

Deep Curvilinear Editing: Commutative and Nonlinear Image Manipulation for Pretrained Deep Generative Model

Takehiro Aoshima, Takashi Matsubara
Osaka University

1-3 Machikaneyama, Toyonaka, Osaka, 560-8531 Japan.

aoshima@hopf.sys.es.osaka-u.ac.jp, matsubara@sys.es.osaka-u.ac.jp

Abstract

Semantic editing of images is the fundamental goal of computer vision. Although deep learning methods, such as generative adversarial networks (GANs), are capable of producing high-quality images, they often do not have an inherent way of editing generated images semantically. Recent studies have investigated a way of manipulating the latent variable to determine the images to be generated. However, methods that assume linear semantic arithmetic have certain limitations in terms of the quality of image editing, whereas methods that discover nonlinear semantic pathways provide non-commutative editing, which is inconsistent when applied in different orders. This study proposes a novel method called deep curvilinear editing (DeCurEd) to determine semantic commuting vector fields on the latent space. We theoretically demonstrate that owing to commutativity, the editing of multiple attributes depends only on the quantities and not on the order. Furthermore, we experimentally demonstrate that compared to previous methods, the nonlinear and commutative nature of DeCurEd facilitates the disentanglement of image attributes and provides higher-quality editing.

1. Introduction

Generating and editing realistic images is one of the fundamental goals in computer vision. Generative adversarial networks (GANs) [20] have emerged as a major image generation approach owing to the quality of generated images [7, 33–36, 48] and provide various real-world applications [1, 13, 18, 53, 63, 68]. Other notable methods include variational autoencoders (VAEs) [24], conditional PixelCNN [61], and diffusion-based models [26, 56], which are collectively called deep generative models. However, deep generative models cannot inherently edit images semantically. They can be viewed as mappings from latent space to image space, and the latent variables determine the

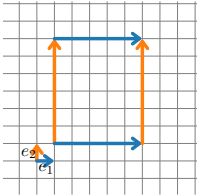
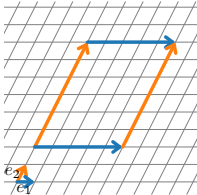
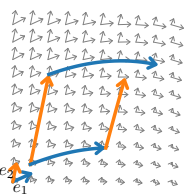
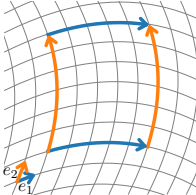
generated images. Therefore, several methods have been developed to train deep generative models such that each semantic attribute the user wants to edit is assigned to each element of the latent variable [10, 45, 46] (see the first column of Table 1). However, this approach requires computationally expensive training and can conflict with the quality of image generation. Other studies have developed image-to-image translation that translates images from one domain to another [28, 64, 68]. However, this approach also requires training from scratch and limits image editing to be discontinuous unless combined with latent variable manipulation. Therefore, a general and promising approach is necessary to identify manipulations on latent variables of already trained models that edit images semantically.

A study reported that adding certain vectors to latent variables can modify the corresponding attributes of the object in the generated images [51]. This indicates that the latent space can be regarded as a linear vector space. Some studies have aimed to identify attribute vectors in a supervised or an unsupervised manner [17, 19, 22, 27, 49, 50, 54, 55, 62, 69]. In any case, these studies introduce the strong assumption of linear semantic arithmetic on the latent space (see the second column of Table 1), which limits the quality of image editing. Other studies have proposed methods to determine attribute vectors depending on the position in the latent space, that is, the attribute vector fields (or local attribute coordinates). [2, 11, 12, 17, 29, 37, 44, 52, 58, 60]; these methods edit an image attribute by moving the latent variable nonlinearly along the corresponding attribute vector field. Although this approach is elegant, the edits of different attributes are generally non-commutative (see the third column of Table 1). That is, what we get is different when we edit attributes (denoted by e_1 and e_2) one after another, or in the reverse order. This property can harm the disentanglement between attributes considering that the relationships among attributes change at different points. In contrast, linear arithmetic ensures that the edits of different attributes are commutative.

To overcome this dilemma, this study proposes *deep*

Table 1. Comparison of Our Proposal against Related Methods.

	Training under constraints [10, 45, 46]	Linear arithmetic [27, 55, 62]	Vector fields/Local basis [12, 52, 60]	DeCurvEd (proposed)
Global coordinate	Cartesian	oblique	(only local)	curvilinear
No retraining	✗	✓	✓	✓
Nonlinear edit	–	✗	✓	✓
Commutative edit	✓	✓	✗	✓

Conceptual diagram				
--------------------	---	---	--	---

curvilinear editing (*DeCurvEd*), a method that determines a set of commuting attribute vector fields in the latent space of a pre-trained deep generative model. The key idea is adopted from the theorem that a set of vector fields is locally expressed as partial derivatives of a coordinate chart if it is linearly independent and commuting [43]. Therefore, we define a curvilinear coordinate system globally [3] by employing a normalizing flow [9, 39], from which we derive commuting vector fields (see the rightmost panel of Table 1). The advantages of *DeCurvEd* are as follows (see also Table 1):

1. Edits of different attributes are always commutative, unlike previous methods that assume attribute vector fields (e.g., [60]). Therefore, an edited image does not depend on the order of editing, but on the amount of editing performed.
2. Edits are nonlinear, which indicates that *DeCurvEd* provides better editing quality than methods that assume linear arithmetic of attribute vectors (e.g., [62]).
3. *DeCurvEd* does not require retraining of deep generative models, but identifies attribute vector fields in the latent space of pre-trained models, unlike image-to-image translation and training under constraints (e.g., [10, 28]).
4. We propose *CurvilinearGANSpace* by combining *DeCurvEd* with GANs and experimentally demonstrate that the nonlinear and commutative nature disentangles attributes and enables high-quality editing.
5. The key idea is not limited to GANs, and is available for any generative models that are conditioned on latent variables, including VAEs [24], conditional PixelCNN [61], and diffusion-based models [26, 56].

2. Related Work

Image Editing by Deep Generative Models Most generative models define a mapping that maps a latent variable

to a data sample (that is, an image in this study). Previous studies on deep learning-based generative models have confirmed that manipulating the latent variable can determine the image to be generated [20, 24]. Radford et al. [51] identified a semantically meaningful vector in the latent space based on the difference between two groups. Then, an attribute can be imposed on an image by adding the attribute vector to the latent variable of the image. This discovery attracted wide attention and spurred research into semantic image editing by manipulating the latent variables.

