

# VecGAN: Image-to-Image Translation with Interpretable Latent Directions

Yusuf Dalva, Said Fahri Altındaş, and Aysegul Dundar

Bilkent University  
{yusuf.dalva, fahri.altindis}@bilkent.edu.tr  
adundar@cs.bilkent.edu.tr

**Abstract.** We propose VecGAN, an image-to-image translation framework for facial attribute editing with interpretable latent directions. Facial attribute editing task faces the challenges of precise attribute editing with controllable strength and preservation of the other attributes of an image. For this goal, we design the attribute editing by latent space factorization and for each attribute, we learn a linear direction that is orthogonal to the others. The other component is the controllable strength of the change, a scalar value. In our framework, this scalar can be either sampled or encoded from a reference image by projection. Our work is inspired by the latent space factorization works of fixed pretrained GANs. However, while those models cannot be trained end-to-end and struggle to edit encoded images precisely, VecGAN is end-to-end trained for image translation task and successful at editing an attribute while preserving the others. Our extensive experiments show that VecGAN achieves significant improvements over state-of-the-arts for both local and global edits.

**Keywords:** Image translation, generative adversarial networks, latent space manipulation, face attribute editing.

## 1 Introduction

There has been a significant progress in image-to-image translation methods [13, 24, 35, 6, 21, 37, 20, 23] especially for facial attribute editing [5, 25, 33, 38, 19] powered with generative adversarial networks (GANs). A main challenge of facial attribute editing methods is to be able to change only one attribute of an image without affecting others such as global lighting parameters of the images, identity of the persons, background, or their other attributes. The other challenge is the interpretability of the style codes so that one can control the attribute intensity of the edit, e.g. increase the intensity of smile or aging.

To achieve the targeted attribute editing while preserving the others, many works set a separate style encoder and an image editing network where modified styles are injected into it [5, 19]. During image-to-image translation, a style encoded from another image or a newly sampled style latent code can be used to output diverse images. To disentangle attributes, works focus on style encoding

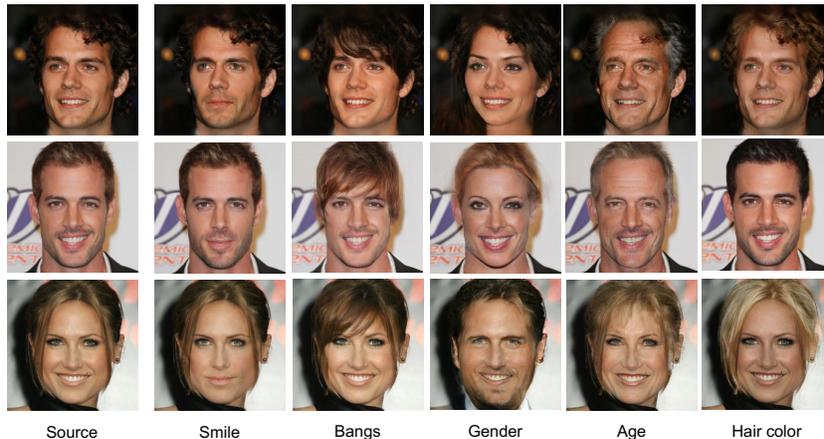


Fig. 1: Attribute editing results of VecGAN. The first column shows the source images, and other columns show the results of editing a specific attribute. Each edited image has an attribute value opposite to that of the source one. For hair color, sources are translated to brown, black, and blonde hair, respectively.

and progress from a shared style code, SDIT [30], to mixed style codes, StarGANv2 [5], to hierarchical disentangled styles, HiSD [19]. Among these works, HiSD independently learn styles of each attribute, bangs, hair color, glasses and introduces a local translator which uses attention masks to avoid global manipulations. HiSD showcases successes on those three local attribute editing task and is not tested for global attribute editing, e.g. age, smile. Furthermore, one limitation of these works is the uninterpretability of style codes as one cannot control the intensity of attribute (e.g. blondness) in a straight-forward manner.

To overcome the challenges of facial attribute editing task, we propose a novel framework, VecGAN, and image-to-image translation framework with interpretable latent directions. Our framework does not require a separate style encoder as in the previous works since we achieve the translation in the encoded latent space directly. The attribute editing directions are learned in the latent space and regularized to be orthogonal to each other for style disentanglement. The other component of our framework is the controllable strength of the change, a scalar value. This scalar can be either sampled from a distribution or encoded from a reference image by projection in the latent space. Our framework not only achieves significant improvements over state-of-the-arts for both local and global edits but also provides a knob to control the editing attribute intensity via its design.

VecGAN is encouraged by the findings that well-trained generative models organize their latent space as disentangled representations with meaningful directions in a completely unsupervised way. Exploring these interpretable directions in latent codes has emerged as an important research endeavor on the fixed pretrained GANs [26, 28, 9, 27, 32]. These works show that images can be

mapped to the GANs latent space and edits can be achieved by manipulations in the latent space. However, since these models are not trained end-to-end, the results are sub-optimal as will also be shown in our experiments.

