

[Re] GANSpace: Discovering Interpretable GAN Controls

Vishnu Asutosh Dasu^{1, ID} and Midhush Manohar T.K.^{2, ID}

¹TCS Research and Innovation, Bangalore, India – ²Akamai Technologies, Inc., Bangalore, India

Edited by

Koustuv Sinha,
Sharath Chandra Raparthy

Reviewed by

Anonymous Reviewers

Received

04 February 2022

Published

23 May 2022

DOI

10.5281/zenodo.6574645

Reproducibility Summary

Scope of Reproducibility

The authors introduce a novel approach to analyze Generative Adversarial Networks (GANs) and create interpretable controls for image manipulation and synthesis. This is done by identifying important latent directions based on Principal Component Analysis (PCA) applied either in the latent space or the feature space. We aim to validate the claims and reproduce the results of the original paper.

Methodology

The code that was provided by the authors in Pytorch was reimplemented in Tensorflow 1.x for the pretrained StyleGAN and StyleGAN2 architectures. This was done with the help of the APIs provided by the original authors of these models.

The experiments were run on a laptop with an Intel(R) Core(TM) i7-8750H CPU @ 2.20GHz processor, 16GB RAM, NVIDIA GeForce GTX 1060 with Max-Q Design (6GB VRAM) GPU, and Ubuntu 18.04.5 LTS.

Results

We were able to reproduce the results and verify the claims made by the authors for the StyleGAN and StyleGAN2 models by recreating the modified images, given the seed and other configuration parameters. Additionally, we also perform our own experiments to identify new edits and extend the truncation trick to images generated using StyleGAN.

What was easy

The paper provides detailed explanations for the different mathematical concepts that were involved in the proposed method. This, augmented with a well-structured and documented code repository, allowed us to understand the major ideas in a relatively short period of time. Running the experiments using the original codebase was straightforward and highly efficient as well, as the authors have taken additional steps to employ batch processing wherever possible.

Copyright © 2022 V.A. Dasu and M.M. T.K., released under a Creative Commons Attribution 4.0 International license.

Correspondence should be addressed to Vishnu Asutosh Dasu (vishnu98dasu@gmail.com)

The authors have declared that no competing interests exist.

Code is available at <https://github.com/midsterx/ReGANSpace> – DOI 10.5281/zenodo.6511501. – SWH

swh:1:dir:4dc4de7856350a4671d97840c5f9ae013c275112.

Open peer review is available at <https://openreview.net/forum?id=BtZVD2f7n0F>.

What was difficult

Originally we were attempting to recreate identical images with zero delta in the RGB values. However, due to differences in the random number generators between PyTorch-CPU, PyTorch-GPU and Numpy, the random values were not the same even with the same seed. This resulted in minute differences in the background artifacts of the generated images. Additionally, there is a lack of open source Tensorflow 1.x APIs to access the intermediate layers of the BigGAN model. Due to time constraints, we were unable to implement these accessors and verify the images that the authors of GANSpace created using BigGAN.

Communication with original authors

While conducting our experiments, we did not contact the original authors. The paper and codebase were organized well and aided us in effectively reproducing and validating the authors' claims.

1 Introduction

A Generative Adversarial Network (GAN)[1] is a machine learning framework where two neural networks, the discriminator and the generator, compete with each other in a zero-sum game. The generator tries to trick the discriminator into believing that artificially generated samples belong to real data.

GANs have proven to be powerful image synthesis tools and are capable of producing high quality images. However, they provide little control over the features of the generated image. Existing solutions[2] that add user control over the generated images require expensive supervised training on latent vectors.

GANSpace[3] proposes a simple technique to discover interpretable GAN controls in an unsupervised manner. This is done by identifying important latent directions based on Principal Component Analysis (PCA) applied either on the latent space or the feature space. The author's experiments on StyleGAN[4], StyleGAN2[5] and BigGAN512-deep[6] demonstrate that layer-wise decomposition of PCA directions leads to many interpretable controls, which affect both low and high level attributes of the output image.

