

Exploring Gradient-based Multi-directional Controls in GANs

Zikun Chen¹, Ruowei Jiang¹, Brendan Duke^{1,2}, Han Zhao³, and Parham Aarabi^{1,2}

¹ ModiFace Inc., {zikun, irene, brendan, parham}@modiface.com

² University of Toronto

³ University of Illinois at Urbana-Champaign, hanzhao@illinois.edu

Abstract. Generative Adversarial Networks (GANs) have been widely applied in modeling diverse image distributions. However, despite its impressive applications, the structure of the latent space in GANs largely remains as a black-box, leaving its controllable generation an open problem, especially when spurious correlations between different semantic attributes exist in the image distributions. To address this problem, previous methods typically learn linear directions or individual channels that control semantic attributes in the image space. However, they often suffer from imperfect disentanglement, or are unable to obtain multi-directional controls. In this work, in light of the above challenges, we propose a novel approach that discovers nonlinear controls, which enables multi-directional manipulation as well as effective disentanglement, based on gradient information in the learned GAN latent space. More specifically, we first learn interpolation directions by following the gradients from classification networks trained separately on the attributes, and then navigate the latent space by exclusively controlling channels activated for the target attribute in the learned directions. Empirically, with small training data, our approach is able to gain fine-grained controls over a diverse set of bi-directional and multi-directional attributes, and we showcase its ability to achieve disentanglement significantly better than state-of-the-art methods both qualitatively and quantitatively. The source code is available at <https://github.com/zikuncshelly/GradCtrl>.

Keywords: GAN, gradient information, latent space, disentanglement, multi-directional

1 Introduction

Generative adversarial networks (GANs) [9] are implicit generative models from which it is straightforward to sample data from the learned distributions. GANs work by learning to generate data from random noise sampled from natural distributions such as isotropic Gaussians. Due to the flexibility of neural networks in modeling complex transformations, GANs and the follow-up variants [3,13,17,26] have been widely applied in modeling image distributions. Furthermore, the remarkable success of GANs' ability to generate realistically looking images has

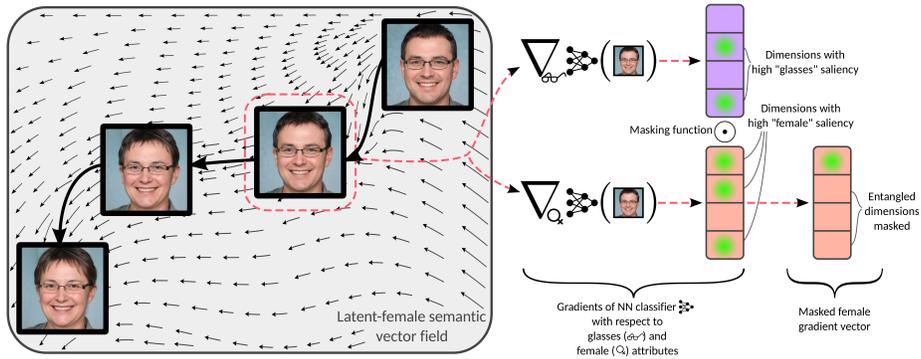


Fig. 1: Our method’s semantic interpolation in a nonlinear vector field. Directions corresponding to the “female” semantic in GAN latent space (left) are determined by differentiating a simple neural network classifier (middle). Disentanglement is achieved by masking channels of the “female” semantic gradient vector that are highly salient for non-target attributes such as glasses (right).

powered many real-world applications, including image inpainting [30,32], super-resolution [5,14], and image-to-image translations [11,33].

However, despite the impressive existing applications, the structure of the latent space in GANs largely remains as a black box. For example, is there a correspondence between interpolation directions in the latent space and visual factors in the generated images? If yes, how can we generate images in a controllable way? Furthermore, are different interpolation directions independent of each other? If not, is it possible to disentangle them in an interpretable fashion?

To answer the above important but challenging questions, existing works have studied the learned latent space by identifying and linearly interpolating along directions that correspond to certain semantic factors in the generated images [10,18,21,24,28]. For example, on facial image data, early works [18,24] use simple linear arithmetic operations in the latent space to control the age of the generated persons. Recent works [21,29] further studied the interpolation directions by learning a linear binary classifier, and illustrated the effects of linearly interpolating different channels in a supervised or unsupervised manner [10,28].

On the other hand, existing works aim to identify directions that correspond to unique semantic factors, and to control generation based on a single factor. However, semantic factors in images, such as age and eyeglasses, are often highly correlated with each other. As a result, artifacts often appear when applying controllable generation to one factor. Another issue arises due to spurious correlations in the training image data. These cause undesirable changes to the generated image to appear even if controllable generation is applied to an independent semantic factor. Furthermore, existing approaches still fail to discover multi-directional controls in the latent space and suffer from imperfect disentanglement between different interpolation directions. In fact, due to the nature of linear interpolation, existing works can only apply linear bi-directional controls

over each semantic attribute. For example, to modify the color of a car, existing approaches need to learn all the pairwise connections among different colors, which is nontrivial [10,28].

In light of the aforementioned challenges in controllable generations using GANs, in this work we propose a novel approach that first explores the multi-directional controls and addresses the disentanglement via gradient-based information in the learned GAN latent space (Fig. 1). Inspired by existing visualization techniques on classification networks [4,20], we navigate the latent space by using gradient-based information learned from a classification network separately trained on the factors and attributes of interest. Using the latent vectors from a pre-trained GAN network [13] as inputs, we train classification networks over all attributes. To disentangle the controllable generation of different attributes, we separate the manipulation into two steps. First, we identify the direction using the gradient on the target attribute of interest. Then, we exclusively control the channels that only affect the probability of the target attribute, based on gradient-weighted attribute activations.

