

This CVPR 2021 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Navigating the GAN Parameter Space for Semantic Image Editing

Anton Cherepkov Yandex, Russia Moscow Institute of Physics and Technology cherepkov.ayu@phystech.edu Andrey Voynov Yandex, Russia an.voynov@yandex.ru

Artem Babenko Yandex, Russia HSE University

artem.babenko@phystech.edu



Abstract

Generative Adversarial Networks (GANs) are currently an indispensable tool for visual editing, being a standard component of image-to-image translation and image restoration pipelines. Furthermore, GANs are especially advantageous for controllable generation since their latent spaces contain a wide range of interpretable directions, well suited for semantic editing operations. By gradually changing latent codes along these directions, one can produce impressive visual effects, unattainable without GANs.

In this paper, we significantly expand the range of visual effects achievable with the state-of-the-art models, like StyleGAN2. In contrast to existing works, which mostly operate by latent codes, we discover interpretable directions in the space of the generator parameters. By several simple methods, we explore this space and demonstrate that it also contains a plethora of interpretable directions, which

Contact author: Andrey Voynov, an.voynov@yandex.ru

are an excellent source of non-trivial semantic manipulations. The discovered manipulations cannot be achieved by transforming the latent codes and can be used to edit both synthetic and real images. We release our code and models and hope they will serve as a handy tool for further efforts on GAN-based image editing.

1. Introduction

Generative Adversarial Networks (GANs) [8] have revolutionized image processing research, significantly pushing the boundaries of machine learning for image enhancement and visual editing. Different research lines currently exploit GANs in several principled ways, e.g. using them as an implicit learnable objective [20, 17, 13, 5, 18, 36, 37], employing them as high-quality image priors [35, 2, 9, 22, 23], manipulating their internal representations for visual editing purposes [3, 6]. Furthermore, the GAN latent spaces often

https://github.com/yandex-research/navigan

encode human-interpretable concepts [26, 28, 7, 14, 25, 29, 10, 24, 30], which makes GANs the dominant paradigm for controllable generation.

Since the seminal paper [26], which has demonstrated the semantic arithmetic of latent vectors in GANs, plenty of methods to discover interpretable directions in the GAN latent spaces have been developed [26, 28, 7, 14, 25, 29, 10, 24, 30]. These methods successfully identify such directions across different GAN models and hold great potential for effective image editing. These days, many impressive visual effects can be achieved by simply moving the image latent codes along these directions.

Our paper demonstrates that a large number of exciting non-trivial visual effects can be produced by gradually *modifying the GAN parameters rather than the latent codes*. In more detail, we show that the GAN parameter space also contains a plethora of directions, corresponding to interpretable image manipulations. Moreover, we describe simple domain-agnostic procedures that discover such directions in an unsupervised fashion. By extensive experiments, we confirm that the discovered visual effects are substantially new and cannot be achieved by the latent code manipulations. Overall, our findings significantly expand the arsenal of GAN-based image editing techniques.

To sum up, our contributions are the following:

- We propose to use the interpretable directions in the space of the generator parameters for semantic editing. Our approach differs from existing works, which operate by the latent codes or the intermediate GAN activations. Our findings demonstrate that remarkable visual effects can be achieved by slightly changing the GAN parameters.
- We develop the methods to discover such directions. The proposed methods are both effective and fast and can work on a single GPU.
- We confirm that the discovered directions are qualitatively new and correspond to semantic manipulations, which existing methods cannot produce.

2. Related work

This section describes the typical ways to exploit GANs for visual editing purposes and positions our work with respect to existing literature.

GAN discriminators as learnable training objectives. In most image generation tasks, it is challenging to explicitly define an objective term that would enforce the produced images' realism. Therefore, many existing approaches employ GANs as implicit learnable objectives, making the output images indistinguishable from the real ones. This approach has become de-facto standard for a wide range of image processing tasks, including superresolution [17], texture transfer [20], image-to-image translation [13, 5, 18, 36, 37].

GANs as high-quality image priors. As the state-ofthe-art GAN models are high-quality approximations of the real image manifold, several recent methods use these models as "hard" priors. In this case, the method's outputs are produced by a large-scale pretrained GAN, thus, the image processing task is reduced to the optimization problem in the GAN latent space. This paradigm has been successfully used for super-resolution [22], visual editing [35, 2], image restoration tasks [9, 23], e.g. colorization and inpainting.