To enable VecGAN, different than previous works of image-to-image translation networks, we use a deeper neural network architecture. Image-to-image translation methods, such as state-of-the-art HiSD [19] uses a network with small receptive fields that decreases the image resolution only by four times in the encoder. However, we want an organization in a latent space such that we can take meaningful linear directions. Therefore, images should be encoded to a spatially smaller feature space and a network should have a full understanding of an image. For that reason, we set a deep encoder and decoder network architecture but then this network faces the challenges of reconstructing all the details from the input image. To solve this problem, we use a skip connection between the encoder and decoder but only at lower resolution to find the optimal equilibrium of the information flow between with and without dimensionality reduction bottleneck. In summary, our main contributions are:

- We propose VecGAN, a novel image-to-image translation network that is trained end to end with interpretable latent directions. Our framework does not employ a separate style network as in the previous works and translations are achieved with a single deep encoder-decoder architecture.
- VecGAN enables both reference attribute copy and attribute strength manipulation. Reference style encoding is designed in a novel way by using the same encoder from the translation pipeline. First, encoder is used to obtain latent codes of a reference image and it is followed by the projection of the codes into learned latent directions for different attributes.
- We conduct extensive experiments to show the effectiveness of our framework and achieve significant improvements over state-of-the-art for both local and global edits. Qualitative results of our framework can be seen in Fig. 1.

## 2 Related Works

**Image to Image Translation.** Image-to-image translation algorithms aim at preserving a given content while changing targeted attributes. Examples range from translating semantic maps into RGB images [29], to translating summer images into winter images [13], to portrait drawing [36] and very popularly to editing faces [5, 25, 33, 38, 19, 31, 7, 11, 2]. These algorithms powered with GAN loss [8] set an encoder-decoder architecture. In models that learn a deterministic mapping from one domain to the other, images are processed with encoder and decoder to output translated images [29, 24]. In multi-modal image-to-image translation methods, style is encoded separately from an another image or sampled from a distribution [12, 5]. In the generator, style and content are either combined with concatenation [41], or combined with a mask [19] or fed separately through instance normalization blocks [12, 42]. The generator also uses an encoder-decoder architecture [19, 34] that is separate than the style encoder. In

our work, we are interested in designing the attribute as a learnable linear direction in the latent space and we do not employ a separate style encoder which results in a more intuitive framework.

**Learning interpretable latent directions.** In another line of research, it is shown that GANs that are trained to synthesize faces can also be used for face attribute manipulations [15, 3, 16]. Initially, these networks are not designed or trained to translate images but rather to synthesize high fidelity images. However, it is shown that one can embed existing images into the GAN’s embedding space [1] and further one can find latent directions to edit those images [26, 28, 9, 27, 32]. These directions are explored in supervised [26] and unsupervised ways [28, 9, 27, 32]. It is quite remarkable when the generative network is only taught to synthesize realistic images, it organizes the use of latent space such that linear shifts on them change a specific attribute. Inspired by these findings, we design our image to image translation such that a linear shift in the encoded features is expected to change a single attribute of an image. Different than previous works, our framework is trained end-to-end for translation task and allows for reference guided attribute manipulation via projection.

### 3 Method

We follow the hierarchical labels defined by [19]. For a single image, its attribute for tag  $i \in \{1, 2, \dots, N\}$  can be defined as  $j \in \{1, 2, \dots, M_i\}$ , where  $N$  is the number of tags and  $M_i$  is the number of attributes for tag  $i$ . For example  $i$  can be tag of hair color, and attribute  $j$  can take the value of black, brown, or blond.

Our framework has two main objectives. As the main task, we aim to be able to perform the image-to-image translation task in a feature (tag) specific manner. While performing this translation, as the second objective, we also want to obtain an interpretable feature space which allows us to perform tag-specific feature interpolation.

#### 3.1 Generator Architecture

For image to image translation task, we set an encoder-decoder based architecture and latent space translation in the middle as given in Fig. 2. We perform the translation in the encoded latent space,  $e$ , which is obtained by  $e = E(x)$  where  $E$  refers to the encoder. The encoded features go through a transformation  $T$  which is discussed in the next section. The transformed features are then decoded by  $G$  to reconstruct the translated images. The image generation pipeline following feature encoding is described in Eq. 1.

$$\begin{aligned} e' &= T(e, \alpha, i) \\ x' &= G(e') \end{aligned} \tag{1}$$

Previous image-to-image translation networks [19, 34, 5] set a shallow encoder decoder architecture to translate an image and a separate deep network for style