To demonstrate the effectiveness of our approach, we present our method’s ability to obtain nonlinear multi-directional controls as well as disentangled manipulations on multiple benchmark datasets. Furthermore, we show that our classification network can gain proper controls by using a tiny dataset of 30 images per class. We also compare our results to the state-of-the-art models both qualitatively and quantitatively. Finally, we present real image manipulation results on several attributes by inverting the image back to a latent code. To our knowledge, this work provides the first nonlinear multi-directional controls that achieve disentanglement successfully. Overall, our contributions are:

- We propose a novel approach to navigate the GAN latent space by gradient-based information that allows nonlinear multi-directional controls.
- We disentangle the controls of different attributes by gradient-based channel selection to exclusively control a single attribute without affecting the others.
- Empirically, we show that our method discovers disentangled controls with small amounts of data, widening the range of applications of our method.

2 Related Work

Conditional Manipulations on GANs. GANs have been widely applied to a variety of tasks that allow conditional manipulations, including image editing [7,11,15,16,17,19,21,26,33], image inpainting [30,32], and super-resolution [5,14], just to name a few. Early works [11,33] explored image translations between two domains by using unpaired and paired training images. Some more recent works focused on improving diversity by translating in multiple domains [7,8] while other works tried to improve the quality of the manipulation by effectively propagating the conditional information [17]. However, these methods often require a huge amount of paired data for training, which limits their applications in settings where it is expensive to manually label these pairs.

Understanding GAN Latent Space. Many attempts have been made to understand and visualize the latent representations of GANs [1,2,10,18,21,24,28,29]. Bau et al. [1,2] used semantic segmentation models to detect and quantify the existence and absence of certain objects at both the individual and population levels. Previous work [18,24] empirically studied linear vector arithmetic in the latent representations. To gain direct controls along attribute manipulations, later works [21,29] further explored linear interpolation in the latent space by finding semantically meaningful directions with explicit binary classifiers. InterFaceGAN [21] hypothesized and showed the existence of a hyperplane serving as a linear separation boundary in the StyleGAN [12] latent space for some binary semantic relationships. In addition, they attempted to solve the entanglement issue on interpolated images via subspace projections between entangled attributes. However, identity change is often observed in simple linear interpolations settings, and the disentangled direction sometimes fails to manipulate the target attribute. A more recent work [28] analyzed StyleGAN’s style space and proposed a method to obtain valid semantic controls by identifying highly localized channels that only affect a specific region, which achieves well-disentangled changes on some attributes. However, this approach lacks the ability to control globally aligned features such as age.

Another line of works aim to identify semantic controls in a self-supervised or unsupervised manner [6,10,25]. For example, GANSpace [10] identifies important latent directions by applying principal component analysis (PCA) to vectors in GAN latent space or feature space. Voynov et al. [25] discover interpretable directions in the GAN latent space by jointly optimizing the directions and a reconstructor that recovers these directions and manipulation strengths from images generated based on the manipulated GAN latent codes. These types of methods usually involve extensive manual examinations of different manipulation directions as well as the identification of meaningful controls. More importantly, in contrast to our approach, controlling over a target semantic is not always guaranteed and attributes are often observed as entangled due to the potential correlations in the training image distribution.

Gradient-based Network Understanding. Gradient-based interpretation methods [4,20,22,23] have been widely explored on discriminative models. Grad-CAM [4,20] used gradient information from the classification outputs to obtain localization maps that visualize evidence in images. In contrast to local sensitivity, Integrated Gradients [23] estimated the global importance of each pixel to address gradient saturation. Integrated Gradients generated so-called gradient-based sensitivity maps, which were further sharpened by SmoothGrad [22].

The generative model LOGAN [27] increased the stability of the training dynamics of GANs by adding an extra step before optimizing the generator and discriminator jointly. In LOGAN, the latent code is first optimized towards regions regarded as more real by the discriminator, and the direction is obtained by calculating the gradient with respect to the latent codes. Intuitively, the gradient points in a direction that obtains a better score from the discriminator. We observed that when using LOGAN the discriminator gradient direction corre-

sponds to meaningful transformations in GAN-synthesized images, for example, removal of artifacts, change in gender, and so on. This observation inspired us to use auxiliary classifier gradients to obtain controls over GAN output semantics. To the best of our knowledge, we are the first to take this approach.

3 Methods

In this section, we describe our proposed method to obtain disentangled controls in the GAN latent space. Our method first finds semantically meaningful directions by calculating gradients of multi-class classifiers that score different semantics given latent code inputs. Next, by selecting the channels in the latent code that correspond to these directions, our technique achieves attribute disentanglement during manipulation through masking.