2 Scope of reproducibility

For our reproduction study, we aim to validate the effectiveness of the proposed technique in offering powerful interpretable controls on the output images in an unsupervised manner.

The following claims of the paper have been verified and tested successfully:

- PCA can be used to highlight important directions in the GAN's latent space.
- The GAN's output can be controlled easily in an unsupervised fashion.
- The earlier components control the higher-level aspects of an image, while the later directions primarily affect the minute details.
- Random directions do not yield meaningful decompositions as compared to the principal components identified using PCA.

3 Methodology

The principal components[7] of a collection of points in real coordinate space are a sequence of p unit vectors, where the i^{th} vector is a direction of the line that best fits the data while being orthogonal to the remaining $i - 1$ vectors. Principal Component Analysis (PCA) is an unsupervised algorithm used to compute the principal components and perform a change of basis of the data using one or more of the computed components, increasing the interpretability of the data while minimizing its information loss[8]. It is commonly used in exploratory data analysis and for dimensionality reduction when dealing with high-dimensional noisy data. The authors of GANSpace propose a technique for identifying interpretable controls in an unsupervised fashion on pretrained GANs using PCA. Specifically, they show that layer-wise perturbations along the principal components generated using PCA on the latent space of StyleGAN based networks can be used to generate human-interpretable transformations on the synthesized images.

Mathematically, a GAN can be expressed as a neural network $G(\mathbf{z})$ that generates an image $I : \mathbf{z} \sim p(\mathbf{z}), I = G(\mathbf{z})$. Here, $p(\mathbf{z})$ is a probability distribution from which the latent vector \mathbf{z} is sampled. The network $G(\mathbf{z})$ can be further decomposed into L intermediate layers $G_1 \dots G_L$. In the StyleGAN/StyleGAN2 models, the input to the first layer is a constant \mathbf{y}_0 . The output and input to the remaining layers is computed as:

$$y_i = G_i(y_{i-1}, \mathbf{w}), \text{ where } \mathbf{w} = M(\mathbf{z}) \quad (1)$$

M is an 8-layer multilayer perceptron which is a non-linear function of \mathbf{z} . The number of layers L depends on the resolution of the generated image. At each layer, the generated image is upsampled by a factor of 2.

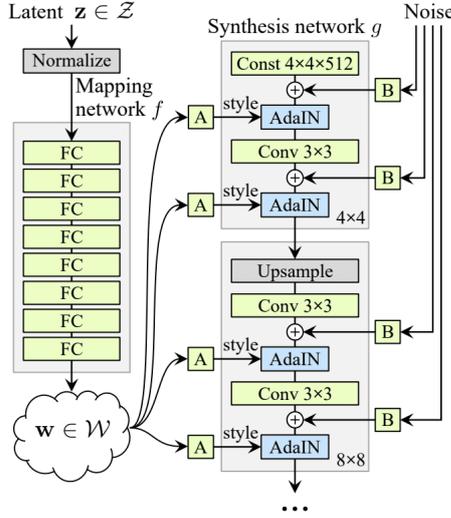


Figure 1. Architecture of StyleGAN[4]

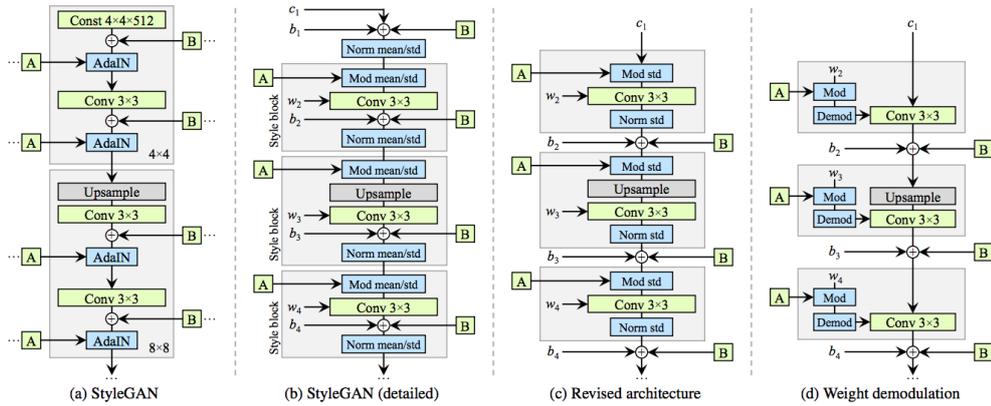


Figure 2. Architecture of StyleGAN2[5]