Latent manipulations in GANs. The prior literature has empirically shown that the GAN latent spaces are endowed with human-interpretable vector space arithmetic [26, 28, 7, 14, 29]. E.g., in GANs trained on face images, their latent spaces possess linear directions corresponding to adding smiles, glasses, gender swap [26, 28]. Since such interpretable directions provide a straightforward route to powerful image editing, their discovery currently receives much research attention. A line of recent works [7, 28] employs explicit human-provided supervision to identify interpretable directions in the latent space. For instance, [28] use the classifiers pretrained on the CelebA dataset [21] to predict certain face attributes. These classifiers are then used to produce pseudo-labels for the generated images and their latent codes. Based on these pseudo-labels, the separating hyperplane is constructed in the latent space, and a normal to this hyperplane becomes a direction, controlling the corresponding attribute. Another work [7] solves the optimization problem in the latent space, maximizing the score of the pretrained model, which predicts image memorability. The result of this optimization is the direction "responsible" for the increase of memorability. Two self-supervised works [14, 25] seek the vectors in the latent space that correspond to simple image augmentations such as zooming or translation. Finally, a bunch of recent methods [29, 10, 24] identify interpretable directions without any form of supervision. [29] learns a set of directions that can be easily distinguished by a separate classification model based on two samples, produced from the original latent codes and their versions shifted along the particular direction. [24] learns the directions by minimizing the sum of squared offdiagonal terms of the generator's Hessian matrix. Another approach, [10], demonstrates that interpretable directions often correspond to the principal components of the activations from the first layer of the generator network. A very recent work [30] aims to unify all unsupervised approaches above. [30] claims that these approaches can be treated as special cases of computing the spectrum of the Hessian for the LPIPS model [33] with respect to latent coordinates. Intuitively, [30] assumes that the interpretable directions correspond to the largest perceptual changes, quantified by the LPIPS model. Our work is partially inspired by the ideas from the latent space exploration, however, we explore the space of the generator parameters rather than the latent space.

Manipulating GAN activations. It was shown [3] that the intermediate activations in the GAN generator often correspond to various semantic concepts. [3] exploits this by controlling the presence or absence of objects at given positions, guided by supervision from a pretrained semantic segmentation model. Another unsupervised approach [6] identifies semantically meaningful activations via simple kmeans clustering and makes local edits for the generated images. In contrast to these works, we operate on the generator weights rather than the activations tensors.

Rewriting a deep generative model. The closest work to ours is [1], which shows that a user-specified editing operation can be achieved by changing the generator weights. In a nutshell, [1] solves the optimization problem that seeks the shift in the generator's parameter space, such that the modified generator would produce the images after a needed editing operation. In contrast, our approach does not require specifying operations a priori and discovers a lot of unexpected non-trivial visual effects, some of which are impossible to formulate explicitly.

3. Methods

3.1. Preliminaries

Assume that we have a pretrained GAN generator G_{θ} , which maps the samples $z \sim \mathcal{N}(0,\mathbb{I})$ from the latent space \mathbb{R}^d into the image space, $\mathbb{R}^{W \times H \times 3} G : z \rightarrow I$, and $\theta \in \Theta$ denotes the set of the generator's parameters. Our goal, then, is to learn a set of vectors $\xi_1, \ldots, \xi_K \in \Theta$, such that changing the generator parameters along these vectors effectively performs continious semantic editing operations. More formally, the mappings $G_{\theta}(z) \rightarrow G_{\theta+t\xi_k}(z), k=1..K$ have to correspond to interpretable visual effects, consistent for all latent codes $z \sim \mathcal{N}(0, \mathbb{I})$. Here $t \in [-T, T]$ is a shift magnitude, which controls the degree of the corresponding visual effect.

Note that the problem statement above resembles the established problem of learning the interpretable latent controls addressed in [28, 14, 29, 10, 30]. However, the existing works perform the editing operations by shifting the latent codes with a fixed generator. In contrast, we operate in a much higher dimensional space of the generator's parameters Θ and do not change the latent codes when editing. In the next subsections, we describe the methods that discover the interpretable directions in Θ .

3.2. Optimization-based approach

The first approach is inspired by the recent unsupervised technique [29], which discovers the interpretable latent di-

rections. Intuitively, [29] assumes that interpretable directions are the ones that are easy to distinguish from each other, observing only the results of the corresponding image manipulations. Here, we build on this intuition to explore the parameter space Θ . While many parts of our protocol have been used in [29], we still present them for selfcontainedness.