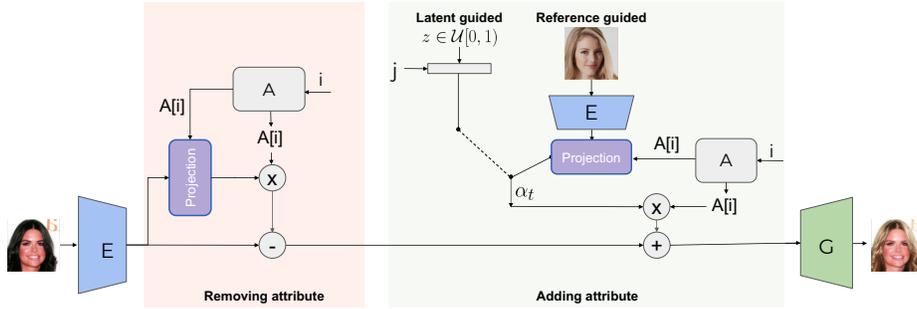


Fig. 2: **VecGAN pipeline.** Our translator is built on the idea of interpretable latent directions. We encode images with an Encoder to a latent representation from which we change a selected tag ( $i$ ), e.g. hair color with a learnable direction  $A_i$  and a scale  $\alpha$ . To calculate the scale, we subtract the target style scale from the source style. This operation corresponds to removing an attribute and adding an attribute. To remove the image’s attribute, source style is encoded and projected from the source image. To add the target attribute, target style scale is sampled from a distribution mapped for the given attribute ( $j$ ), e.g. blonde, brown or encoded and projected from a reference image.

encoding. In most cases, the style encoder includes separate branches for each tag. The shallow architecture that is used to translate images prevents the model from making drastic changes in the images and this helps preserving the identity of the persons. Our framework is different as we do not employ a separate style encoder and instead have a deep encoder-decoder architecture for translation. That is because to be able to organize the latent space in an interpretable way, our framework requires a full understanding of the image and therefore a larger receptive field; deeper network architecture. A deep architecture with decreasing size of feature size, on the other hand, faces the challenges of reconstructing all the fine details from the input image.

With the motivation of helping the network to preserve tag independent features such as the fine details from background, we use skip connections between our encoder and decoder. However, we observe that the flow of information should be limited to force the encoder-decoder architecture learn facial attributes and well-organized latent representations. Because of that reason, we only allow skip connection at low resolution. This design is extensively justified in our Ab-lation Studies.

### 3.2 Translation Module

To achieve a style transformation, we perform the tag-based feature manipulation in a linear fashion in the latent space. First, we set a feature direction matrix  $A$  which contains learnable feature directions for each tag. In our formulation  $A_i$  denotes the learned feature direction for tag  $i$ . Direction matrix  $A$  is randomly initialized and learned during the training process.

Our translation module is formulated in Eq. 2, which adds the desired shift on top of the encoded features  $e$  similar to [28].

$$T(e, \alpha, i) = e + \alpha \times A_i \quad (2)$$

We compute the shift by subtracting target style from the source style as given in Eq 3.

$$\alpha = \alpha_t - \alpha_s \quad (3)$$

Since the attributes are designed as linear steps in the learnable directions, we find the style shift by subtracting the target attribute scale from source attribute scale. This way the same target attribute  $\alpha_t$  can have the same impact on the translated images no matter what the attributes were of the original images. For example, if our target scale corresponds to brown hair, the source scale can be coming from an image with blonde or back hair but since we take a step for difference of the scales, they can be both translated to an image with the same shade of brown hair.

To extract the target shifting scale for feature (tag)  $i$ ,  $\alpha_t$ , there are two alternative pathways. The first pathway, named as latent-guided path, samples a  $z \in \mathcal{U}[0, 1)$  and applies a linear transformation  $\alpha_t = w_{i,j} \cdot z + b_{i,j}$ , where  $\alpha_t$  denotes sampled shifting scale for tag  $i$  and attribute  $j$ . Here tag  $i$  can be hair color and attribute  $j$  can be blonde, brown, or back hair. For each attribute we learn a different transformation module which is denoted as  $M_{i,j}(z)$ . Since we learn a single direction for every tag for example for hair color, this transformation module can put the initially sampled  $z$ 's into correct scale in the linear line based on the target hair color attribute. As the other alternative pathway, we encode the scalar value  $\alpha_t$  in a reference-guided manner. We extract  $\alpha_t$  for tag  $i$  from a provided reference image by first encoding it into the latent space,  $e_r$ , and projecting  $e_r$  via by  $A_i$  as given in Eq. 4.

$$\alpha_t = P(e_r, A_i) = \frac{e_r \cdot A_i}{\|A_i\|} \quad (4)$$

In the reference guidance set-up, we do not use the information of attribute  $j$ , since it is encoded by the tag  $i$  features of the image.

The source scale,  $\alpha_s$ , is obtained by the same way we obtain  $\alpha_t$  from reference image. We perform the projection for the corresponding tag we want to manipulate,  $i$ , by  $P(e, A_i)$ . We formulate our framework with the intuition that the scale controls the amount of feature to be added. Therefore, especially when the attribute is copied over from a reference image, the amount of features that will be added will be different based on the source image. It is for this reason, we find the amount of shift by subtraction as given in Eq. 3. Our framework is intuitive and relies on a single encoder-decoder architecture. Fig. 2 shows the overall pipeline.

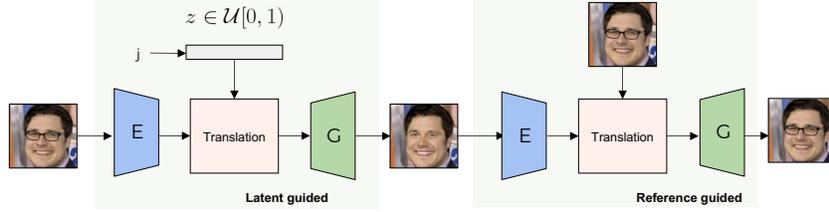


Fig. 3: Overview of cycle translation path.