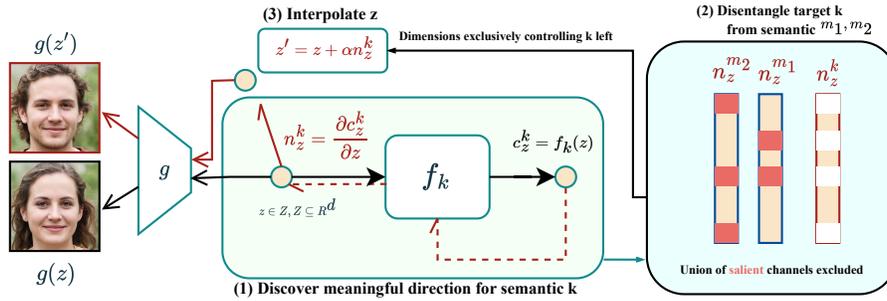


Fig. 2: Overview of our method. We first obtain the control for a semantic attribute k , e.g., gender, by calculating the gradient direction n_z^k from the mapping function f_k in (1). To disentangle direction n_z^k from attributes m_1, m_2 , e.g., smile and age, the union of salient dimensions in control directions $n_z^{m_1}, n_z^{m_2}$ for the corresponding semantics are masked to 0 in n_z^k , as shown in (2). We then interpolate the latent code following the direction computed at each step in (3).

3.1 Discovery of Semantically Meaningful Directions

Learning Semantics in GAN Latent Space. GANs learn a mapping $g : Z \rightarrow X$ from a latent code z in its latent space $Z \subseteq \mathbb{R}^d$ to an image $x = g(z)$ in the image space $X \subseteq \mathbb{R}^{H \times W \times 3}$, where $H \times W$ is the spatial resolution. Human-interpretable semantics such as age of a person exist in the image space. Given a series of scoring functions (s_1, \dots, s_K) for K semantics, we can obtain K mappings from the GAN latent space Z to the semantic spaces C_1, \dots, C_K respectively, where

$$c_z^k = s_k(g(z)) \in C_k. \quad (1)$$

Here c_z^k denotes the k th semantic score of the image generated based on the latent code z , and $C_k \in \mathbb{R}^{d_k}$, where d_k denotes the number of classes the k th semantic has, e.g., for binary semantics, $d_k = 1$. We parametrize each mapping function $s_k(g(z))$ with a neural network f_k trained on paired samples of GAN latent codes and corresponding semantic labels. The labels are generated by pre-trained image classifiers or human supervision. With (f_1, \dots, f_k) , we are able to estimate human-interpretable semantics based on GAN latent codes, without inferring from the image space.

Manipulating Semantics in GAN Latent Space. Provided the accuracy of (f_1, \dots, f_k) , it is expected that when controlling the k th semantic by interpolating z along meaningful directions, its score c_z^k changes accordingly. We hypothesize that such changes correspond to the changes in image space semantics. Inspired by LOGAN [27], which hypothesizes that the gradient from a GAN discriminator points in the direction of better samples, as the discriminator can be seen as a scoring function for “realness”, we propose to control the k semantics using the Jacobians of (f_1, \dots, f_k) . In particular, our method controls c_z^k by interpolating the latent code following gradient directions that are rows of the Jacobian of f_k with respect to z :

$$n_z^k = \nabla f_k(z) = \frac{df_k(z)}{dz} \in \mathbb{R}^{d_k \times d}, \quad (2)$$

where the matrix n_z^k denotes the Jacobian of the k th latent-semantic mapping c_z^k at z . Note that for binary attributes, $d_k = 1$ and we use $n_z^k = n_z^k[0]$ to denote a gradient vector. For multi-class attributes, when optimizing towards the j th class for semantic k , we take the corresponding row in the Jacobian and $n_z^k = n_z^k[j]$. For notational simplicity, $n_z^k[j]$ is abbreviated as n_z^k in the following sections, so that n_z^k always refers to a gradient vector.

Intuitively, for each target semantic, interpolating the latent code following the gradient from its scoring function changes the semantic score, then the semantic is manipulated accordingly. More specifically, to manipulate the semantic k with respect to its i th class once, we update z with

$$z' = z + \alpha n_z^k, \quad (3)$$

where α is a hyperparameter that controls the interpolation step size and sign of change in the target semantic value. In contrast to works that discovered linear semantically meaningful directions applicable to all coordinates in GAN latent space [21], our method finds unique local and data-dependent directions for each coordinate due to the nonlinear nature of the latent-semantic mapping (Eqn. 1).

3.2 Disentanglement of Attributes During Manipulation

Entanglements such as changing the gender affecting one’s age can emerge during latent code interpolation following the gradient of the latent-semantic mapping (Eqn. 3). In this section, we propose a technique to minimize this effect.

Algorithm 1 Disentangle attributes

-
- 1: Target attribute k is entangled with attributes m_1, \dots, m_j in direction n_z^k
 - 2: Scoring function f_{m_1}, \dots, f_{m_j} , numbers of top channels excluded $\{c_1, \dots, c_j\}, j \neq k$
 - 3: Excluded set of dimensions $E = \{\}$
 - 4: For each $(m, f_m, c) \in (\{m_1, \dots, m_j\}, \{f_{m_1}, \dots, f_{m_j}\}, \{c_1, \dots, c_j\})$:
 - 5: $n_z^m = \nabla f_m(z) = \frac{df_m(z)}{dz}, L_i^m = |(n_z^m)_i|$
 - 6: $t = \text{sorted}(L^m)[-c], e = \{i\}$ where $L_i^m \geq t$
 - 7: $E = E \cup e$
 - 8: $(n_z^k)_i = 0$ for $i \in E$
 - 9: Return n_z^k
-

Target-specific Dimensions in Direction Vectors. In early experiments, we observed that interpolation along the gradient direction alters non-target semantics, but disentanglement is occasionally achieved by randomly excluding dimensions in the direction vector used for interpolation. Therefore, we hypothesize that within the d -dimensional direction vector discovered, only some dimensions are responsible for the change in the target semantic. Conversely, the remaining channels denote bias learned from the training data. For example, when interpolating in the direction that increases a person’s age, eyeglasses appear since eyeglasses and age are correlated in the dataset.