The images generated by StyleGAN and StyleGAN2 can be controlled by identifying the principal axes of $p(\mathbf{w})$, which is the probability distribution of the output of the mapping network M . First, we sample N latent vectors $\mathbf{z}_{1:N}$ and compute the corresponding $\mathbf{w}_i = M(\mathbf{z}_i)$. The PCA of these $\mathbf{w}_{1:N}$ values gives us the basis \mathbf{V} for \mathcal{W} . The output attributes of a new image given by \mathbf{w} can then be controlled by varying the PCA coordinates of \mathbf{x} before feeding them into the synthesis network:

$$\mathbf{w}' = \mathbf{w} + \mathbf{V}\mathbf{x} \quad (2)$$

Each entry x_k of \mathbf{x} is a separate control parameter which can be modified to update the desired attributes of the output image.

We follow the same notation used by the authors to denote edit directions in this report. $E(\mathbf{v}_i, j - k)$ means moving along component \mathbf{v}_i from layers j to k . Identifying specific edits, for example "changing the color of a car", is done via exploratory analysis using

a trial-and-error method. The authors have created a GUI-based application for this purpose.

3.1 Model descriptions

We use NVIDIA’s official implementation of StyleGAN¹ and StyleGAN2² models. The original code uses a PyTorch/NumPy implementation of StyleGAN and StyleGAN2 which creates a PyTorch model and copies the weights from NVLabs’ implementations which are in Tensorflow. However, we directly use the NVLabs’ APIs with NumPy and make changes to the official GANSpace codebase to support the same.

3.2 Datasets

The experiments in the paper were performed using the FFHQ, LSUN Car, CelebA-HQ, Wikiart, Horse and Cat datasets. The official Tensorflow implementation of StyleGAN contains links to download pretrained models on FFHQ, LSUN Car, Wikiart, Horse and Cat. The models trained on Wikiart were downloaded from awesome-pretrained-stylegan³. In addition to the datasets used by the authors, we also perform our own experiments on the Beetles dataset which was downloaded from awesome-pretrained-stylegan2⁴.

3.3 Experimental setup

All the experiments were conducted on a laptop with an Intel(R) Core(TM) i7-8750H CPU @ 2.20GHz processor, 16GB RAM, NVIDIA GeForce GTX 1060 with Max-Q Design (6GB VRAM) GPU, and Ubuntu 18.04.5 LTS. The generated images from our experiments were evaluated visually to determine whether the edits were working as expected.

4 Results

We were able to reproduce the results and verify the claims (mentioned in Section 2) made by the authors for the StyleGAN and StyleGAN2 models by recreating the modified images, given the configuration parameters. Additionally, we also perform our own experiments to provide additional results that validate the effectiveness of the technique employed by GANSpace.

4.1 Effectiveness of PCA



Figure 3. Sequences of image edits performed using control discovered with StyleGAN2 cars: “Initial Image” → “Change Color” → “Add Grass” → “Rotate” → “Change Type”

¹<https://github.com/NVLabs/stylegan>

²<https://github.com/NVLabs/stylegan2>

³<https://github.com/justinpinkney/awesome-pretrained-stylegan>

⁴<https://github.com/justinpinkney/awesome-pretrained-stylegan2>

Figure 3 highlights the effectiveness of PCA on changing the low and high level attributes of the image. We are able to control object shape, colour and pose as well as nuanced landscape attributes.

The edit directions corresponding to each of the edits are: $E(\mathbf{v}_{22}, 9-10)$ ("Change Color"), $E(\mathbf{v}_{11}, 9-10)$ ("Add Grass"), $E(\mathbf{v}_0, 0-4)$ ("Rotate") and $E(\mathbf{v}_{16}, 3-5)$ ("Change type").