Several studies have developed methods for training models under constraints to easily determine attribute vectors, rather than identifying attribute vectors after training [10, 25, 45, 46]. These methods often make each element of the latent variable as independent as possible [25]. Then, one element is assigned to a group (an attribute) that changes collectively and is independent of other groups in the image [6]. These methods can be viewed as introducing a Cartesian coordinate system to the latent space and assigning one attribute to each axis. However, constraints often conflict with other training criteria (such as likelihood and Gaussian prior) and result in models with inferior quality and diversity. Additionally, the generated models needed to be trained from scratch, which incurs computation costs. See the column “Training under constraints” in Table 1.

Moreover, some studies have developed image-to-image transformation, which maps an image from one domain to another rather than manipulating a latent variable [28, 64, 68]. This approach limits image editing only between domains or needs to be combined with a latent space.

Discovering Linear Attribute Arithmetic Several studies have investigated linear manipulations on latent variables in the already trained deep generative models [27, 62].

SeFa and related methods attempted to find semantic directions by analyzing the weight parameters [55, 66, 67]. Voynov and Babenko [62] proposed an unsupervised framework that learns a set of semantic directions. According

to this framework, changing a latent variable along a semantic direction will edit one attribute of the corresponding image, and the degree of change in the attribute will be proportional to the amount of change in the latent variable. GANSpace [27] applied a principal component analysis (PCA) to extract a lower-dimensional subspace of the latent space, assuming each principal component corresponds to an attribute. These methods assume linear arithmetic of the attribute vectors; that is, they introduce an oblique coordinate system to the latent space. See the column “Linear arithmetic” in Table 1.

Because the distribution of real-world data is often biased and skewed, it is unlikely that the latent space is flat and homogeneous. Khulkov et al. [37] found that different directions correspond to the same attribute at different locations in the latent space. Therefore, the above methods are limited in terms of image editing quality.

Discovering Semantic Vector Fields The direction corresponding to an attribute varies depending on the location in the latent space, thereby indicating that a set of directions corresponding to attributes forms a set of vector fields, rather than linear arithmetic. If so, one can edit an attribute of an image by moving the latent variable nonlinearly along the vector field corresponding to the attribute instead of adding an attribute vector. Tzelepis et al. [60] proposed WarpedGANSpace, which learns a set of vector fields, each of which is defined as a gradient flow of an RBF function on the latent space. Choi et al. [12] learned a local basis at every point of the latent space such that each element of the local basis corresponds to an attribute. StyleFlow [2] and SSFlow [44] used normalizing flows to define a local coordinate system. On an N -dimensional manifold, a local basis, local coordinate system, and a set of N linearly independent vector fields are compatible; such vector fields are called coordinate vector fields (see Example 8.2, [43]). However, because the coordinate system in the above studies is defined only locally, multiple edits may be inconsistent globally. We will demonstrate this in the following section. See the column “Vector fields/Local basis” in Table 1.

Some studies have attempted to define a (Riemannian) metric on the latent space [4, 5, 8]. These methods successfully interpolate between two images by nonlinearly moving latent variables along the geodesic; however, they are insufficient for attribute editing. Some others attempted complex and dynamic editing specified by text rather than attributes [59]; nevertheless, such methods cannot be directly compared to ours.

3. Theoretical Background

We introduce the theoretical background of the proposed and related methods introduced in Section 2. Theorems in

this paper are basic knowledge about manifolds; readers unfamiliar with this topic are referred to the reference [43]. Remarks are our findings.

Let \mathcal{X} and \mathcal{Z} denote an image space and a latent space of a deep generative model, respectively. The generator G (also called decoder) of the deep generative model is a mapping from the latent space \mathcal{Z} to the image space \mathcal{X} ; given a latent variable $z \in \mathcal{Z}$, the generator produces an image $x \in \mathcal{X}$ as $x = G(z)$. We assume the latent space \mathcal{Z} to be an N -dimensional space diffeomorphic to a Euclidean space.

Let $\{z^i\}_{i=1}^N$ denote the coordinate system (i.e., the basis) on a neighborhood of the point $z \in \mathcal{Z}$. Let Z_k denote a vector field on the latent space \mathcal{Z} indexed by k , that is, $Z_k : \mathcal{Z} \rightarrow \mathcal{T}_z\mathcal{Z}$, where $\mathcal{T}_z\mathcal{Z}$ is the tangent space (i.e., the space of tangent vectors or velocities) of the latent space \mathcal{Z} at point z . At point z , the coordinate system on tangent space $\mathcal{T}_z\mathcal{Z}$ is denoted by $\{\frac{\partial}{\partial z^i}\}_{i=1}^N$, and a vector field Z_k is expressed as $Z_k(z) = \sum_{i=1}^N Z_k^i(z) \frac{\partial}{\partial z^i}$ for smooth functions $Z_k^i : \mathcal{Z} \rightarrow \mathbb{R}$. A method that assumes attribute vector fields [12, 52, 60] edits an attribute k of an image x by integrating a latent variable z along the vector field Z_k that corresponds to attribute k ; the edited image is $x' = G(z')$ for

$$z' = z + \int_0^t Z_k(z(\tau)) d\tau, \quad (1)$$

where $z(0) = z$, and $t \in \mathbb{R}$ denotes the change amount of attribute k . t may be positive or negative. We rewrite the above equation using a flow, denoted by $\phi_k^t : \mathcal{Z} \rightarrow \mathcal{Z}$ for t as:

$$z' = \phi_k^t(z). \quad (2)$$

Then, we define the commutativity of editing as follows:

Definition 1 (Commutativity). *Edits of attributes k and l are commutative if and only if the corresponding flows ϕ_k^t and ϕ_l^s are commuting, that is, it holds that $\phi_l^s \circ \phi_k^t = \phi_k^t \circ \phi_l^s$ for any $s, t \in \mathbb{R}$ at any point z on the latent space \mathcal{Z} .*

Intuitively, making a person smile and then wear sunglasses results in the same image as making the person wear sunglasses and then smile if the vector fields corresponding to smiling and wearing sunglasses are commuting. Else, edits in different orders produce different images.

Remark 1. *A method that assumes linear attribute arithmetic (e.g., [27, 55, 62]) is a special case of a method that assumes attribute vector fields, and its edits are commutative.*

See Appendix A for formal proofs of any remarks in this manuscript. Therefore, we can discuss a method that assumes linear attribute arithmetic in the same context. We introduce the following theorem.

Theorem 1 (Commuting Vector Fields, Theorem 9.44, [43]). *Two flows ϕ_k and ϕ_l are commuting if and only if the underlying vector fields Z_k and Z_l are commuting.*




Figure 1. The concept diagram. The left part shows how DeCurvEd edits an attribute. The right part shows its combination with GANs called CurvilinearGANSpace.