Filtering Dimensions to Control Target Attributes Exclusively. To construct disentangled controls, we propose to filter out dimensions based on gradient magnitudes from the semantic scoring function in a technique inspired by Grad-CAM [20]. Grad-CAM measures the importance of neurons by

$$a_k^c = \frac{1}{Z} \sum_i \sum_j \frac{\partial y^c}{\partial A_{ij}^k}, \quad (4)$$

where the derivative of the score y^c for class c with respect to the activation map A is average pooled over spatial dimensions. The final heatmap is then

$$L_{\text{Grad-CAM}}^c = \text{ReLU}\left(\sum_k a_k^c A^k\right). \quad (5)$$

For our approach, gradient n_z^k is regarded as the only activation map, and the saliency of the i th dimension is calculated by

$$L_i^k = |(n_z^k)_i|, \quad (6)$$

where $|(n_z^k)_i|$ denotes the absolute value of the i th element of the gradient vector. Contrary to the original classification setting for Grad-CAM where ReLU is chosen as the activation, here the absolute value is used because in our classification setting dimensions with negative influence on the semantic score also contain meaningful information.

By definition, the value of gradient n_z^k at its i th dimension indicates the rate of change in c_z^k induced by a small change in latent vector z along dimension i .

Intuitively, dimensions in n_z^k with greater saliency L_i^k are more relevant and have more impact on c_z^k , whereas less relevant dimensions for c_z^k could have high saliency in another semantic m 's gradient $n_z^m, m \neq k$. Hence even small changes in these dimensions while optimizing for the k th semantic could affect c_z^m significantly, causing attribute entanglement. Note that such effects could be either positive or negative, hence we use the absolute values as the saliency. To alleviate this issue, we exclude the top- k salient dimensions for predicting an entangled attribute from the target direction, with k in $\{50,100,150,200,250\}$. Algo. 1 describes the full process to calculate the new direction with disentangled attributes before interpolating the latent code (Eqn. 3).

4 Experiments

In this section, we apply our method to learned latent representations by state-of-the-art GAN models, and compare it with existing state-of-the-art latent code interpolation methods both qualitatively and quantitatively. In particular, we analyze the ability to achieve disentangled controls using each method.

4.1 Qualitative Results

Image Bank. We perform our experiments on StyleGAN2 [13] pretrained on FFHQ [12], LSUN Car [31], LSUN Cat [31], and StyleGAN [12] pretrained on LSUN Bedroom [31]. To prepare an image bank, we synthesize images by randomly sampling the latent codes and obtain labelled pairs for each attribute value, via pre-trained classifiers or human supervision.

Training the Classifier. We train a simple classifier with only 30 examples for each attribute value. The network consists of two fully connected layers, followed by the corresponding prediction head: sigmoid for binary classification and softmax for multi-class classification. In the Appendix, we also demonstrate the effects of using different numbers of training samples.

Interpolating in the Latent Space. We interpolate the latent codes following the gradient directions of our trained classifiers by a fixed step size of 0.6. We perform layer-wise edits by applying the learned directions to certain layers in $\mathbf{W}+$ space in StyleGAN/StyleGAN2. More implementation details can be found in the Appendix.

Disentangled Attribute Manipulation. In this section, we evaluate our method on StyleGAN2 models pretrained on the FFHQ [12], LSUN-Cars [31], LSUN-Cats datasets [31] and the StyleGAN model [12] pretrained on the LSUN-Bedroom dataset [31]. We demonstrate that our approach obtains multi-directional controls over GAN-generated images in a disentangled manner. We show a qualitative comparison among key attributes in the FFHQ dataset [12]: smile, gender, age, and eyeglasses (Fig. 3). For the age and gender attributes, not only is our approach able to generate realistic localized target effects such as wrinkles or acne for age, but it is also capable of editing the image on a global level while