Figure 1 demonstrates our training scheme. It includes two learnable modules:

- 1. Direction matrix $\Xi = [\xi_1, \dots, \xi_K] \in \mathbb{R}^{dim(\Theta) \times K}$.
- 2. Reconstructor R, which recieves the results of editing operations represented by pairs of the form $\{G_{\theta}(z); G_{\theta+t\xi_k}(z)\}$, and predicts both k and t from its input. More formally, R is a function $(I_1, I_2) \rightarrow (\{1, \ldots, K\}, \mathbb{R})$ parameterized by a deep convolutional network.



Figure 1. Our learning protocol that discover the interpretable shifts in the space of the generator's parameters Θ . A training sample consists of two images, produced by the generators with original and shifted parameters. The images are given to a reconstructor R that predicts a direction index k and a shift magnitude t.

The learning is performed by minimizing the expected reconstructor's prediction error:

$$\min_{\substack{[\xi_1,\dots,\xi_K],R\\k\sim\mathcal{U}\{1,K\}\\t\sim\mathcal{U}[-T,T]}} \mathbb{E}\left[L_{cl}(k,\widehat{k}) + \lambda L_r(t,\widehat{t})\right]$$
(1)

where \hat{k} and \hat{t} denote the reconstructor's output:

$$(\widehat{k};\widehat{t}) = R\left(G_{\theta}(z); G_{\theta+t\xi_k}(z)\right) \tag{2}$$

For the classification objective term $L_{cl}(\cdot, \cdot)$ we use cross-entropy, and for the regression term $L_r(\cdot, \cdot)$ mean absolute error is used. Since all the components of our scheme are differentiable, it can be optimized jointly by stochastic gradient descent.

Reducing the dimensionality of the optimization space. State-of-the-art generators (e.g., StyleGAN2) typically have millions of parameters θ , making it infeasible to learn the full $\Xi \in \mathbb{R}^{dim(\Theta) \times K}$ matrix. To simplify the optimization, in experiments, we minimize (1) considering only the shifts ξ_1, \ldots, ξ_K applied to a particular generator's layer, with all other parameters being fixed. Namely, we add a shift to the convolutional kernel of one StyleGAN2 block (see Figure 2). This design choice is motivated by the evidence from [3] that different generator layers typically affect the different aspects of the image.



Figure 2. The additive shift ξ is added to the convolutional kernel weight in the StyleGAN2 demodulation block.

However, only considering a single layer does not fully solve the problem of high dimensionality. To reduce the optimization space even further, we perform the following. Inspired by the recent work on few-shot GAN adaptation [27], we compute the SVD decomposition of the chosen convolutional layer, flattened to a 2D matrix, as in [27]. Then we optimize over the shifts ξ_1, \ldots, ξ_K applied only to the singular values of the diagonal matrix from SVD. Such parametrization makes the optimization problem feasible but is also expressive enough to represent various visual effects, as shown experimentally. While we tried some alternative parameterizations, they typically resulted in worse visual performance.

Practical details.

- 1. In all experiments we use the Resnet-18 architecture [11] for the reconstructor, with two heads, predicting t and k respectively. We downscale the input images to the resolution 256×256 to decrease memory consumption.
- 2. Assuming that SVD of a flattened convolutional kernel equals $U \cdot \text{diag}(\sigma_1, \ldots, \sigma_n) \cdot V$, we apply an additive shift $\xi = (\xi^{(1)}, \ldots, \xi^{(n)})$ to the singular values, mapping $\text{diag}(\sigma_1, \ldots, \sigma_n)$ to $\text{diag}(\sigma_1 + \xi^{(1)}, \ldots, \sigma_n + \xi^{(n)})$. We also normalize the shift ξ to a unit length to avoid parameter explosion.
- 3. In all experiments we use K=64, $\lambda=2.5 \cdot 10^{-3}$ and T=3500. For optimization, we use Adam with a constant learning rate 0.0001 and perform 10^5 learning iterations with a batch size 32.

3.3. Spectrum-based approach

Our second approach originates from an alternative premise presented in [30], which claims that the interpretable directions in the latent space should correspond to as large perceptual changes of the generated images as possible. [30] formalizes this intuition by the following. They calculate the latent directions that are the eigenvectors of the Hessian of the LPIPS model [33] computed with respect to the latent codes z. [30] also proposes an efficient way to compute top-k eigenvectors, corresponding to the largest

eigenvalues, avoiding the explicit computation of the whole Hessian, which is impractical.