4.2 Unsupervised vs. Supervised methods

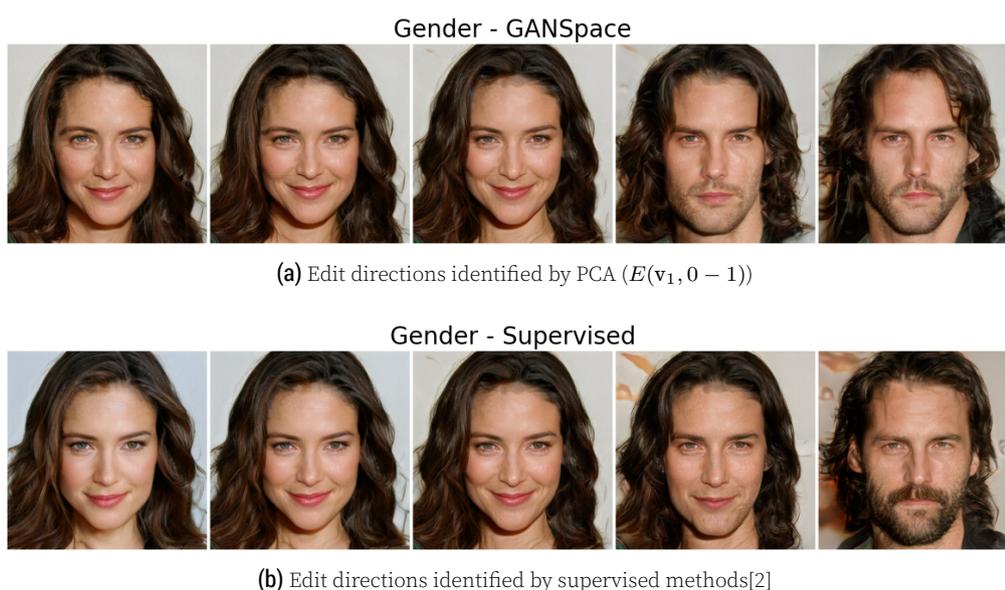


Figure 4. Comparison of edits using unsupervised and supervised methods

Previous methods for finding interpretable directions in GAN latent spaces require external supervision, such as labeled training images or pretrained classifiers. GANSpace, on the other hand, automatically identifies variations intrinsic to the model without supervision. This has been validated using the CelebA-HQ Faces dataset by comparing the edit directions found through PCA to those found in previous works using supervised methods.

Figure 4 shows that comparable edits can be obtained in a completely unsupervised fashion. Additionally, GANSpace can be used to identify new edits which have not been previously demonstrated. Supervised methods are not viable for this task as supervising each new edit would be costly. It is also difficult to know in advance which edits are even possible in supervised approaches.

4.3 PCA components vs. Random directions

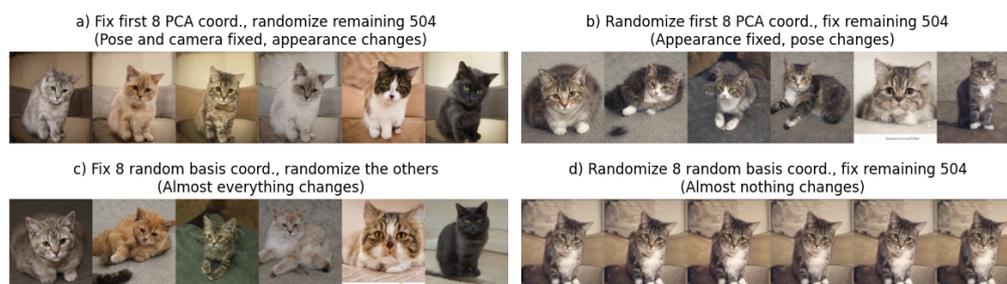


Figure 5. Illustration of the significance of the principal components as compared to random directions in the intermediate latent space of StyleGAN2.