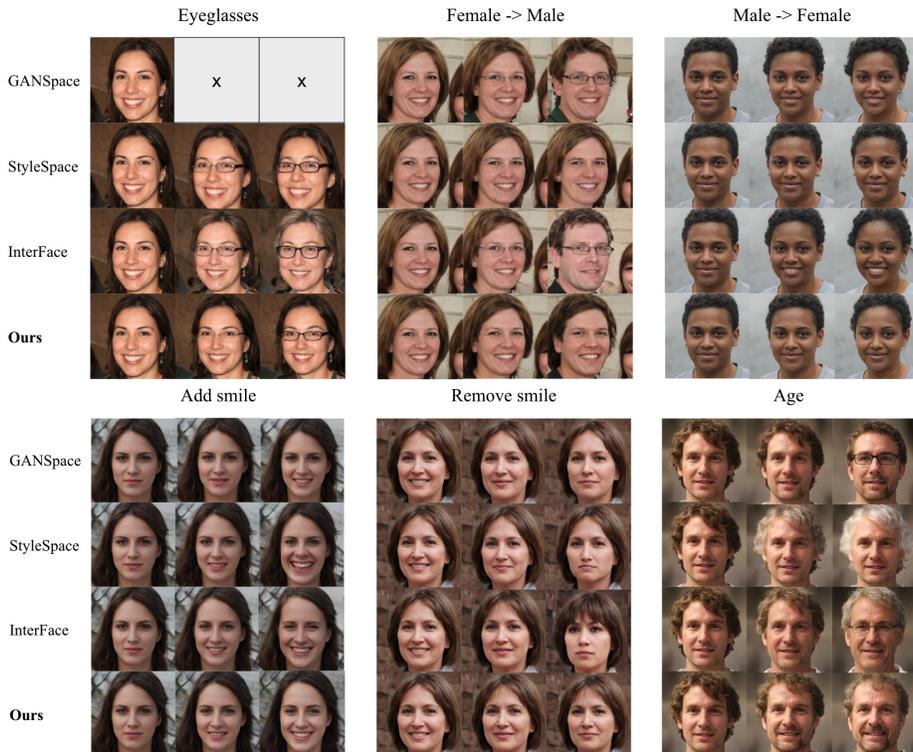


Fig. 3: Comparison with state-of-the-art methods on StyleGAN2 [13] pretrained on FFHQ [12]. For each group of images, left: source; middle, right: small and large step size. Our method edits eyeglasses and smiles naturally, and for global-level attributes, ours is able to produce more visually distinguishable editing results compared to StyleSpace [28] and GANSpace, and achieves more disentangled control than InterFaceGAN [21].

preserving the person’s identity and irrelevant attributes disentangled, such as keeping the person’s smile when changing from female to male.

We further show that our method works on a variety of attributes on multiple datasets with qualitative results on LSUN car and LSUN bedroom [31] (Fig. 4).

Multiclass Manipulation. In this section, we showcase our method’s ability to manipulate multiclass attributes. Based on the latent-semantic Jacobian (Eqn. 2), regardless of the original state of the image, we identify the control direction and the activated channels with respect to the target class. Using only one classifier, our method performs manipulations in any direction (Fig. 5).

Real Image Manipulation. We present real image editing results to verify that the semantic controls learned by our method can be applied to real images. We first invert real human face images into the W space of StyleGAN2 [13]

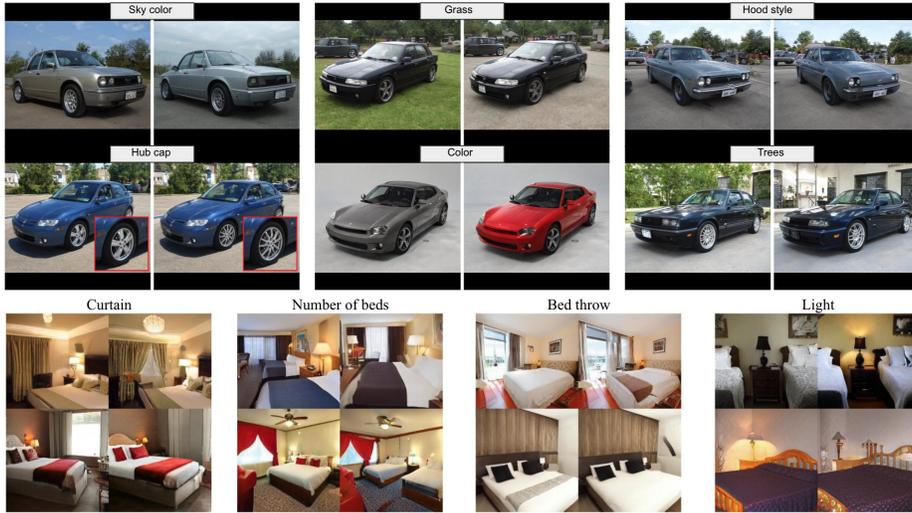


Fig. 4: LSUN-Cars and LSUN-Bedroom [31] editing results only trained on 30 pairs of examples. For each group of 4 images, left: source, right: modified.

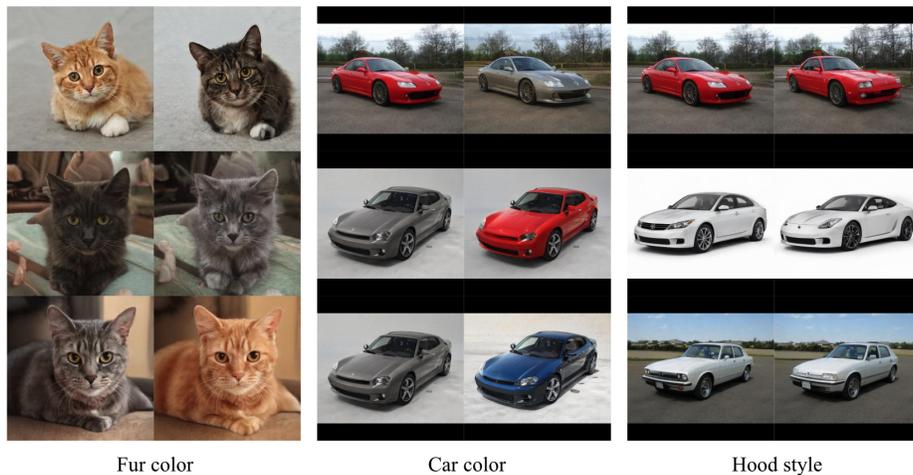
trained on FFHQ [12] via an optimization-based framework [13]. We then interpolate the latent codes along directions learned by our method. In the resulting manipulations (Fig. 6), our method manipulates all attributes successfully.