This theorem suggests the following remark:

Remark 2. In general, edits by a method that assumes attribute vector fields (e.g., [12, 52, 60]) are non-commutative.

In addition, we introduce the following theorem.

Theorem 2 (Canonical Forms for Commuting Vector Fields, Theorem 9.46, [43]). Let vector fields Z_1, Z_2, \dots, Z_N on an N -dimensional space \mathcal{Z} be linearly independent and commuting on an open set $\mathcal{U} \subset \mathcal{Z}$. At each $z \in \mathcal{U}$, there exists a smooth coordinate chart $\{\frac{\partial}{\partial s^i}\}_{i=1}^N$ centered at z such that $\frac{\partial}{\partial s^i} = Z_i$.

Furthermore, given a smooth coordinate chart $\{\frac{\partial}{\partial s^i}\}_{i=1}^N$, vector fields $Z_i = \frac{\partial}{\partial s^i}$ are commuting. A coordinate chart is a nonlinear bijective mapping to Euclidean space. Therefore, intuitively, a set of N linearly independent and commuting vector fields on an N -dimensional space \mathcal{Z} is equivalent to a set of N vector fields along the axes of a coordinate system up to geometric transformation.

4. Method

4.1. DeCurvEd

Given the theoretical background, we propose DeCurvEd, as shown in Fig. 1. Intuitively, we consider the case where the open set \mathcal{U} in Theorem 2 is not a proper subset but equal to the latent space \mathcal{Z} .

We prepare an N -dimensional Euclidean space \mathcal{V} and call it the Cartesianized latent space, whose coordinate system $\{v^i\}_{i=1}^N$ is a global Cartesian coordinate system. Let e_k denote the k -th element of the standard basis of the tangent space, that is, $e_k := \frac{\partial}{\partial v^k}$. Then, the vector field \tilde{Z}_k corresponding to attribute k is defined as

$$\tilde{Z}_k := e_k. \quad (3)$$

As discussed in the previous section, vector fields \tilde{Z}_k and \tilde{Z}_l defined in this way are commuting for any k and l . The flow $\psi_k^t : \mathcal{V} \rightarrow \mathcal{V}$ that arises from the vector field \tilde{Z}_k is given by

$$\psi_k^t(v) := v + \int_0^t e_k d\tau = v + t e_k. \quad (4)$$

The flows ψ_k are commuting because $(\psi_l^s \circ \psi_k^t)(v) = v + t e_k + s e_l = (\psi_k^t \circ \psi_l^s)(v)$. We introduce a smooth bijective mapping $f : \mathcal{Z} \rightarrow \mathcal{V}, z \mapsto v$ that corresponds to the coordinate chart in Theorem 2. The mapping f can be implemented using a normalizing flow; however, it is not limited to [9, 21, 39]. We define a flow ϕ_k that edits attribute k on the latent space \mathcal{Z} as:

$$\phi_k^t := f^{-1} \circ \psi_k^t \circ f. \quad (5)$$

See also the left half of Fig. 1. We redefine the edit as Algorithm 1 in Appendix B.

Subsequently, one can generate an edited image $x' = G(z')$ using generator G . Deep generative models such as GANs do not have an inherent way of inferring a latent variable z from an image x ; this is outside the scope of this study. Interested readers can refer to this survey [65].

4.2. Theoretical Analysis

The pushforward f_* is a mapping naturally induced by the mapping f , which maps a tangent vector (or a basis) on the latent space \mathcal{Z} to that on the Cartesianized latent space \mathcal{V} . Also, the pushforward $(f^{-1})_*$ maps the Cartesian coordinate system on the Cartesianized latent space \mathcal{V} and implicitly defines a coordinate system on the latent space \mathcal{Z} [43]. A coordinate system defined by a bijective transformation of a Cartesian coordinate is called a curvilinear coordinate [3]. Therefore, we name this method deep curvilinear editing (DeCurvEd). Because the mapping f is defined globally between spaces \mathcal{Z} and \mathcal{V} , the curvilinear coordinate system is also defined globally. The pushforward $(f^{-1})_*$ can define commuting vector fields on \mathcal{Z} by push-forwarding the coordinate vector fields on \mathcal{V} . Therefore, we make the following remarks.

Remark 3. Using DeCurvEd, any edits of attributes in the latent space \mathcal{Z} can be nonlinear and commutative.

Remark 4. DeCurvEd can define vector fields on the latent space \mathcal{Z} and is a special case of a method that assumes attribute vector fields (e.g., [12, 52, 60]).

Remark 5. A method that assumes linear attribute arithmetic (e.g., [27, 55, 62]) is a special case of DeCurvEd, with a linear mapping f .

Therefore, DeCurvEd enjoys the advantages of both attribute arithmetic and vector fields. All theories and remarks are not dependent on the properties of particular models. Thus, we make the following remark.

Remark 6. DeCurvEd offers attribute editing for any generative models conditioned on latent variables, including GANs [20], VAEs (see Fig. 4 of [24]), conditional Pixel-CNN [61], and diffusion-based models (see Fig. 8 of [26] and Fig. 4 of [56]).

4.3. CurvilinearGANSpace

Attribute editing by DeCurvEd is available for any deep generative models and for both supervised and unsupervised learning. This study adopted the unsupervised training framework for GANs proposed by Voynov and Babenko [62], as shown in the right half of Fig. 1. Following previous studies, we call it *CurvilinearGANSpace*.

Given a latent variable z , CurvilinearGANSpace randomly edits index k by ϵ and produces an edited one z' . In some cases, only the first N' indices of all N indices are candidates for editing. We prepare a neural network called reconstructor R , which accepts the pair of generated images $x = G(z)$ and $x' = G(z')$ and regresses the edited index k and the change amount ϵ . In particular, one output is an N' -dimensional vector \hat{k} to regress the edited index k : the loss function is the classification error $L_{\text{cls}}(k, \hat{k})$, which is defined as the cross-entropy. As the mapping f minimizes this error, image editing of index k becomes easier for the reconstructor R to distinguish from image editing of other indices $l \neq k$, thereby assigning one attribute to each vector field and facilitating the disentanglement between attributes. The other output $\hat{\epsilon}$ is a scalar regressing the change amount ϵ ; the loss function is the regression error $L_{\text{reg}}(\epsilon, \hat{\epsilon})$ defined as the absolute error. As this error is minimized, the change in the latent variable continuously matches the semantic change in the image.