### 3.3 Training pathways

Modifying the translation paths defined by [19], we train our network using two different paths. For each iteration to optimize our model, we sample a tag  $i$  for shift direction, a source attribute  $j$  as the current attribute and a target attribute  $\hat{j}$ .

**Non-translation path.** To ensure that the encoder-decoder structure preserves details of the images, we perform a reconstruction of the input image without applying any style shifts. The resulting image is denoted as  $x_n$  as given in Eq. 5.

$$x_n = G(E(x)) \quad (5)$$

**Cycle-translation path.** We apply a cyclic translation to ensure that we get a reversible translation from a latent guided scale. In this path, as shown in Fig. 3, we first apply a style shift by sampling  $z \in \mathcal{U}(0, 1)$  and obtaining target  $\alpha_t$  with  $M_{i,\hat{j}}(z)$  for target attribute  $\hat{j}$ . The translation uses  $\alpha$  that is obtained by subtracting  $\alpha_t$  from the source style. Decoder generates an image,  $x_t$  as given in Eq. 6 where  $e$  is encoded features from input image  $x$ ,  $e = E(x)$ .  $x_t$  refers to the image without glasses in Fig. 3.

$$x_t = G(T(e, M_{i,j}(z) - P(e, i), i)) \quad (6)$$

Then by using the original image,  $x$ , as a reference image, we aim to reconstruct the original image by translating  $x_t$ . Overall, this path attempts to reverse a latent-guided style shift with a reference-guided shift. The second translation is given in Eq. 7 where  $e_t = E(x_t)$ .

$$x_c = G(T(e_t, P(e, i) - P(e_t, i), i)) \quad (7)$$

In our learning objectives, we use  $x_n$  and  $x_c$  for reconstruction and  $x_t$  and  $x_c$  for adversarial losses, and  $M_{i,j}(z)$  for the shift reconstruction loss. Details about the learning objectives are given in the next section.

### 3.4 Learning objectives

Given an input image  $x_{i,j} \in \mathcal{X}_{i,j}$ , where  $i$  is the tag to manipulate and  $j$  is the current attribute of the image, we optimize our model with the following objectives. In our equations,  $x_{i,j}$  is shown as  $x$ .

**Adversarial Objective.** During training, our generator performs a style-shift either in a latent-guided way or a reference-guided way, which results in a translated image. In our adversarial loss, we receive feedback from the two steps of cycle-translation path. As the first component of the adversarial loss, we feed a real image  $x$  with tag  $i$  and attribute  $j$  to the discriminator as the real example. To give adversarial feedback to latent-guided path, we use the intermediate image generated in cycle-translation path,  $x_t$ . Finally, to provide adversarial feedback to reference-guided path, we use the final outcome of the cycle-translation path  $x_c$ . Only  $x$  acts as real image, both  $x_t$  and  $x_c$  are translated images, and they are treated as fake images with different attributes. The discriminator aims at classifying whether an image, given its tag and attribute, is real or not. The objective is given in Eq 8.

$$\mathcal{L}_{adv} = 2\log(D_{i,j}(x)) + \log(1 - D_{i,\hat{j}}(x_t)) + \log(1 - D_{i,j}(x_c)) \quad (8)$$

**Shift Reconstruction Objective.** As the cycle-consistency loss performs reference-guided generation followed by latent-guided generation, we utilize a loss function to make these two methods consistent with each other [17, 12, 18, 19]. Specifically, we would like to obtain the same target scale,  $\alpha_t$ , both from the mapping and from the encoded reference image generated by the mapped  $\alpha_t$ . The loss function is given in Eq. 9.

$$\mathcal{L}_{shift} = \|M_{i,j}(z) - P(e_t, i)\|_1 \quad (9)$$

Those parameters,  $M_{i,j}(z)$  and  $P(e_t, i)$ , are calculated for the cycle-translation path as given in Eq. 6 and 7.

**Image Reconstruction Objective.** In all of our training paths, the purpose is to be able to re-generate the original image again. To supervise this desired behavior, we use  $L_1$  loss for reconstruction loss. In our formulation  $x_n$  and  $x_c$  are outputs of non-translation path and cycle-translation path, respectively. Formulation of this objective is provided in Eq. 10.

$$\mathcal{L}_{rec} = \|x_n - x\|_1 + \|x_c - x\|_1 \quad (10)$$

**Orthogonality Objective.** To encourage the orthogonality between directions, we use soft orthogonality regularization based on Frobenius norm, which is given in Eq. 11. This orthogonality further encourages a disentanglement in the learned style directions.

$$\mathcal{L}_{ortho} = \|A^T A - I\|_F \quad (11)$$

**Full Objective.** Combining all of the loss components described, we reach to the overall objective for optimization as given in Eq. 12. We additionally add L1 loss on the matrix  $A$  parameters to encourage its sparsity.

$$\min_{E,G,M,A} \max_D \lambda_a \mathcal{L}_{adv} + \lambda_s \mathcal{L}_{shift} + \lambda_r \mathcal{L}_{rec} + \lambda_o \mathcal{L}_{ortho} + \lambda_{sp} \mathcal{L}_{sparse} \quad (12)$$

To control the dominance of each loss component, we use  $\lambda_a, \lambda_s, \lambda_r, \lambda_o$ , and  $\lambda_{sp}$  hyperparameters. These hyperparameter values and training details are given in Supplementary.