4.2 Quantitative Evaluation

We compare our method with the related works [10,21,28] by measuring the accuracy and the level of attribute entanglement. We test on StyleGAN2 [13] pretrained on the FFHQ [12] dataset. We generate a test image bank with 500k samples using the StyleGAN2 model and score 4 attributes (age, gender, smile and eyeglasses) with the corresponding pre-trained CelebA classifiers [12]. From our test image bank, we sample 3k images with logits around the decision boundaries for 4 attributes, and for each image, we sample a target attribute k uniformly from the 4 to optimize until the modified attribute logit crosses the boundary, and we include the details regarding the algorithm in the Appendix.

In a nutshell, we manipulate the latent code until the target attribute classification changes. We stop after crossing the boundary as for attributes like smile and eyeglasses, further optimization results in regions in the latent space irrelevant to the original target, e.g., removing eyeglasses from a person not wearing eyeglasses generally makes the person younger.

Attribute Manipulation Accuracy. We measure whether the attributes of an optimized image match the intended targets, i.e., change in the target attribute label and unchanged labels for the rest (Tab. 1). High accuracy corresponds to a low level of attribute entanglement and high success rate, i.e., target attribute being present. Our approach achieves the best results among all.



(a) Multi-class manipulation results for LSUN cats and cars [31]. For each group of images, left: source, right: modified. Regardless of the original color of the car/cat, our proposed method can independently identify any target control direction.



(b) Multi-class manipulation results for LSUN cars rotation. To avoid 3D-unaware changes, starting from the initial angle class, the latent code is optimized towards increments of 45 degrees which correspond to different classes.

Fig. 5: Multi-class manipulation for LSUN cats and LSUN cars [31]. Our method can easily apply multi-directional changes without learning each direction between each pair of classes.

Attribute Disentanglement. Attribute Dependency (AD) [28] measures the level of entanglements for semantic controls in GAN latent space, and here we follow a similar paradigm. On a high level, for images with target attribute $k \in A$ being manipulated, where A stands for all attributes, we group the modified images by $x = \frac{\Delta l_k}{\sigma_k}$. We then compute mean-AD with $E(\frac{1}{|A|-1} \sum_{i \in A \setminus k} \frac{\Delta l_i}{\sigma l_i})$ for each group. The full algorithm can be found in the appendix.

(Fig. 7a) plots the mean AD for each target attribute, where methods with extremely low accuracy are excluded, including StyleSpace for gender and age and GANSpace for eyeglasses as it doesn't offer such control. For attributes that only require localized changes, including eyeglasses and smile, our method performs similarly to StyleSpace and GANSpace. Our eyeglasses control shows higher AD than StyleSpace, as StyleSpace only modifies channels responsible for the eye area, yet for ours, the image is edited on a global level, and the pre-trained classifiers might be sensitive to some subtle changes. Nevertheless,

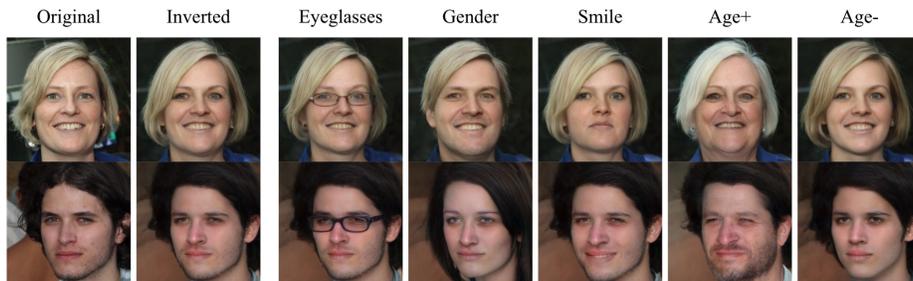
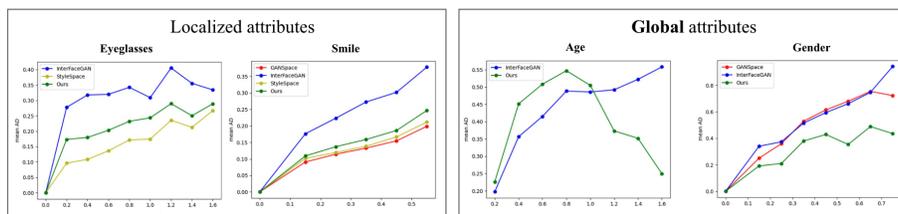
Fig. 6: Manipulating real faces by first inverting them into W space.

Table 1: Attribute manipulation accuracy in a disentangled manner. A resulting image is considered a true positive if only the target attribute is successfully changed. Among all, our method achieves the highest scores for all 4 attributes, corresponding to fewer changes in attributes other than the target during manipulation, and higher success rates, i.e., target attribute being edited successfully.

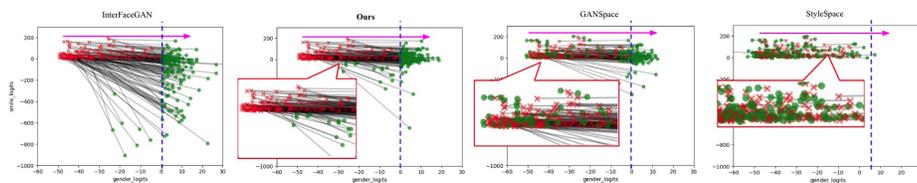
Accuracy	Gender	Smile	Eyeglasses	Age
InterFaceGAN	0.5859	0.9033	0.5946	0.5620
GANSpace	0.2770	0.6146	x	0.0063
StyleSpace	0.0812	0.7390	0.3510	0.0034
Ours	0.6937	0.9254	0.6626	0.6139

our method significantly outperforms the rest when editing gender and age, as the amount of change in the target increases. As we traverse longer distances in the latent space to increase a person’s age, the linear directions learned by the SVMs become inaccurate with lots of entanglement with gender and eyeglasses, whereas our non-linear method continues to disentangle the target direction from the others.

