu, Yukang Ding, Min Xia, Xiao Liu, Errui Ding, Wangmeng Zuo, and Shilei Wen. Stgan: A unified selective transfer network for arbitrary image attribute editing. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3673–3682, 2019. 2
- [25] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. 6, 8
- [26] Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. In *international conference on machine learning*, pages 4114–4124. PMLR, 2019. 2
- [27] SC Martin Arjovsky and Leon Bottou. Wasserstein generative adversarial networks. In *Proceedings of the 34 th International Conference on Machine Learning, Sydney, Australia*, 2017. 2

- [28] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018. [2](#)
- [29] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *International conference on machine learning*, pages 2642–2651, 2017. [2](#)
- [30] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014. [1](#)
- [31] Hang Shao, Abhishek Kumar, and P Thomas Fletcher. The riemannian geometry of deep generative models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 315–323, 2018. [2](#)
- [32] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9243–9252, 2020. [1](#), [2](#), [4](#), [7](#)
- [33] Yujun Shen, Ping Luo, Junjie Yan, Xiaogang Wang, and Xiaoou Tang. Faceid-gan: Learning a symmetry three-player gan for identity-preserving face synthesis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 821–830, 2018. [2](#)
- [34] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. Interfacegan: Interpreting the disentangled face representation learned by gans. *arXiv preprint arXiv:2005.09635*, 2020. [2](#)
- [35] Yujun Shen, Bolei Zhou, Ping Luo, and Xiaoou Tang. Facefeat-gan: a two-stage approach for identity-preserving face synthesis. *arXiv preprint arXiv:1812.01288*, 2018. [2](#)
- [36] Luan Tran, Xi Yin, and Xiaoming Liu. Disentangled representation learning gan for pose-invariant face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1415–1424, 2017. [2](#)
- [37] Arash Vahdat and Jan Kautz. Nvae: A deep hierarchical variational autoencoder. *Advances in Neural Information Processing Systems*, 33, 2020. [2](#)
- [38] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, pages 6306–6315, 2017. [2](#)
- [39] Joe H Ward Jr. Hierarchical grouping to optimize an objective function. *Journal of the American statistical association*, 58(301):236–244, 1963. [7](#), [8](#)
- [40] Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Unsupervised learning of probably symmetric deformable 3d objects from images in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1–10, 2020. [3](#), [4](#), [6](#), [7](#)
- [41] Ceyuan Yang, Yujun Shen, and Bolei Zhou. Semantic hierarchy emerges in deep generative representations for scene synthesis. *arXiv preprint arXiv:1911.09267*, 2019. [2](#)
- [42] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014. [5](#), [6](#)
- [43] Xi Yin, Xiang Yu, Kihyuk Sohn, Xiaoming Liu, and Manmohan Chandraker. Towards large-pose face frontalization in the wild. In *Proceedings of the IEEE international conference on computer vision*, pages 3990–3999, 2017. [2](#)
- [44] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *International Conference on Machine Learning*, pages 7354–7363. PMLR, 2019. [2](#)
- [45] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. [2](#)