Our second approach exploits the same intuition as in [30] but operates with the Hessian of LPIPS computed with respect to the generator's parameters. For completeness, we briefly describe the main steps to compute the Hessian's top-eigenvectors, though they are almost the same as in [30]. Let us denote by $d(\cdot, \cdot)$ the LPIPS model, which is shown to capture the perceptual distance between two images [34]. Let θ be the weights of the pretrained GAN generator. Then, we consider the quantity $\mathbb{E}_z d^2(G_\theta(z), G_{\theta+\alpha}(z))$, which is the expected perceptual change induced by shifting the generator parameters by α . Assuming the LPIPS smoothness, we can write

$$\mathbb{E}_{z}d^{2}(G_{\theta}(z), G_{\theta+\delta\alpha}(z)) =$$

$$\mathbb{E}_{z}d^{2}(G_{\theta}(z), G_{\theta}(z)) + \frac{\partial d^{2}(G_{\theta}(z), G_{\theta+\alpha}(z))}{\partial \alpha}|_{\alpha=0} \cdot \delta\alpha +$$

$$\delta\alpha^{T} \cdot \frac{\partial^{2}d^{2}(G_{\theta}(z), G_{\theta+\alpha}(z))}{\partial \alpha^{2}}|_{\alpha=0} \cdot \delta\alpha + \bar{o}(\|\delta\alpha\|_{2}^{2}) \quad (3)$$

The first two terms from the right side of (3) are equal to zero since d^2 achieves its global minimum at $\alpha = 0$. Thus, we focus on the eighenvectors of the Hessian

$$H = \frac{\partial^2 \mathbb{E}_z d^2(G_\theta(z), G_{\theta+\alpha}(z))}{\partial \alpha^2}|_{\alpha=0}$$

For efficient computation, define the gradient function $g(a) = \frac{\partial \mathbb{E}_z d^2(G_{\theta}(z), G_{\theta+\alpha}(z))}{\partial \alpha}|_{\alpha=a}$. Following [30], we sample $v \sim \mathcal{N}(0, \mathcal{I})$ and iteratively update it by the rule:

$$v \mapsto \frac{g(\varepsilon v) - g(-\varepsilon v)}{2\varepsilon \|v\|} \tag{4}$$

where ε is a small constant set to 0.1. This process converges to the leading eigenvector of the hessian H [19]. Once the top k - 1 eigenvectors are found, we repeat this procedure, restricted to their orthogonal complementary, to obtain the k-th eigenvector. Namely, on every step of (4) we project v into the orthogonal complement of the already found top-(k - 1) eigenvectors.

In the experiments, we compute the top-64 eigenvectors. We approximate the expectation of the gradient function g with a minibatch of 512 randomly sampled $z \sim \mathcal{N}(0, \mathbb{I})$. We always perform ten iterations of (4) since it is sufficient for convergence. As in the optimization-based approach, we explore the parameter subspaces corresponding to different layers separately and compute the eigenvectors only with respect to parameters from a particular layer.

3.4. Hybrid scheme

We also propose a simple hybrid scheme that combines both the optimization-based and the spectrum-based approaches in a single procedure. First, we compute the top-k eigenvectors of the LPIPS Hessian v_1, \ldots, v_k , as described in Section 3.3. Then we solve the optimization problem (1) considering only the shifts ξ applied to v_1, \ldots, v_k . Informally, this hybrid scheme is equivalent to the optimizationbased approach that operates in the parameter subspace that captures the maximal perceptual differences in the generated images. For the hybrid scheme, we use the same hyperparameter values as in Section 3.2 except for T, which was set to 80 and batch size 16 in purpose to be able to run on a single Tesla V100. This change is due to the Hessian directions produce more noticeable visual effects compared to SVD.

3.5. Inspecting directions.

The methods from Section 3.2, Section 3.3, Section 3.4 all produce a set of directions ξ_1, \ldots, ξ_K . These K directions are then inspected manually by observing the image sequences $\{G_{\theta+t\xi_k}(z)|t \in [-T,T]\}$, for several latent codes z. If a particular sequence maintains the image realism and corresponds to interpretable manipulation, it then can be used in editing applications. Since K is typically small (e.g., 64), this procedure takes only several minutes for a single person.