To provide further insights for AD, we visualize the change in smile logits when changing from non-smiling male to female in (Fig. 7b), and big negative slopes indicate entanglement between smile and gender, i.e., smile being added after optimization. InterFaceGAN suffers from the entanglement issue with a lot of big negative slopes, whereas GANSpace and StyleSpace have a lot of failure cases (the green points lying on the left of the blue vertical threshold line after interpolation in (Fig. 7b)), indicating manipulation in the wrong/unclear semantic direction. Our method has much less entanglement (slopes are more similar to the flat pink arrow) with fewer failure cases, and for those that failed, the direction is mostly correct as the gender logits are still moving towards the blue dash line.



(a) X-axis: Change in logits of the target attribute. Y-axis: Mean AD. Lower AD indicates better disentanglement. Our method (green) is significantly less entangled than InterFaceGAN and GANSpace for gender, which involves global-level changes, and achieves similar levels of disentanglement for editing smile compared to StyleSpace and GANSpace. Note that unnatural looking and identity change will not affect this calculation.



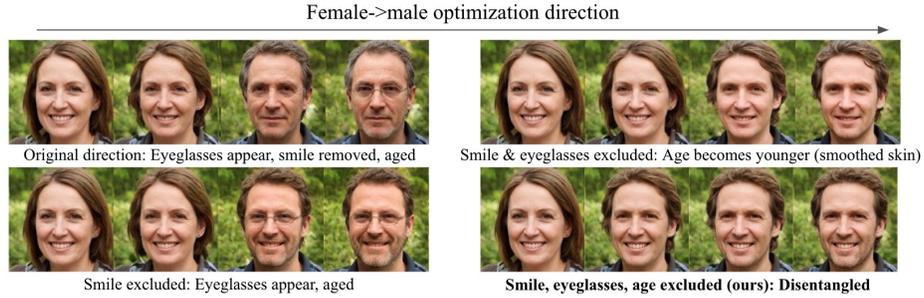
(b) Disentanglement visualization for gender and smile. X-axis: gender logits; Y-axis: smile logits. Red crosses: original logits; Green dots: logits after modification. Pink arrow: change in logits when attributes are perfectly disentangled. Blue dash line: success threshold. Flatter lines indicate less disentanglement. More green dots lying to the right of the blue dash line indicate a higher success rate. Ours exhibits the best disentanglement with a high success rate.

Fig. 7: Disentanglement analysis

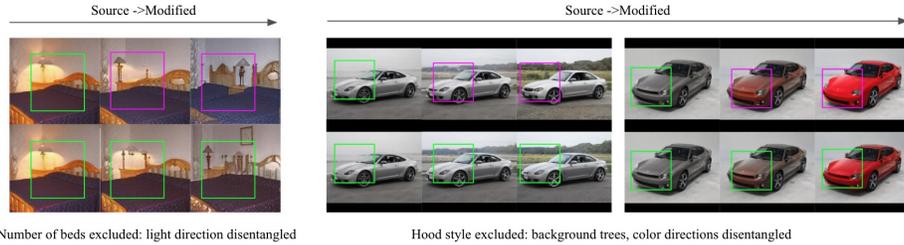
4.3 Ablation Studies: Attribute Disentanglement

In this section, we demonstrate how controls over different target attributes computed by our method can be disentangled using gradient information from each entangled attribute. We first demonstrate how our method enables disentanglement from multiple attributes for StyleGAN2 [13] pre-trained on the FFHQ [12] dataset in (Fig. 8a), where the original gender direction is entangled with smile, eyeglasses and age. By filtering out activated channels for the entangled attribute(s) calculated from the gradient weights in our classification network, we achieve different levels of disentanglement. We further present disentanglement results for StyleGAN2 [13] pretrained on the LSUN-Cars [31] dataset and StyleGAN [12] pretrained on the LSUN-Bedrooms dataset [31] in (Fig. 8b).

Next, we study the **robustness** of attribute disentanglement with respect to the number of top- k channels selected. When setting k to 100 for all attributes as shown in (Fig. 8c), our method produces disentangled results, and the level of disentanglement for editing each attribute can be further improved by manually adjusting based on the results on a small meta-validation set of 30 examples with k in $\{50, 100, 150, 200, 250\}$.



(a) Disentangle gender from other attributes one by one. The upper left row shows entanglement with eyeglasses, smile and age. At the bottom right row, by excluding the union of the activated channels for all entangled attributes, we successfully achieve disentanglement, where only the gender attribute is changed.



(b) Disentanglement results for LSUN-Cars and LSUN-Bedrooms [31]. The first row and second row shows original direction and disentangled direction respectively.



(c) Effect of number of channels excluded for smile, age, eyeglasses when editing gender. Tweaked: results shown in (Fig. 3).

Fig. 8: Ablation studies on disentanglement.

5 Discussion

Limitations. Despite the overall success of our method, we also notice some failures which lack 3D understanding (e.g. car rotation) or result in identity change. Examples and discussions can be found in the Appendix.