## 4 Experiments

### 4.1 Dataset and Settings

We train our model on CelebA-HQ dataset [22] which contains 30,000 face images. To extensively compare with state-of-the-arts, we follow two training-evaluation protocols as follows:

**Setting A.** In our first setting, we follow the set-up from HiSD [19]. Following HiSD, we use the first 3000 images of CelebA-HQ dataset as the test set and 27000 as the training set. These images include annotations for different attributes from which we use hair color, presence of glass, and bangs attributes for translation task in this setting. Hair color attribute includes 3 tags, black, brown, and blonde whereas the other attributes are binary. The images are resized to  $128 \times 128$ . Following the evaluation protocol proposed by HiSD [19], we compute FID scores on bangs addition task. For each test image without bangs, we translate them to images with bangs with latent and reference guidance. In latent guidance, 5 images are generated for each test image by randomly sampling scale from a uniform distribution. Then this generated set of images are compared with images that have attribute bangs in terms of their FIDs. FIDs are calculated for these 5 sets and averaged. For reference guidance, we randomly pick 5 references images to extract the style scale. FIDs are calculated for these 5 sets separately and averaged.

**Setting B.** In this setting, we follow the set-up from L2M-GAN [34]. The training/test split is obtained by re-indexing each image in CelebA-HQ back to the original CelebA and following the standard split of CelebA. This results in 27,176 training and 2,824 test images. Models are trained for hair color, presence of glasses, bangs, age, smiling, and gender attributes. Images are resized to  $256 \times 256$  resolution. For evaluation, smiling attribute is used following L2M-GAN [34]. It is noted that smiling is one of the most challenging among the CelebA facial attributes because adding/removing a smile requires high-level understanding of the input face image for modifying multiple facial components simultaneously. FIDs are calculated for adding and removing the smile attribute.

Method	Lat.	Ref.	Method	FID (+)	FID (-)	Avg
SDIT [30]	33.73	33.12	StarGAN [4]	32.6	38.6	35.6
StarGANv2 [5]	26.04	25.49	CycleGAN [40]	22.5	24.4	23.5
Elegant [33]	-	22.96	Elegant [33]	39.7	42.9	41.3
HiSD [19]	21.37	21.49	PA-GAN [10]	20.5	21.4	21.0
VecGAN (Ours)	<b>20.17</b>	<b>20.72</b>	InterFaceGAN [26]	24.8	24.9	24.9
			L2M-GAN [34]	17.9	23.3	20.6
			VecGAN (Ours)	<b>17.7</b>	<b>20.3</b>	<b>19.0</b>

(a) Quantitative results for Setting A. Lat: Latent guided, Ref: Reference guided. FID scores are given. Lower is better.

(b) Quantitative results for Setting B. FID (+) (or FID (-)) denotes the FID score for adding (or removing) a smile.

Table 1: Comparisons with state-of-the-art competing methods. Please refer to Section 4 for details on training and evaluation protocol of Setting A and B.

Comparisons	Smiling (+)		Smiling (-)		Smiling (Avg)	
	Quality	Fidelity	Quality	Fidelity	Quality	Fidelity
VecGAN (Ours) vs L2M-GAN	57.96%	70.94%	60.93%	77.50%	59.45%	74.22%
VecGAN (Ours) vs InterFaceGAN	88.13%	91.56%	77.50%	90.62%	82.82%	91.09%

Table 2: User study results conducted with smiling attribute. Smiling (+) denotes the results of adding a smile, Smiling (-) refers to the results of removing a smile, and Smiling (avg) denotes the average of Smiling (+) and Smiling (-). Percentages show the preference rates of our method versus the other competing method.

## 4.2 Results

We extensively compare our results with other competing methods in Table 1. In Setting A, as given in Table 1a, we compare with SDIT [30], StarGANv2 [5], Elegant [33], and HiSD [19] models. Among these methods, HiSD learns a hierarchical style disentanglement whereas StarGANv2 learns a mixed style code. Therefore, StarGANv2 when translating images also does other unnecessary manipulations and does not strictly preserve the identity. Our work is most similar to HiSD as we also learn disentangled style directions. However, HiSD learns feature based local translators which is an approach known to be successful on local edits, e.g. bangs. Ours results show that VecGAN achieves significantly better quantitative results than HiSD both in latent guided and reference guided evaluations even though they are compared on a local edit task.