4. Experiments

In this section, we evaluate the proposed techniques on the state-of-the-art StyleGAN2 generators [16] pretrained on the FFHQ [15], LSUN-Cars [32], LSUN-Horse [32] and LSUN-Church [32] datasets. Our goal is to demonstrate that the StyleGAN2 parameter space contains the linear directions that can be exploited for high-fidelity image editing.

4.1. Non-trivial visual effects

First, we present several visual effects induced by the discovered directions. Figure 3 and Figure 4 present three examples for each dataset. More examples are provided in supplementary and at the repository.

The typical manipulations affect scales and aspect ratio of certain object parts, like "Belly size" and "Nose length", or mutual arrangement of these parts, e.g., "Moving legs" and "Wheels rotation". In supplementary, we discuss that shifting the parameters of various generator layers results in different effects, i.e., earlier layers induce global geometrical transformations. The middle layers correspond to localized transformations of object parts, and the final layers are responsible for coloring manipulations.

4.2. Comparison

In this subsection, we compare different techniques of parameter space navigation, both qualitatively and quantitatively. Specifically, we compare the following methods:

https://github.com/yandex-research/navigan



A A A A

- Moving legs +



- Head size +



- Roof height +



- Trees +



Figure 3. Visual effects achieved by navigating the StyleGAN2 parameter space. Rows 1-3 correspond to the LSUN-Horse dataset and rows 4-6 correspond to the LSUN-Church dataset.

- **SVD** performs a smooth changing of a particular singular value in the SVD of a flattened generator layer. We include this method as a baseline inspired by the evidence from [27], which shows that varying the singular values often result in interpretable manipulations.
- **Optimization-based** method navigates Θ along the directions discovered as described in Section 3.2.
- Spectrum-based method navigates Θ along the directions discovered as described in Section 3.3.
- Hybrid method navigates Θ along the directions discovered as described in Section 3.4.



Figure 4. Visual effects achieved by navigating the StyleGAN2 parameter space. Rows 1-3 correspond to the FFHQ dataset and rows 4-6 correspond to the LSUN-Cars dataset.

To compare the methods, we apply all of them to the particular layer. Interestingly, in some cases, all methods discover directions, roughly corresponding to the same visual effect. Given such a direction, we can qualitatively compare the methods by observing how this direction modifies the images produced from the same latent noise z. The typical visualizations are presented on Figure 5 and Figure 6. These examples show that the **SVD** baseline dramatically corrupts the image realism. In contrast, the techniques described in Section 3 produce more clear and visible effects.

For a quantitative comparison, we perform the following experiment. For each of four methods, we consider a direction that corresponds to the "Wheel size" visual effect. This effect is chosen since it was identified by all the methods. Then, for the generators $G_{\theta+t\xi}(z)$, we plot the FID curves by varying the shift magnitudes t. For each t the FID is



Figure 5. Comparison of the "Thickness" direction discovered by different methods applied to the 3-rd conv layer of StyleGAN2.



Figure 6. Comparison of the "Wheel size" direction discovered by different methods applied to the 5-th conv layer of StyleGAN2.

always computed with $5 \cdot 10^4$ real and generated images. The plots are presented on Figure 7. The SVD baseline is inferior to three advanced methods from Section 3, while there is no clear winner among them. Overall, in practice we recommend to use all three proposed methods to explore the parameter space of a particular GAN since we observed that each method can reveal effects, missed by the others.

Since different methods have to use different scales of t values to

achieve the same degree of the visual effect, we first calibrate these scales manually to guarantee that the max values of t for all methods correspond to the same wheel size.



Figure 7. FID plots for the "Wheel size" visual effect produced by different methods.

4.3. Latent transformations cannot produce these visual effects

Here, we show that the manipulations induced by parameter shifts cannot be achieved by changing the latent codes. Specifically, we perform the following experiment. Given a shift ξ modifying a generator $G_{\theta} \to G_{\theta+\xi}$, we look for a latent shift h such that $G_{\theta}(z+h) = G_{\theta+\xi}(z)$. To this end, we solve the optimization problem

$$\min_{h} \mathbb{E}_{z} ||G_{\theta}(z+h) - G_{\theta+\xi}(z)||_{2}^{2}$$
(5)

For G we take the StyleGAN2 pretrained on the LSUN-Car dataset and ξ is the "Wheels size" direction. We optimize over h from the input \mathbb{Z} -space, the W-space and the space of the activation tensors before the modified convolutional layer. In all three settings, the optimization failed to obtain a latent shift that would properly mimic the visual effect. The results of the optimization are shown on Figure 8. Similar results were obtained for other generators and directions reported in the paper, see supplementary. Thus, our approach reveals new semantic manipulations unattainable by existing methods.