Additionally, we introduce a regularization term L_{nl} to be minimized for the mapping f ;

$$L_{\text{nl}}(z) = (\log \det |\frac{\partial f}{\partial z}|)^2. \quad (6)$$

The Jacobian determinant $\det \frac{\partial f}{\partial z}$ of the mapping f indicates the extent to which the latent space \mathcal{Z} is stretched by the mapping f ; when it is 1.0, the mapping f is isometric. Subsequently, this term L_{nl} avoids extreme deformation of the latent space \mathcal{Z} by the mapping f . The final objective function is defined as:

$$\mathbb{E}_{z, k, \epsilon} \left[L_{\text{cls}}(k, \hat{k}) + \lambda L_{\text{reg}}(\epsilon, \hat{\epsilon}) + \alpha L_{\text{nl}}(z) \right], \quad (7)$$

where $\lambda, \alpha \in \mathbb{R}$ are hyperparameters weighing objectives. See also Algorithm 2 in Appendix B for more details.

Table 2. Datasets and Settings.

Dataset	GANs	Reconstructor	N	N'
MNIST [42]	SNGAN [48]	LeNet [41]	128	128
AnimeFaces [31]	SNGAN [48]	LeNet [41]	128	128
ILSVRC [14]	BigGAN [7]	ResNet-18 [23]	120	120
CelebA-HQ [47]	ProgGAN [33]	ResNet-18 [23]	512	200
CelebA-HQ [47]	StyleGAN2 [36]	ResNet-18 [23]	512	200
LSUN Car [40]	StyleGAN2 [36]	ResNet-18 [23]	512	200

5. Experiments

5.1. Experimental Setting

Datasets, Backbones, and Comparison Methods We examined CurvilinearGANSpace and related methods using combinations of datasets, GANs, and reconstructors, as summarized in Table 2. N denotes the number of dimensions of the latent space \mathcal{Z} , and N' denotes the number of dimensions used for training. For StyleGAN2, W space was used as the latent space. For ILSVRC and CelebA-HQ, we used pre-trained models from their official repositories. These experimental settings are identical to those in previous studies [60, 62]. See Appendix C.1 and the references [60, 62] for more details.

For comparison, we used a method that assumes linear arithmetic [62] and a method that assumes attribute vector fields called WarpedGANSpace [60]. To clarify the difference, we hereafter refer to the former method as LinearGANSpace. We used their pre-trained models for all but the LSUN Car dataset and used our own trained models for the LSUN Car dataset, each trained in the same framework. ^{1 2}

Architectures and Hyperparameters As the bijective mapping f , we used a continuous normalizing flow with six concatsquash layers [21]. We set the number of hidden units equal to the input dimension and used hyperbolic tangent function as its activation function. See Appendix C.2 for more introduction. We used Adam optimizer [38] with a constant learning rate of 10^{-4} . We used $\lambda = 0.25$, which is equivalent to that used by previous studies [60, 62]. For simplicity, we used $\alpha = 1$.

Evaluation Metrics For CelebA-HQ, we measured the attribute scores of generated images using separate pre-trained attribute predictors $A_k : \mathcal{X} \rightarrow \mathbb{R}$.² FairFace measured age, gender, and race (skin color) attributes [32], and CelebA-HQ attributes classifier measured smile, beard, and bangs attributes [30] from 0 to 1; Hopenet measured face

¹<https://github.com/anvoynov/GANLatentDiscovery> for LinearGANSpace.

²<https://github.com/chi0tztz/WarpedGANSpace> for WarpedGANSpace and attribute predictors.




Figure 2. Results of sequential editing of attributes from left to right. (left) CelebA-HQ+ProgGAN. The change amount was set to $\tilde{t} = 0.3$ for smile, bangs, and pitch, and $\tilde{t} = 1.0$ for yaw. (right) AnimeFaces+SNGAN. Each row shows the results of LinearGANSpace, WarpedGANSpace, and CurvilinearGANSpace, from top to bottom. The signs + and - denote the addition and the subtraction of the corresponding attributes, respectively. O: original, S: “smile”, B: “bangs”, P: “pitch”, Y: “yaw”. C: “hair color”, L: “hair length”.

Table 3. Commutativity Errors [%] of StyleGAN2.

	A+G	R+P	B+Y
LinearGANSpace [62]	0.01 / 0.05	0.02 / 0.07	0.02 / 0.15
WarpedGANSpace [60]	11.40 / 6.62	3.15 / 3.46	1.28 / 2.22
CurvilinearGANSpace (ours)	<u>0.07 / 0.35</u>	<u>0.05 / 0.62</u>	<u>0.08 / 0.55</u>

A: “age”, G: “gender”, R: “race” B: “bangs”, P: “pitch”, Y: “yaw”.

directions, yaw and pitch, in degree [16]. Additionally, ArcFace measured the identity score $I(\cdot, \cdot) : \mathcal{X} \times \mathcal{X} \rightarrow [0, 1]$ to evaluate whether two images are of the same person [15]. We selected index k as the one corresponding to that attribute if the measured attribute score has the largest covariance with the change amounts of index k . The above procedure is identical to that of the previous study [60].

The amount t by which attribute k of latent variable z is edited differs from the amount by which the corresponding attribute score A_k of the generated image $x = G(x)$ is changed. For a fair comparison, we normalized the change amount t such that the measured attribute score changes by 5 degrees for the pitch and yaw attributes and 0.1 for others, and denoted the amount by $\tilde{t} = 0.1$.

After index identification and normalization, we used several evaluation metrics. We defined *commutativity error* of attributes $k+l$ to evaluate how commutative the image editing is by measuring the difference in the attribute score between images with edits of attributes k and l applied in different orders. We defined *side effect error* to evaluate the disentanglement between attributes by measuring how much an edit of the target attribute k changes the other attributes $l \neq k$ as undesired side effects. We also defined *identity error* to evaluate the disentanglement by measuring how much an edit of the target attribute k reduces the identity score. The errors in edits of multiple attributes are defined similarly.

Owing to the availability of attribute predictors, these evaluations were performed only for CelebA-HQ. See Appendix C.3 for the detailed procedures and definition. For other datasets, we manually selected the index k , following previous studies [60, 62].

5.2. Experimental Results

Commutativity of Editing Table 3 shows the commutativity errors for CelebA-HQ+StyleGAN2 with $\tilde{t} = 0.1$. Those of LinearGANSpace and CurvilinearGANSpace were always less than 0.7 %; even though they were not exactly zero due to numerical and rounding errors, the errors were negligible. The errors of WarpedGANSpace were between 1.2 % and 11.4 %. Therefore, as expected, the image editing by WarpedGANSpace is non-commutative, whereas that by LinearGANSpace and CurvilinearGANSpace is commutative.