The original authors claim that the earlier PCA components primarily control the geometry and other high-level aspects (pose and style), while the lower components capture minute details. Additionally, they claim that fixing and randomizing randomly-chosen directions do not yield PCA-like meaningful decompositions, thus showing the importance of identifying good directions using PCA. This has been illustrated in Figure 5, where different subsets of principal coordinates and random coordinates are randomized while keeping the latent vector constant. In Figure 5a, the first eight principal coordinates $\mathbf{x}_{0:7}$ are fixed and the remaining 504 coordinates $\mathbf{x}_{8:512}$ are randomized. This changes the background and appearance of the cat while keeping the cat’s pose and camera angle constant. Conversely, Figure 5b shows that fixing the last 504 coordinates and randomizing the first eight yields images where the camera and orientation vary, but the color and appearance are held roughly constant. Figure 5c and Figure 5d shows the results of the same process applied to random directions. The images illustrate that any given 8 directions have no distinctive effect on the output.

4.4 Additional results not present in the original paper

New edits – We identify new edits on the Stylegan2 Beetles dataset. Edit $E(\mathbf{v}_2, 0 - 17)$, referred to as “Patterns”, adds a pattern on the shell of the beetle as well as increasing the overall size of the beetle. The generated pattern varies depending on the seed used to sample \mathbf{w} .

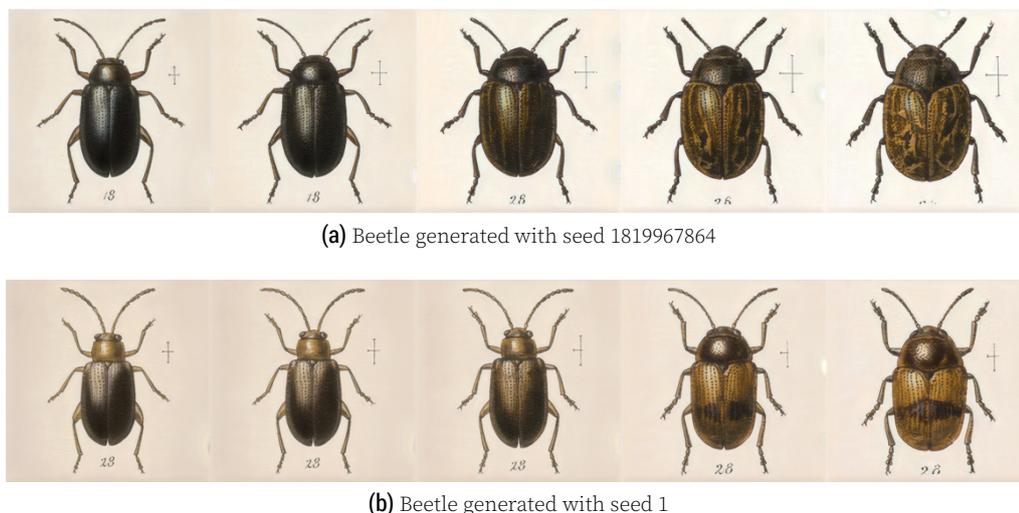


Figure 6. "Patterns" edit applied on the output images of StyleGAN2 Beetles

Truncation Trick on StyleGAN – The "Truncation Trick" is a procedure applied to the latent vectors to improve the quality of the generated images at the expense of variety in the images. It does this by sampling the latent vectors from a truncated distribution that is closer to the average of the latent vectors sampled during training, thereby reducing the variance of the latent vectors used during inference. The authors of [6] show that using the truncation trick improves the Fréchet Inception Distance (FID) and Inception Score (IS).