4.4. Editing real images

Importantly, the discovered image manipulations can be naturally applied to real images using the GAN inversion techniques, which embed a given image to the latent space of a pretrained GAN. Furthermore, the manipulations can be applied to the images embedded into the extended W+ space of StyleGAN2. We argue that this is an appealing feature of our transformations since the conventional latent shifts require embeddings to the lower-dimensional Wspace, which can harm the inversion quality. To illustrate how the discovered visual effects perform on real images domain, we invert the real images with the standard Style-GAN2 projector in the W+ space [16]. On the Figure 9 we manipulate real images via the "Nose length" direction. Figure 10 illustrates different manipulations applied to real car images.





activations space



Figure 8. Unsatisfactory reproduction of the "Wheel size" manipulation by the shifts in the latent spaces and in the space of the intermediate activations.



Figure 9. Interpolations along the "Nose length" direction for real images embedded in W+ space.

4.5. Maintaining realism

To verify that modifying the generator parameters does not significantly harm the visual quality, we compute the Frèchet Inception Distance (FID) [12] values for the modified generators. Namely, we plot the FID for the generators with parameters shifted with different magnitudes along the "Nose length" and "Wheel size" directions, see Figure 11.

Notably, even under extreme shift magnitudes, the manipulated samples have high visual quality. We also observe that the FID plot is not symmetric. We attribute this behavior to the fact that the real data can be inherently asymmetric with respect to specific attributes. For instance, the generated cars FID grows much faster once the wheels become small rather than large, probably, because LSUN-Cars con-



Figure 10. Manipulations of real car images embedded in $\mathcal{W}+$ space.



Figure 11. FID for the generators shifted along the "Wheel size" direction (*top*) and the "Nose length" direction (*bottom*).

tains cars with big wheels, but there are no cars without wheels. The same behavior appears with a human nose as tiny noses are rarer than long noses. We also present similar plots for other directions and datasets in the supplementary.

4.6. Locality of visual effects

To illustrate that the navigation along the discovered directions results in an isolated and "disentangled" effect, we perform the following. For a particular direction ξ , we compute the per-pixel differences $||G_{\theta+t\cdot\xi}(z) - G_{\theta}(z)||_2^2$ aver-



Figure 12. Averaged heatmaps of the pixel differences between the original and the edited images for the "Eyes distance" direction (*left*) and the "Nose length" direction (*right*).



Figure 13. *Left*: original image; *center*: a shift in the direction "Wheel rotation"; right: the squared distances between the original and edited image averaged over different shift magnitudes.

aged over 1600 z-samples and 20 shift magnitudes t from a uniform grid in a range [-100, 100]. Figure 12 shows the averaged heatmap for the "Eyes distance" and "Nose length" directions. Notably, the "Nose length" slightly affects the eyes and mouth since the extreme shift pushes the nose to overlap these face areas.

As another example, Figure 13 shows the averaged square distance between an image generated by the original StyleGAN2 and its shifts along the "wheel rotation" direction. Only the "wheel regions" are affected.

4.7. Alternative GANs models

The proposed method is model-agnostic and can be applied even to image-to-image models. In the supplementary we present qualitative results of our approach been applied to pix2pixHD [31] pretrained on Cityscapes dataset and BigGAN [4] pretrained on Imagenet.

5. Conclusion

In this paper, we have investigated the possibilities for GAN-based semantic editing via smooth navigation in the space of the generator's parameters. In particular, we have shown that this space contains various interpretable controls, which can be identified via efficient and straightforward procedures. Given the simplicity and universality of the proposed techniques, these controls become a valuable complement to existing visual editing tools.