We edited image attributes sequentially so that the total amount of change is zero and summarized the results in Fig. 2. The images generated after sequential editing by LinearGANSpace or CurvilinearGANSpace look identical to the originals, which indicates that their image editing is commutative. When editing a human face by WarpedGANSpace, the face’s yaw rotation and image brightness were not restored. Also for an AnimeFaces image, the hair color was not restored. These results indicate that image editing by WarpedGANSpace is non-commutative.

A closer look at each edit reveals that the editing of a human face by LinearGANSpace does not properly edit the smile attribute, and the edit of yaw rotation changes the hairstyle as well. When WarpedGANSpace edits the pitch or yaw rotation of the human face, it changes hair color, skin color, and brightness as well. For the AnimeFaces image, the editing by LinearGANSpace is of inferior quality. WarpedGANSpace’s edits of the hair color and hair length change the face (i.e., identity). CurvilinearGANSpace’s editing is of excellent quality without severe side effects.

These results indicate that CurvilinearGANSpace provides commutative editing and significantly improves the disentanglement between attributes. The training framework used [62] leads the editing methods to learn disentanglement between attributes by classifying indices. WarpedGANSpace takes advantage of nonlinearity to allow better editing; however, there is no mechanism to further facilitate disentanglement. CurvilinearGANSpace assumes

Table 4. Side Effect Errors [%] of StyleGAN2.

	Target k	Side Effect Errors l [%]					
		A	G	R	B	P	Y
LinearGANSpace [62]	A	100	59	37	63	41	61
	G	28	100	16	78	20	17
	R	61	52	100	71	24	19
	B	175	172	78	100	70	64
	P	71	90	43	76	100	57
	Y	58	55	43	94	36	100
WarpedGANSpace [60]	A	100	51	63	111	59	23
	G	75	100	94	124	236	57
	R	63	64	100	131	73	25
	B	23	27	22	100	15	21
	P	41	44	30	80	100	41
	Y	30	30	22	97	23	100
CurvilinearGANSpace (ours)	A	100	80	45	137	60	37
	G	62	100	50	84	61	40
	R	65	56	100	60	37	23
	B	40	38	15	100	14	19
	P	60	52	36	76	100	44
	Y	41	62	21	79	21	100

A : “age”, G : “gender”, R : “race” B : “bangs”, P : “pitch” Y : “yaw”.
Severe side effects (more than 90 %) are highlighted in bold red.

Table 5. Identity Errors [%] of StyleGAN2.

	A	G	R	B	P	Y	Avg.
LinearGANSpace [62]	26.1	5.5	19.1	47.4	26.4	24.7	29.9
WarpedGANSpace [60]	27.6	56.2	33.6	6.3	14.6	8.4	29.3
CurvilinearGANSpace (ours)	21.1	15.4	25.3	6.0	18.9	9.6	19.2

A: “age”, G: “gender”, R: “race” B: “bangs”, P: “pitch”, Y: “yaw”,
Avg.: average.

that attribute vector fields are locally linearly independent, and hence, always assigns a different direction to each attribute, which facilitates the disentanglement. We describe disentanglement in the following section.

Disentanglement of Attributes Table 4 shows the side effect errors. We defined a “severe side effect” as a change in another attribute by 0.09 or more when editing a target attribute by $\tilde{t} = 0.1$, and highlighted it in bold red. All three editing methods confounded the age and bangs attributes (i.e., editing one impacted the other) owing to their high correlation, caused by unsupervised learning. CurvilinearGANSpace has no other severe side effects, whereas LinearGANSpace and WarpedGANSpace have many severe side effects; for example, WarpedGANSpace’s edit of the race attribute rather changes the bangs attribute.

Table 5 shows the identity errors. CurvilinearGANSpace produced the lowest errors for two of the six attributes, the second lowest errors for the remaining, and the lowest average error. LinearGANSpace and WarpedGANSpace produced severe identity errors in some cases.

These results indicate that only CurvilinearGANSpace selectively edits target attributes and preserves as much

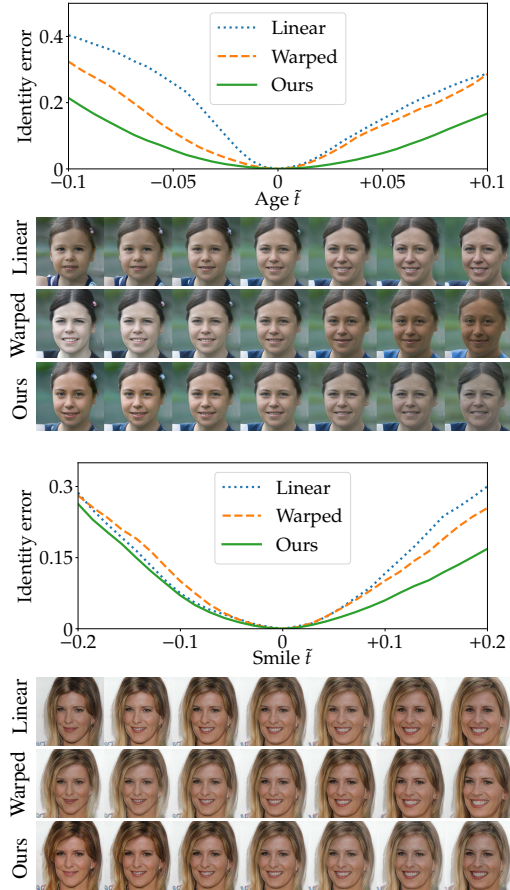


Figure 3. Identity errors when editing an attribute gradually. (top) Editing the age attribute of a StyleGAN2 image. (bottom) Editing the smile attribute of a ProgGAN image. Each row shows the results of LinearGANSpace, WarpedGANSpace, and CurvilinearGANSpace, from top to bottom.

other information as possible, that is, it is of excellent quality for the disentanglement of attributes.

Visualization of Disentanglement We summarized the identity errors and generated images when editing attributes in Fig. 3. The larger the attribute editing, the greater the increase in identity error. However, CurvilinearGANSpace has the lowest identity errors. When WarpedGANSpace edited the age attribute of a StyleGAN2 image, it also altered the skin color, hair color, and facial expression. LinearGANSpace also lost the identity. WarpedGANSpace’s edit of the smile attribute of a ProgGAN image altered the pitch and yaw rotations. CurvilinearGANSpace can edit only specific attributes with a smaller loss of identity.