Fig. 4 shows reference guided results of our model versus HiSD. We compare with HiSD since it provides with the best results after ours. As can be seen from Fig. 4, both methods achieve attribute disentanglement, they do not change any other attribute of the image than the bangs tag. However, HiSD outputs artifacts especially for the reference image from the last column. On the other hand, VecGAN outputs higher quality results. As the second example, we pick a very challenging example to compare these methods. Even though, our results

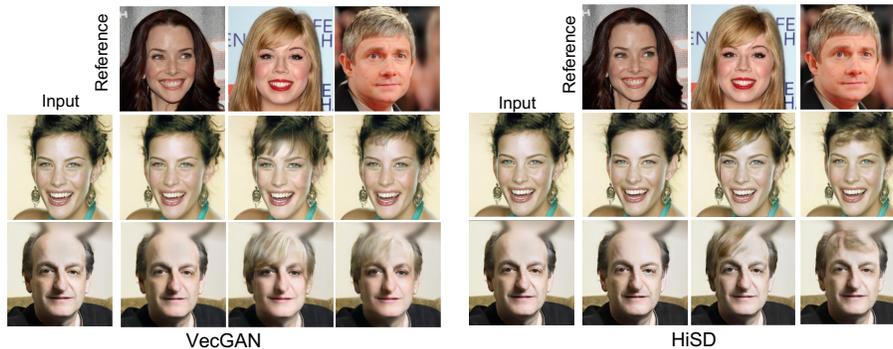


Fig. 4: Qualitative results of bangs attribute of our model (VecGAN) and HiSD. In the second example, we provide a very challenging sample where VecGAN even though not perfect achieves significantly better results than HiSD.

can be further improved to look more realistic, it achieves significantly better outputs with no artifacts compared to HiSD.

In our second set-up of evaluation, we compare our method with many state-of-the-art methods as given in Table 1b. We compare with StarGAN [4], CycleGAN [40], Elegant [33], PA-GAN [10], InterFaceGAN [26], and L2M-GAN [34]. For InterFaceGAN, we use the GAN Inversion [39] as the encoder and pre-trained StyleGAN [15] as the generator backbone. As can be seen from Table 1b, we achieve significantly better scores on both settings and in average.

In our visual comparisons, we mainly focus on L2M-GAN and InterFaceGAN since L2M-GAN is the second best model after ours and InterFaceGAN shares the same intuition with our model and performs edits by latent code manipulation. The results are shown in Fig. 5 where the first four examples show smile addition and the other four examples show smile removal manipulations. The most prominent limitation of L2M-GAN and InterFaceGAN is that they do not preserve the other attributes of images, especially on the background whereas VecGAN does a very good job at that. Smile attribute addition and removal of L2M-GAN is better than InterFaceGAN, however, worse than ours. VecGAN is the only method among them that can produce manipulated images with high fidelity to the originals with only targeted attribute manipulated in a natural and realistic way.

We also conduct a user study on the first 64 images of validation set among 10 users. We set an A/B test and provide users with input images and translations obtained by VecGAN and other competing methods. The left-right order is randomized to ensure fair comparisons. We perform two separate tests. 1) Quality: We ask users to select the best result according to i) whether the smile attribute is correctly added, ii) whether irrelevant facial attributes preserved, iii) and overall whether the output image looks realistic and high quality. 2) Fidelity: We ask users to pay attention if details from the input image is preserved in addition to the quality. When only asked for quality, users pay attention to



Fig. 5: Qualitative results of smile attribute of our model (VecGAN), L2M-GAN, and InterFaceGAN. The first four examples show smile addition and the other four shows smile removal manipulations.

facial attributes and do not pay much attention to the background, ornament, details of hair of the image, and so on. In this test, we remind the users to pay attention to those as well. Table 2 shows the results of the user study. Users preferred our method as opposed to L2M-GAN 59.45% of the time (50% is tie), and as opposed to InterFaceGAN 82.82% of the time for the quality measure in average of smile addition and removal results. When users asked to pay attention to non-facial attributes as well, they preferred our method as opposed to L2M-GAN 74.22% of the time, and as opposed to InterFaceGAN 91.09% of the time in average.

### 4.3 Ablation Study

We conduct ablation studies for network architecture and loss objectives as given in Table 3. We first experiment with a shallower architecture where encoder decreases the input dimension of  $128 \times 128$  to a spatial dimension of  $8 \times 8$ . This version gives reasonable scores, however, we are interested in a better latent space organization. For that, we use a deeper encoder-decoder architecture where encoded latent space goes as low as  $1 \times 1$  which we refer as deep architecture. Deep architecture without skip connections is not able to minimize the reconstruction objective and results in a high FID. On the other hand, deep architecture with a skip connection at each resolution from encoder to decoder can minimize the reconstruction loss however the latent space is not well organized since the model tends to pass all the information from the encoder which instabilizes the training. Our architecture with single skip layer at resolution  $32 \times 32$  provides a good balance between the information flow from encoder-decoder and the latent space bottleneck.



Fig. 6: Qualitative results of ablation study of orthogonality loss. Bangs tag transferred from the reference image.

Method	Lat.	Ref.
Shallow	21.30	20.94
Deep w/o skip	88.62	127.65
Deep all skip	273.80	273.97
Ours	<b>20.17</b>	<b>20.72</b>
w/o Orthogonality	21.98	22.50
w/o Sparsity	24.07	22.43

Table 3: FID results of ablation study with Setting A. Lat: Latent guided, Ref: Reference guided.