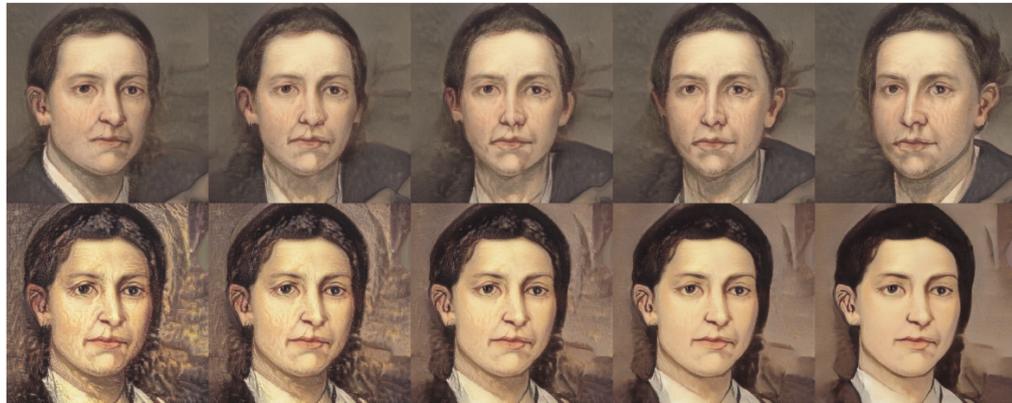
In the StyleGAN/StyleGAN2 models, the truncation trick is applied on the latent space \mathbf{w} , which is the output of the mapping network M . During the training process, a running average \mathbf{w}_{avg} of the latents is computed. Later, the latents sampled during inference are truncated to lie close to \mathbf{w}_{avg} . Equation 3 shows the truncation process on StyleGAN/StyleGAN2 models:

$$\mathbf{w}' = \mathbf{w}_{avg} + \psi(\mathbf{w} - \mathbf{w}_{avg}) \quad (3)$$

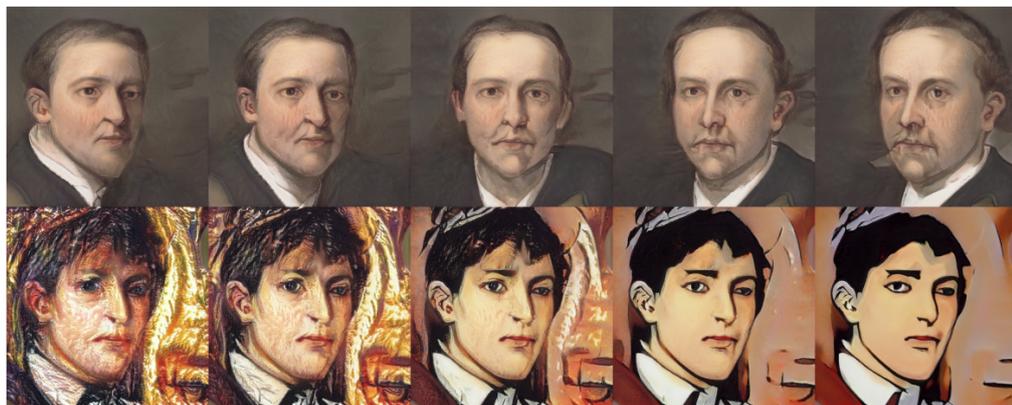
During our experiments, we noticed that the original authors use the truncation trick on images generated using StyleGAN2 to reduce the number of artifacts. However, this is not enabled for StyleGAN images. We found that enabling truncation while applying edits on StyleGAN images improved their quality as well. We demonstrate this using the Wikiart dataset through the "Head Rotation" ($E(\mathbf{v}_7, 0-1)$) and "Simple Strokes" ($E(\mathbf{v}_9, 8-14)$) edits. In Figure 7, we can see that the generated faces contain less noise and artifacts when the truncation trick is used. For example, the lower half of the person's face in the "Head Rotation" image does not contain as much noise as their counterpart which does not employ the truncation trick. Here, we can also observe the change in the generated images as truncation ψ is decreased a lower value of 0.25. This truncates the sampled latents to lie very close to the average and results in images that look very similar to each other. If truncation ψ is set to 0, then according to Equation 3, we can see that the truncated latent \mathbf{w}' is always equal to \mathbf{w}_{avg} .

5 Discussion

After performing our experiments, we feel that the results justify the claims of the paper. This is further bolstered by the fact that the proposed method worked on different datasets which were not covered by the original authors.