We summarized the results of other models when editing specific attributes in Fig. 4. The image editing by CurvilinearGANSpace was as intended with the least side effects; however, the image editing by LinearGANSpace

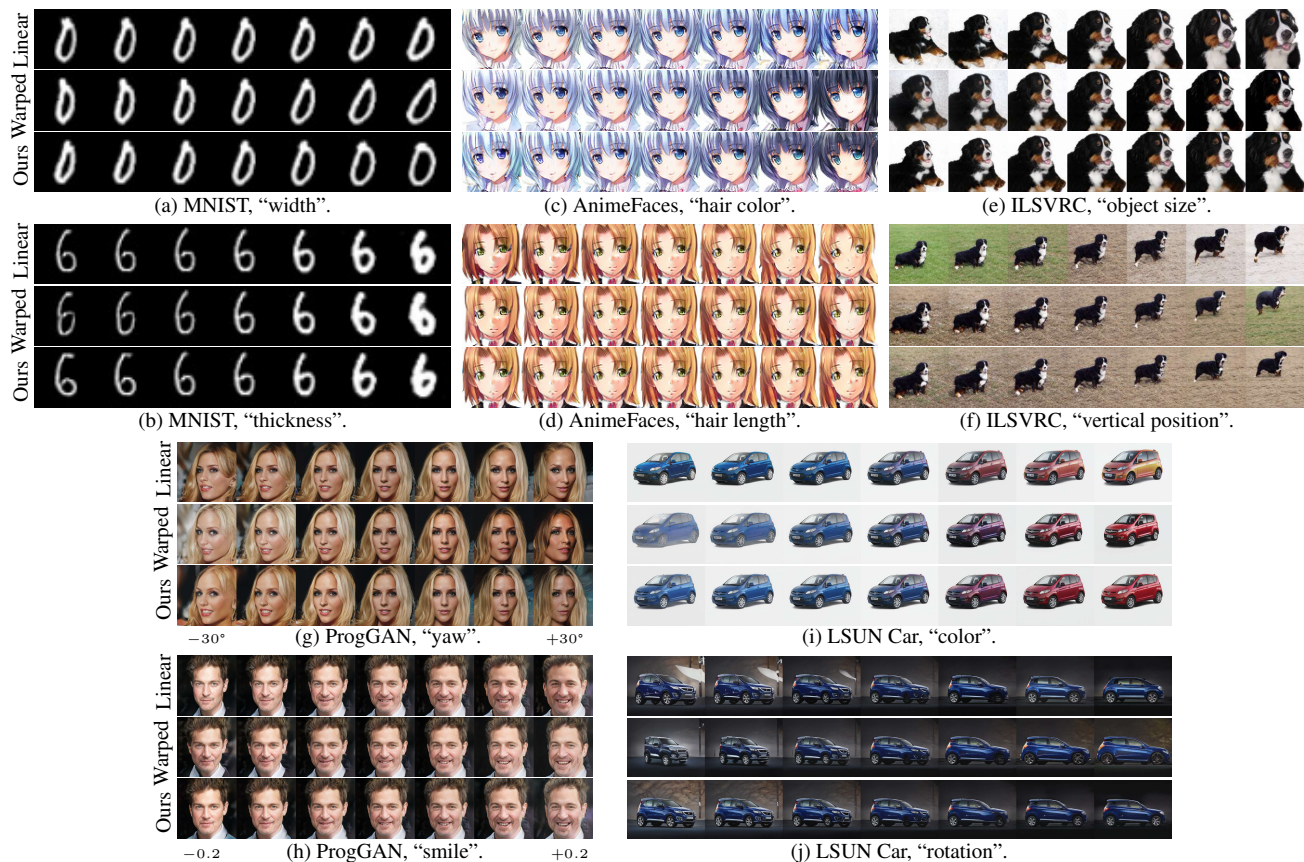


Figure 4. Visualization results. The models and edited attributes k are shown below the panels. The image in the center is the original, the images on the right have attributes added, and the images on the left have attributes subtracted, as is the case with Fig. 3.

or WarpedGANSpace exhibited severe side effects. Panels (a) and (b) show the edits of the width of digit 0 and the thickness of digit 6 in the MNIST dataset, respectively. LinearGANSpace and WarpedGANSpace additionally rotated the digits, while CurvilinearGANSpace maintained the original direction. Panels (c) and (d) show the results of the AnimeFaces dataset. The edit of the hair color by LinearGANSpace or WarpedGANSpace altered the face (i.e., loses the identity). When LinearGANSpace and WarpedGANSpace lengthened the hair, they paradoxically reduced the face’s shading. Panels (e) and (f) demonstrate that, when LinearGANSpace and WarpedGANSpace enlarged or vertically moved dogs in photos, they changed the orientations and backgrounds. Panel (g) shows that, when editing the yaw attribute, LinearGANSpace also edits the hairstyle, and WarpedGANSpace edits the skin color. Panel (h) shows that WarpedGANSpace’s edit of the smile attribute altered the pitch and yaw rotations, as in Fig. 3. Panels (i) and (j) show similar tendencies for LSUN Car.

Therefore, we conclude that the nonlinear and commutative nature of DeCurvEd contributes to the disentanglement between attributes and high-quality editing. We also provide additional results in Appendix D.

6. Conclusion

This study proposed deep curvilinear editing (DeCurvEd), which defines a curvilinear coordinate on the latent space of generative models and edits images along axes of the coordinate. DeCurvEd’s edits of semantic attributes are theoretically nonlinear and commutative. Combined with pre-trained GANs, we proposed CurvilinearGANSpace and experimentally demonstrated that it is superior to previous methods whose edits are linear or non-commutative in terms of the disentanglement between attributes and the preservation of identity. Future work will focus on a combination of DeCurvEd with other deep generative models in supervised and unsupervised learning, such as Khrukov et al. [37].

Limitations: Because DeCurvEd assumes a continuous change in attribute, it is unavailable for a discrete attribute, such as “wearing sunglasses.” A combination with discrete attributes remains a topic for future research.

Acknowledgements: This work was supported by JST PRESTO (JPMJPR21C7), CREST (JPMJCR1914), and JSPS KAKENHI (19K20344, 19H04172), Japan.