Fig. 7: Results of changing the strength of a manipulation gradually. Each example shows a different attribute manipulation. Rows show bangs, hair color, gender, smile, glasses, and age manipulations in this order.

Next, we experiment the effect of loss functions. First, we remove the orthogonality loss of  $A$  directions. This results in worse FID scores but more importantly we observe that the styles are not disentangled, e.g. changing bangs attribute changes the gender as can be seen in Fig. 6. Even without this loss function, we observe that during training the orthogonality loss of  $A$  decreases but to a higher value than when this loss is added to the final objective. That is because the framework and other loss objectives also encourage the disentanglement of attribute manipulations and it shows in the orthogonality of direction vectors. This also shows the importance of orthogonality in style disentanglement and this targeted loss helps improve that significantly. We also observe that sparsity loss applied on the directional vectors stabilizes the training and without that FIDs are much higher.

#### 4.4 Other Capabilities of VecGAN

**Gradually Increased Scale.** We translate images with gradually increased attribute strength as shown in Fig. 7. We plot the manipulation results on six

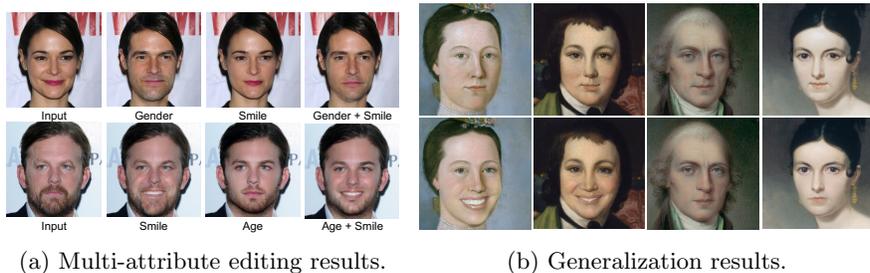


Fig. 8: Results of multi-attribute editing and cross-dataset generalization results of VecGAN.

different attributes. These results show that attributes that are designed as linear transformations are disentangled, and changing one attribute does not affect the other components. In these results, as scales are gradually increased, the strength of the tag smoothly increases with the identity of the person preserved.

**Multi-tag Edits.** We additionally experiment with multi-tag manipulation. To change two attributes, instead of encoding and decoding the image twice with a translation in between each time, we perform two translation operations in the latent code simultaneously. That is we apply Eq. 2 twice for two different  $i$ . Fig. 8a shows results of the multi-tag edits. In the first row, we consider gender and smile tags, and first edit those attributes individually. In the last column, we edit the image with these two tags simultaneously. The second row shows a similar experiment with smile and age tags. We observe that VecGAN provides with disentangled tag control and can successfully edit tags independently.

**Generalization to other domains.** We apply VecGAN model to MetFace dataset [14] without any retraining. The results are provided in Fig. 8b. The first row shows source images, and the second row shows outputs of our model. In the first two examples, we increase the smile attribute, and in the other two, we decrease it. The results show that VecGAN has a good generalization ability and works reasonably well across datasets.

## 5 Conclusion

This paper introduces VecGAN, an image-to-image translation framework with interpretable latent directions. This framework includes a deep encoder and decoder architecture with latent space manipulation in between. Latent space manipulation is designed as vector arithmetic where for each attribute, a linear direction is learned. This design is encouraged by the finding that well-trained generative models organize their latent space as disentangled representations with meaningful directions in a completely unsupervised way. Each change in the architecture and loss functions is extensively studied and compared with state-of-the-arts. Experiments show the effectiveness of our framework.

## References

1. Abdal, R., Qin, Y., Wonka, P.: Image2stylegan: How to embed images into the stylegan latent space? In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4432–4441 (2019)
2. Abdal, R., Zhu, P., Mitra, N.J., Wonka, P.: Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows. *ACM Transactions on Graphics (ToG)* **40**(3), 1–21 (2021)
3. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. arXiv preprint arXiv:1809.11096 (2018)
4. Choi, Y., Choi, M., Kim, M., Ha, J.W., Kim, S., Choo, J.: Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8789–8797 (2018)
5. Choi, Y., Uh, Y., Yoo, J., Ha, J.W.: Stargan v2: Diverse image synthesis for multiple domains. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2020)
6. Dundar, A., Sapra, K., Liu, G., Tao, A., Catanzaro, B.: Panoptic-based image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8070–8079 (2020)
7. Gao, Y., Wei, F., Bao, J., Gu, S., Chen, D., Wen, F., Lian, Z.: High-fidelity and arbitrary face editing. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16115–16124 (2021)
8. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. *Advances in neural information processing systems* **27** (2014)
9. Härkönen, E., Hertzmann, A., Lehtinen, J., Paris, S.: Ganspace: Discovering interpretable gan controls. *Advances in Neural Information Processing Systems* **33** (2020)
10. He, Z., Kan, M., Zhang, J., Shan, S.: Pa-gan: Progressive attention generative adversarial network for facial attribute editing. arXiv preprint arXiv:2007.05892 (2020)
11. Hou, X., Zhang, X., Liang, H., Shen, L., Lai, Z., Wan, J.: Guidedstyle: Attribute knowledge guided style manipulation for semantic face editing. *Neural Networks* **145**, 209–220 (2022)
12. Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. *Eur. Conf. Comput. Vis.* (2018)
13. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2017)
14. Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., Aila, T.: Training generative adversarial networks with limited data. *Advances in Neural Information Processing Systems* **33**, 12104–12114 (2020)
15. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4401–4410 (2019)
16. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8110–8119 (2020)
17. Lee, H.Y., Tseng, H.Y., Huang, J.B., Singh, M., Yang, M.H.: Diverse image-to-image translation via disentangled representations. In: Proceedings of the European conference on computer vision (ECCV). pp. 35–51 (2018)