References

- David Bau, Steven Liu, Tongzhou Wang, Jun-Yan Zhu, and Antonio Torralba. Rewriting a deep generative model. 2020.
 3
- [2] David Bau, Hendrik Strobelt, William Peebles, Jonas Wulff, Bolei Zhou, Jun-Yan Zhu, and Antonio Torralba. Semantic photo manipulation with a generative image prior. ACM Transactions on Graphics (TOG), 2019. 1, 2
- [3] David Bau, Jun-Yan Zhu, Hendrik Strobelt, Bolei Zhou, Joshua B Tenenbaum, William T Freeman, and Antonio Torralba. Gan dissection: Visualizing and understanding generative adversarial networks. In *International Conference on Learning Representations*, 2019. 1, 3, 4
- [4] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019. 8
- [5] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8789–8797, 2018. 1, 2
- [6] Edo Collins, Raja Bala, Bob Price, and Sabine Susstrunk. Editing in style: Uncovering the local semantics of gans. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 5771–5780, 2020. 1, 3
- [7] Lore Goetschalckx, Alex Andonian, Aude Oliva, and Phillip Isola. Ganalyze: Toward visual definitions of cognitive image properties. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5744–5753, 2019. 2
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, 2014. 1
- [9] Jinjin Gu, Yujun Shen, and Bolei Zhou. Image processing using multi-code gan prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3012–3021, 2020. 1, 2
- [10] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. *NeurIPS*, 2020. 2, 3
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [12] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In Advances in Neural Information Processing Systems, pages 6626–6637, 2017. 7
- [13] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017. 1, 2

- [14] Ali Jahanian, Lucy Chai, and Phillip Isola. On the "steerability" of generative adversarial networks. Proceedings of the International Conference on Learning Representations (ICLR), 2020. 2, 3
- [15] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 5
- [16] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*, 2020. 5, 7
- [17] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photorealistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017. 1, 2
- [18] Hsin-Ying Lee, Hung-Yu Tseng, Qi Mao, Jia-Bin Huang, Yu-Ding Lu, Maneesh Singh, and Ming-Hsuan Yang. Drit++: Diverse image-to-image translation via disentangled representations. *International Journal of Computer Vision*, pages 1–16, 2020. 1, 2
- [19] Richard B Lehoucq, Danny C Sorensen, and Chao Yang. ARPACK users' guide: solution of large-scale eigenvalue problems with implicitly restarted Arnoldi methods. SIAM, 1998. 4
- [20] Chuan Li and Michael Wand. Precomputed real-time texture synthesis with markovian generative adversarial networks. In *European conference on computer vision*, pages 702–716. Springer, 2016. 1, 2
- [21] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, pages 3730–3738, 2015. 2
- [22] Sachit Menon, Alexandru Damian, Shijia Hu, Nikhil Ravi, and Cynthia Rudin. Pulse: Self-supervised photo upsampling via latent space exploration of generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2437–2445, 2020. 1, 2
- [23] Xingang Pan, Xiaohang Zhan, Bo Dai, Dahua Lin, Chen Change Loy, and Ping Luo. Exploiting deep generative prior for versatile image restoration and manipulation. 2020. 1, 2
- [24] William Peebles, John Peebles, Jun-Yan Zhu, Alexei Efros, and Antonio Torralba. The hessian penalty: A weak prior for unsupervised disentanglement. ECCV, 2020. 2
- [25] Antoine Plumerault, Hervé Le Borgne, and Céline Hudelot. Controlling generative models with continuos factors of variations. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020. 2
- [26] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015. 2
- [27] Esther Robb, Wen-Sheng Chu, Abhishek Kumar, and Jia-Bin Huang. Few-shot adaptation of generative adversarial networks. arXiv preprint arXiv:2010.11943, 2020. 4, 5

- [28] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9243–9252, 2020. 2, 3
- [29] Andrey Voynov and Artem Babenko. Unsupervised discovery of interpretable directions in the gan latent space. *ICML*, 2020. 2, 3
- [30] Binxu Wang and Carlos R Ponce. A geometric analysis of deep generative image models and its applications. In *International Conference on Learning Representations*, 2021. 2, 3, 4
- [31] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018. 8
- [32] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 5
- [33] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 2, 4
- [34] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 4
- [35] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A Efros. Generative visual manipulation on the natural image manifold. In *European conference on computer* vision, pages 597–613. Springer, 2016. 1, 2
- [36] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycleconsistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223– 2232, 2017. 1, 2
- [37] Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A Efros, Oliver Wang, and Eli Shechtman. Toward multimodal image-to-image translation. In Advances in neural information processing systems, pages 465–476, 2017. 1, 2