Conclusions. In this work, we propose a simple but powerful approach based on gradients to obtain multi-directional controls and disentangle one attribute from the others. Our findings show that attributes’ controls can be separated even with strongly correlated training datasets via gradient information. This work shows considerable potential to discover rich controls in the learned latent space via gradient-based information, as well as its effectiveness in small data regimes.

References

1. Bau, D., Zhu, J.Y., Strobel, H., Zhou, B., Tenenbaum, J.B., Freeman, W.T., Torralba, A.: Gan dissection: Visualizing and understanding generative adversarial networks. In: Proceedings of the International Conference on Learning Representations (ICLR) (2019) [4](#)
2. Bau, D., Zhu, J.Y., Wulff, J., Peebles, W., Strobel, H., Zhou, B., Torralba, A.: Seeing what a gan cannot generate. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4502–4511 (2019) [4](#)
3. Brock, A., Donahue, J., Simonyan, K.: Large scale GAN training for high fidelity natural image synthesis. In: International Conference on Learning Representations (2019), <https://openreview.net/forum?id=B1xsqj09Fm> [1](#)
4. Chattopadhyay, A., Sarkar, A., Howlader, P., Balasubramanian, V.N.: Gradcam++: Generalized gradient-based visual explanations for deep convolutional networks. In: 2018 IEEE winter conference on applications of computer vision (WACV). pp. 839–847. IEEE (2018) [3](#), [4](#)
5. Chen, Y., Liu, S., Wang, X.: Learning continuous image representation with local implicit image function. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8628–8638 (2021) [2](#), [3](#)
6. Cherepkov, A., Voynov, A., Babenko, A.: Navigating the gan parameter space for semantic image editing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3671–3680 (2021) [4](#)
7. 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) [3](#)
8. Choi, Y., Uh, Y., Yoo, J., Ha, J.W.: Stargan v2: Diverse image synthesis for multiple domains. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8188–8197 (2020) [3](#)
9. 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) [1](#)
10. Härkönen, E., Hertzmann, A., Lehtinen, J., Paris, S.: Ganspace: Discovering interpretable gan controls. *Advances in Neural Information Processing Systems* **33**, 9841–9850 (2020) [2](#), [3](#), [4](#), [10](#)
11. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (July 2017) [2](#), [3](#)
12. 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) [4](#), [8](#), [9](#), [10](#), [13](#)
13. 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) [1](#), [3](#), [8](#), [9](#), [10](#), [13](#)
14. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4681–4690 (2017) [2](#), [3](#)

15. Li, Z., Jiang, R., Aarabi, P.: Continuous face aging via self-estimated residual age embedding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15008–15017 (2021) [3](#)
16. Liu, M., Ding, Y., Xia, M., Liu, X., Ding, E., Zuo, W., Wen, S.: Stgan: A unified selective transfer network for arbitrary image attribute editing. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3673–3682 (2019) [3](#)
17. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2337–2346 (2019) [1](#), [3](#)
18. Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks (2016) [2](#), [4](#)
19. Saha, R., Duke, B., Shkurti, F., Taylor, G., Aarabi, P.: Loho: Latent optimization of hairstyles via orthogonalization. In: CVPR (2021) [3](#)
20. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE international conference on computer vision. pp. 618–626 (2017) [3](#), [4](#), [7](#)
21. Shen, Y., Yang, C., Tang, X., Zhou, B.: Interfacegan: Interpreting the disentangled face representation learned by gans. *IEEE transactions on pattern analysis and machine intelligence* (2020) [2](#), [3](#), [4](#), [6](#), [9](#), [10](#)
22. Smilkov, D., Thorat, N., Kim, B., Viégas, F., Wattenberg, M.: Smoothgrad: removing noise by adding noise. arXiv preprint arXiv:1706.03825 (2017) [4](#)
23. Sundararajan, M., Taly, A., Yan, Q.: Axiomatic attribution for deep networks. In: International conference on machine learning. pp. 3319–3328. PMLR (2017) [4](#)
24. Upchurch, P., Gardner, J., Pleiss, G., Pless, R., Snavely, N., Bala, K., Weinberger, K.: Deep feature interpolation for image content changes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7064–7073 (2017) [2](#), [4](#)
25. 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) [4](#)
26. 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) [1](#), [3](#)
27. Wu, Y., Donahue, J., Balduzzi, D., Simonyan, K., Lillicrap, T.: Logan: Latent optimisation for generative adversarial networks. arXiv preprint arXiv:1912.00953 (2019) [4](#), [6](#)
28. 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) [2](#), [3](#), [4](#), [9](#), [10](#), [11](#)
29. Yang, C., Shen, Y., Zhou, B.: Semantic hierarchy emerges in deep generative representations for scene synthesis. *International Journal of Computer Vision* **129**(5), 1451–1466 (2021) [2](#), [4](#)
30. Yeh, R.A., Chen, C., Yian Lim, T., Schwing, A.G., Hasegawa-Johnson, M., Do, M.N.: Semantic image inpainting with deep generative models. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5485–5493 (2017) [2](#), [3](#)

31. Yu, F., Seff, A., Zhang, Y., Song, S., Funkhouser, T., Xiao, J.: Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. arXiv preprint arXiv:1506.03365 (2015) [8](#), [9](#), [10](#), [11](#), [13](#), [14](#)
32. Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., Huang, T.S.: Free-form image inpainting with gated convolution. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4471–4480 (2019) [2](#), [3](#)
33. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2223–2232 (2017) [2](#), [3](#)