18. Li, X., Hu, J., Zhang, S., Hong, X., Ye, Q., Wu, C., Ji, R.: Attribute guided unpaired image-to-image translation with semi-supervised learning. arXiv preprint arXiv:1904.12428 (2019)
19. Li, X., Zhang, S., Hu, J., Cao, L., Hong, X., Mao, X., Huang, F., Wu, Y., Ji, R.: Image-to-image translation via hierarchical style disentanglement. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8639–8648 (2021)
20. Liu, M.Y., Breuel, T., Kautz, J.: Unsupervised image-to-image translation networks. In: Adv. Neural Inform. Process. Syst. (2017)
21. Liu, M.Y., Tuzel, O.: Coupled generative adversarial networks. *Advances in neural information processing systems* **29**, 469–477 (2016)
22. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: Proceedings of International Conference on Computer Vision (ICCV) (December 2015)
23. Mardani, M., Liu, G., Dundar, A., Liu, S., Tao, A., Catanzaro, B.: Neural ffts for universal texture image synthesis. *Advances in Neural Information Processing Systems* **33**, 14081–14092 (2020)
24. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2337–2346 (2019)
25. Shen, W., Liu, R.: Learning residual images for face attribute manipulation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4030–4038 (2017)
26. Shen, Y., Gu, J., Tang, X., Zhou, B.: Interpreting the latent space of gans for semantic face editing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9243–9252 (2020)
27. Shen, Y., Zhou, B.: Closed-form factorization of latent semantics in gans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1532–1540 (2021)
28. Voynov, A., Babenko, A.: Unsupervised discovery of interpretable directions in the gan latent space. In: International Conference on Machine Learning. pp. 9786–9796. PMLR (2020)
29. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional gans. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8798–8807 (2018)
30. Wang, Y., Gonzalez-Garcia, A., van de Weijer, J., Herranz, L.: Sdit: Scalable and diverse cross-domain image translation. In: Proceedings of the 27th ACM International Conference on Multimedia. pp. 1267–1276 (2019)
31. Wu, P.W., Lin, Y.J., Chang, C.H., Chang, E.Y., Liao, S.W.: Relgan: Multi-domain image-to-image translation via relative attributes. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 5914–5922 (2019)
32. Wu, Z., Lischinski, D., Shechtman, E.: Stylespace analysis: Disentangled controls for stylegan image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12863–12872 (2021)
33. Xiao, T., Hong, J., Ma, J.: Elegant: Exchanging latent encodings with gan for transferring multiple face attributes. In: Proceedings of the European conference on computer vision (ECCV). pp. 168–184 (2018)
34. Yang, G., Fei, N., Ding, M., Liu, G., Lu, Z., Xiang, T.: L2m-gan: Learning to manipulate latent space semantics for facial attribute editing. In: Proceedings of the

- IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2951–2960 (2021)
35. Yi, R., Liu, Y.J., Lai, Y.K., Rosin, P.L.: Apdrawinggan: Generating artistic portrait drawings from face photos with hierarchical gans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10743–10752 (2019)
  36. Yi, R., Liu, Y.J., Lai, Y.K., Rosin, P.L.: Unpaired portrait drawing generation via asymmetric cycle mapping. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8217–8225 (2020)
  37. Yi, Z., Zhang, H., Tan, P., Gong, M.: Dualgan: Unsupervised dual learning for image-to-image translation. In: Int. Conf. Comput. Vis. (2017)
  38. Zhang, G., Kan, M., Shan, S., Chen, X.: Generative adversarial network with spatial attention for face attribute editing. In: Proceedings of the European conference on computer vision (ECCV). pp. 417–432 (2018)
  39. Zhu, J., Shen, Y., Zhao, D., Zhou, B.: In-domain gan inversion for real image editing. In: European conference on computer vision. pp. 592–608. Springer (2020)
  40. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Int. Conf. Comput. Vis. (2017)
  41. Zhu, J.Y., Zhang, R., Pathak, D., Darrell, T., Efros, A.A., Wang, O., Shechtman, E.: Multimodal image-to-image translation by enforcing bi-cycle consistency. In: Advances in neural information processing systems. pp. 465–476 (2017)
  42. Zhu, P., Abdal, R., Qin, Y., Wonka, P.: Sean: Image synthesis with semantic region-adaptive normalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5104–5113 (2020)