References

- [1] Rameen Abdal, Peihao Zhu, John C. Femiani, Niloy Jyoti Mitra, and Peter Wonka. CLIP2StyleGAN: Unsupervised Extraction of StyleGAN Edit Directions. In *SIGGRAPH*, 2022. [1](#)
- [2] Rameen Abdal, Peihao Zhu, Niloy J. Mitra, and Peter Wonka. StyleFlow: Attribute-Conditioned Exploration of StyleGAN-Generated Images Using Conditional Continuous Normalizing Flows. *ACM Transactions on Graphics*, 40(3):1–21, 2021. [1](#), [3](#)
- [3] George B. Arfken, Hans J. Weber, and Frank E. Harris. *Mathematical Methods for Physicists: A Comprehensive Guide*. Academic Press, Jan. 2012. [2](#), [4](#)
- [4] Georgios Arvanitidis, Lars Kai Hansen, and Søren Hauberg. Latent Space Oddity: On the Curvature of Deep Generative Models. In *International Conference on Learning Representations (ICLR)*, page 15, 2018. [3](#)
- [5] Georgios Arvanitidis, Søren Hauberg, and Bernhard Schölkopf. Geometrically Enriched Latent Spaces. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1–23, 2020. [3](#)
- [6] Yoshua Bengio, Aaron C. Courville, and Pascal Vincent. Representation Learning: A Review and New Perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35:1798–1828, 2013. [2](#)
- [7] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In *International Conference on Learning Representations*, pages 1–35, 2019. [1](#), [5](#), [13](#)
- [8] Nutan Chen, Alexej Klushyn, Richard Kurle, Xueyan Jiang, Justin Bayer, and Patrick van der Smagt. Metrics for deep generative models. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1540–1550, 2018. [3](#)
- [9] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural Ordinary Differential Equations. In *Advances in Neural Information Processing Systems*, pages 1–18, 2018. [2](#), [4](#), [13](#)
- [10] 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 1–14, 2016. [1](#), [2](#)
- [11] Zikun Chen, Ruowei Jiang, Brendan Duke, Han Zhao, and Parham Aarabi. Exploring Gradient-based Multi-directional Controls in GANs. In *European Conference on Computer Vision*, pages 1–23, 2022. [1](#)
- [12] Jaewoong Choi, Changyeon Yoon, Junho Lee, Jung Ho Park, Geonho Hwang, and Myung joo Kang. Do Not Escape From the Manifold: Discovering the Local Coordinates on the Latent Space of GANs. In *International Conference on Learning Representations*, pages 1–24, 2022. [1](#), [2](#), [3](#), [4](#)
- [13] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. StarGAN v2: Diverse Image Synthesis for Multiple Domains. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8185–8194, 2020. [1](#)
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, K. Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition*, pages 1–8, 2009. [5](#), [13](#)
- [15] Jiansheng Deng, J. Guo, and Stefanos Zafeiriou. ArcFace: Additive Angular Margin Loss for Deep Face Recognition. In *Conference on Computer Vision and Pattern Recognition*, pages 4685–4694, 2019. [6](#)
- [16] Bardia Doosti, Shujon Naha, Majid Mirbagheri, and David Crandall. HOPE-Net: A Graph-based Model for Hand-Object Pose Estimation. In *Computer Vision and Pattern Recognition*, pages 6607–6616, 2020. [6](#)
- [17] Nurit Spingarn Eliezer, Ron Banner, and Tomer Michaeli. GAN “Steerability” without optimization. In *International Conference on Learning Representations*, pages 1–58, 2021. [1](#)
- [18] Maayan Frid-Adar, Idit Diamant, Eyal Klang, Michal Amitai, Jacob Goldberger, and Hayit Greenspan. GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification. *Neurocomputing*, 321:321–331, 2018. [1](#)
- [19] Lore Goetschalckx, Alex Andonian, Aude Oliva, and Phillip Isola. GANalyze: Toward Visual Definitions of Cognitive Image Properties. In *International Conference on Computer Vision*, pages 5743–5752, 2019. [1](#)
- [20] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, pages 1–9, 2014. [1](#), [2](#), [5](#)
- [21] Will Grathwohl, Ricky T. Q. Chen, Jesse Bettencourt, Ilya Sutskever, and David Kristjansson Duvenaud. FFLORD: Free-form Continuous Dynamics for Scalable Reversible Generative Models. In *International Conference on Learning Representations*, pages 1–13, 2019. [4](#), [5](#), [13](#)
- [22] Ren’e Haas, Stella Grasshof, and Sami S. Brandt. Tensor-based Emotion Editing in the StyleGAN Latent Space. In *CVPR 2022 Workshop on AI for Content Creation Workshop*, pages 1–10, 2022. [1](#)
- [23] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Computer Vision and Pattern Recognition*, pages 770–778, 2016. [5](#), [13](#)
- [24] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. β -VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. In *International Conference on Learning Representations (ICLR)*, pages 1–14, 2017. [1](#), [2](#), [5](#)
- [25] Irina Higgins, Loic Matthey, Arka Pal, Christopher P. Burgess, Xavier Glorot, Matthew M. Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. In *International Conference on Learning Representations*, pages 1–22, 2017. [2](#)
- [26] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 6840–6851, 2020. [1](#), [2](#), [5](#)