(a) "Head Rotation" and "Simple Strokes" edits on StyleGAN Wikiart with truncation psi set to 0.25



(b) "Head Rotation" and "Simple Strokes" edits on StyleGAN Wikiart with truncation psi set to 0.7



(c) "Head Rotation" and "Simple Strokes" edits on StyleGAN Wikiart without truncation psi

Figure 7. Quality of images generated by StyleGAN before and after applying the "truncation trick".

5.1 What was easy

The paper provides detailed explanations for the different mathematical concepts that were involved in the proposed method. This, augmented with a well-structured and documented code repository, allowed us to understand and verify the major ideas in a relatively short period of time. Additionally, the paper provided a lot of examples on various datasets to demonstrate exactly how their algorithm works. The authors ensured that all the figures in the paper had accompanying code to recreate them.

NVIDIA's implementation of StyleGAN and StyleGAN2 provided access to well written API's which we could integrate easily into the author's codebase.

5.2 What was difficult

While running our experiments, we noticed that there was a small difference in the RGB values of the recreated images. This was due to the difference in the random values generated by PyTorch-CPU, PyTorch-GPU and Numpy random number generators even when seeded with the same seed. The noise variables in the StyleGAN networks were not identical because of this. This resulted in minute differences in background artifacts of the images.

Python Library	Random Number
PyTorch 1.3.1 (CPU)	0.3367
PyTorch 1.3.1 (GPU)	0.1940
Numpy 1.20.1	0.49671415

Table 1. Random values generated using different Python libraries seeded with 42

We were not able to replicate the author's experiments on BigGAN512-deep due to time constraints.

5.3 Communication with original authors

While conducting our experiments, we did not contact the original authors. The paper and codebase were organized well and aided us in effectively reproducing and validating the authors' claims.

References

1. I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. C. Courville, and Y. Bengio. "Generative Adversarial Nets." In: *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*. Ed. by Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger. 2014, pp. 2672–2680. URL: <https://proceedings.neurips.cc/paper/2014/hash/5ca3e9b122f61f8f06494c97b1afccf3-Abstract.html>.
2. Y. Shen, J. Gu, X. Tang, and B. Zhou. "Interpreting the Latent Space of GANs for Semantic Face Editing." In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*. IEEE, 2020, pp. 9240–9249. doi: 10.1109/CVPR42600.2020.00926. URL: <https://doi.org/10.1109/CVPR42600.2020.00926>.
3. E. Härkönen, A. Hertzmann, J. Lehtinen, and S. Paris. "GANSpace: Discovering Interpretable GAN Controls." In: *Proc. NeurIPS*. 2020.
4. T. Karras, S. Laine, and T. Aila. "A Style-Based Generator Architecture for Generative Adversarial Networks." In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*. Computer Vision Foundation / IEEE, 2019, pp. 4401–4410. doi: 10.1109/CVPR.2019.00453. URL: http://openaccess.thecvf.com/content%5C_CVPR%5C_2019/html/Karras%5C_A%5C_Style-Base%5C_Generator%5C_Architecture%5C_for%5C_Generative%5C_Adversarial%5C_Networks%5C_CVPR%5C_2019%5C_paper.html.

5. T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila. "Analyzing and Improving the Image Quality of StyleGAN." In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*. IEEE, 2020, pp. 8107–8116. doi: 10.1109/CVPR42600.2020.00813. URL: <https://doi.org/10.1109/CVPR42600.2020.00813>.
6. A. Brock, J. Donahue, and K. Simonyan. "Large Scale GAN Training for High Fidelity Natural Image Synthesis." In: *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL: <https://openreview.net/forum?id=B1xsqj09Fm>.
7. H. Hotelling. "Analysis of a complex of statistical variables into principal components." In: *Journal of educational psychology* 24.6 (1933), p. 417.
8. I. T. Jolliffe and J. Cadima. "Principal component analysis: a review and recent developments." In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 374.2065 (2016), p. 20150202.