- [27] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. GANSpace: Discovering Interpretable GAN Controls. In *Advances in Neural Information Processing Systems*, pages 1–29, 2020. [1](#), [2](#), [3](#), [5](#)
- [28] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-Image Translation with Conditional Adversarial Networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5967–5976, 2017. [1](#), [2](#)
- [29] Ali Jahanian, Lucy Chai, and Phillip Isola. On the “steerability” of generative adversarial networks. In *International Conference on Learning Representations*, 2020. [1](#)
- [30] Yuming Jiang, Ziqi Huang, Xingang Pan, Chen Change Loy, and Ziwei Liu. Talk-to-Edit: Fine-Grained Facial Editing via Dialog. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13799–13808, 2021. [5](#)
- [31] Yanghua Jin, Jiakai Zhang, Minjun Li, Yingtao Tian, and Huachun Zhu. Towards the High-quality Anime Characters Generation with Generative Adversarial Networks. In *Proceedings of the Machine Learning for Creativity and Design Workshop at NeurIPS*, pages 1–13, 2017. [5](#), [13](#)
- [32] Kimmo Kärkkäinen and Jungseock Joo. FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age for Bias Measurement and Mitigation. In *Winter Conference on Applications of Computer Vision*, pages 1548–1558, 2021. [5](#)
- [33] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *International Conference on Learning Representations*, pages 1–26, 2018. [1](#), [5](#), [13](#)
- [34] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-Free Generative Adversarial Networks. In *Advances in Neural Information Processing Systems*, 2021. [1](#)
- [35] Tero Karras, Samuli Laine, and Timo Aila. A Style-Based Generator Architecture for Generative Adversarial Networks. *Conference on Computer Vision and Pattern Recognition*, pages 4396–4405, 2019. [1](#)
- [36] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and Improving the Image Quality of StyleGAN. In *Conference on Computer Vision and Pattern Recognition*, pages 8107–8116, 2020. [1](#), [5](#), [13](#)
- [37] Valentin Khulkov, Leyla Mirvakhabova, I. Oseledets, and Artem Babenko. Latent Transformations via NeuralODEs for GAN-based Image Editing. In *International Conference on Computer Vision*, pages 14408–14417, 2021. [1](#), [3](#), [8](#)
- [38] Diederik P Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, pages 1–15, 2015. [5](#)
- [39] Diederik P. Kingma and Prafulla Dhariwal. Glow: Generative Flow with Invertible 1x1 Convolutions. In *Advances in Neural Information Processing Systems*, pages 1–10, 2018. [2](#), [4](#), [13](#)
- [40] Tin Kramberger and Božidar Potočnik. Lsun-stanford car dataset: Enhancing large-scale car image datasets using deep learning for usage in gan training. *Applied Sciences*, 10(14), jul 2020. [5](#), [13](#)
- [41] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. [5](#), [13](#)
- [42] Yann LeCun and Corinna Cortes. The mnist database of handwritten digits, 2005. [5](#), [13](#)
- [43] John M. Lee. *Introduction to Smooth Manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer, New York, NY, 2012. [2](#), [3](#), [4](#), [12](#)
- [44] Hanbang Liang, Xianxu Hou, and Linlin Shen. SSFlow: Style-guided Neural Spline Flows for Face Image Manipulation. *ACM International Conference on Multimedia*, pages 1–9, 2021. [1](#), [3](#)
- [45] Zinan Lin, Kiran Koshy Thekumparampil, Giulia C. Fanti, and Sewoong Oh. InfoGAN-CR and ModelCentrality: Self-supervised Model Training and Selection for Disentangling GANs. In *International Conference on Machine Learning*, pages 1–45, 2020. [1](#), [2](#)
- [46] Bingchen Liu, Yizhe Zhu, Zuohui Fu, Gerard de Melo, and A. Elgammal. OOGAN: Disentangling GAN with One-Hot Sampling and Orthogonal Regularization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4836–4843, 2020. [1](#), [2](#)
- [47] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep Learning Face Attributes in the Wild. *International Conference on Computer Vision*, pages 3730–3738, 2015. [5](#), [13](#)
- [48] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*, pages 1–26, 2018. [1](#), [5](#), [13](#)
- [49] James Oldfield, Markos Georgopoulos, Yannis Panagakis, Mihalis A. Nicolaou, and I. Patras. Tensor Component Analysis for Interpreting the Latent Space of GANs. In *British Machine Vision Conference*, pages 1–18, 2021. [1](#)
- [50] Antoine Plummerault, Hervé Le Borgne, and Céline Hudelot. Controlling generative models with continuous factors of variations. In *International Conference on Learning Representations*, pages 1–17, 2020. [1](#)
- [51] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. In *International Conference on Learning Representations*, pages 1–10, 2016. [1](#), [2](#)
- [52] Aditya Ramesh, Youngduck Choi, and Yann LeCun. A Spectral Regularizer for Unsupervised Disentanglement. *ArXiv*, pages 1–17, 2018. [1](#), [2](#), [3](#), [4](#)
- [53] Thomas Schlegl, Philipp Seeböck, Sebastian M. Waldstein, Ursula Margarethe Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International Conference on Information Processing in Medical Imaging*, 2017. [1](#)
- [54] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the Latent Space of GANs for Semantic Face Editing. In *Computer Vision and Pattern Recognition*, pages 9240–9249, 2020. [1](#)

- [55] Yujun Shen and Bolei Zhou. Closed-Form Factorization of Latent Semantics in GANs. In *Computer Vision and Pattern Recognition*, pages 1532–1540, 2021. [1](#), [2](#), [3](#), [5](#)
- [56] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations (ICLR)*, pages 1–32, 2021. [1](#), [2](#), [5](#)
- [57] Takeshi Teshima, Koichi Tojo, Masahiro Ikeda, Isao Ishikawa, and Kenta Oono. Universal Approximation Property of Neural Ordinary Differential Equations. In *NeurIPS Workshop on Differential Geometry Meets Deep Learning (DiffGeo4DL)*, 2020. [13](#)
- [58] Ayush Tewari, Mohamed A. Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhöfer, and Christian Theobalt. StyleRig: Rigging StyleGAN for 3D Control over Portrait Images. In *Computer Vision and Pattern Recognition*, pages 6141–6150, 2020. [1](#)
- [59] Christos Tzelepis, James Oldfield, Georgios Tzimiropoulos, and Ioannis Patras. ContraCLIP: Interpretable GAN generation driven by pairs of contrasting sentences. *arXiv*, June 2022. [3](#)
- [60] Christos Tzelepis, Georgios Tzimiropoulos, and Ioannis Patras. WarpedGANSpace: Finding Non-Linear RBF Paths in GAN Latent Space. In *International Conference on Computer Vision*, pages 6393–6402, 2021. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [12](#), [13](#), [14](#)
- [61] Aaron van den Oord, Nal Kalchbrenner, Oriol Vinyals, Lasse Espeholt, Alex Graves, and Koray Kavukcuoglu. Conditional Image Generation with PixelCNN Decoders. In *Advances in Neural Information Processing Systems (NIPS)*, pages 4790–4798, 2016. [1](#), [2](#), [5](#)
- [62] Andrey Voynov and Artem Babenko. Unsupervised Discovery of Interpretable Directions in the GAN Latent Space. In *International Conference on Machine Learning*, pages 9786–9796, 2020. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [12](#), [14](#)
- [63] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-Video Synthesis. In *Advances in Neural Information Processing Systems*, pages 1–14, 2018. [1](#)
- [64] Po-Wei Wu, Yu-Jing Lin, Che-Han Chang, Edward Y. Chang, and Shih-Wei Liao. RelGAN: Multi-Domain Image-to-Image Translation via Relative Attributes. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5913–5921, 2019. [1](#), [2](#)
- [65] W. Xia, Y. Zhang, Y. Yang, J. Xue, B. Zhou, and M. Yang. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. [4](#)
- [66] Jiapeng Zhu, Ruili Feng, Yujun Shen, Deli Zhao, Zhengjun Zha, Jingren Zhou, and Qifeng Chen. Low-Rank Subspaces in GANs. In *Advances in Neural Information Processing Systems*, pages 1–23, 2021. [2](#)
- [67] Jiapeng Zhu, Yujun Shen, Yinghao Xu, Deli Zhao, and Qifeng Chen. Region-Based Semantic Factorization in GANs. In *International Conference on Machine Learning*, 2022. [2](#)
- [68] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2242–2251, 2017. [1](#), [2](#)
- [69] P. Zhuang, O. Koyejo, and A. G. Schwing. Enjoy Your Editing: Controllable GANs for Image Editing via Latent Space Navigation. In *International Conference on Learning Representations*, pages 1–16, 2021. [1](#)