

Can We Find Neurons that Cause Unrealistic Images in Deep Generative Networks?

Hwanil Choi¹, Wonjoon Chang¹ and Jaesik Choi^{1,2,*}

¹Graduate School of AI, KAIST, Seongnam, Republic of Korea

²INEEJI, Republic of Korea

{hwanil.choi, one_jj, jaesik.choi}@kaist.ac.kr

Abstract

Even though Generative Adversarial Networks (GANs) have shown a remarkable ability to generate high-quality images, GANs do not always guarantee the generation of photorealistic images. Occasionally, they generate images that have defective or unnatural objects, which are referred to as ‘artifacts’. Research to investigate why these artifacts emerge and how they can be detected and removed has yet to be sufficiently carried out. To analyze this, we first hypothesize that rarely activated neurons and frequently activated neurons have different purposes and responsibilities for the progress of generating images. In this study, by analyzing the statistics and the roles for those neurons, we empirically show that rarely activated neurons are related to the failure results of making diverse objects and inducing artifacts. In addition, we suggest a correction method, called ‘Sequential Ablation’, to repair the defective part of the generated images without high computational cost and manual efforts.

1 Introduction

Generative Adversarial Networks (GANs) have been exhibited excellent performance for creating photo-realistic images, such as human faces, cars, and buildings. [Karras *et al.*, 2019; Brock *et al.*, 2019; Karras *et al.*, 2020; Chen and Yang, 2021; Esser *et al.*, 2021; Karras *et al.*, 2021]. Additionally, GANs have been used to synthesize a new realistic image from several real images. For example, a target human image can be transformed to obtain another’s face [Kim *et al.*, 2021]. These various uses of GANs have been utilized in many applications [Hung *et al.*, 2021; Chen and Yang, 2021; Liu *et al.*, 2021; Wang *et al.*, 2021].

Even though it is known that GANs have sufficient representative capacity to generate high-quality images, recent studies show that GANs still produce unrealistic images [Odena *et al.*, 2016; Bau *et al.*, 2019] or only generate limited types of samples. In this context, even PGGAN [Karras *et al.*, 2018] and StyleGAN [Karras *et al.*, 2020], which

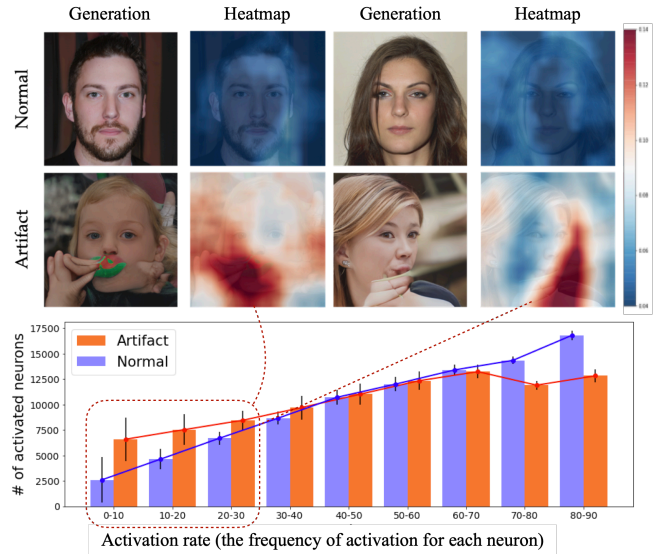


Figure 1: Heatmaps of rarely activated neurons for generated images. The first and second rows show the discrepancy between normal and artifact images respectively. The second and fourth columns describe the ratio of rarely activated neurons in featuremaps. In the red-colored area, where rarely activated neurons are activated a lot, there exist defective parts in the generated images. The bar plot shows that artifact images are highly related to rarely activated neurons compared to well-generated images.

are recognized as the high-performance models, may generate unrealistic images that include defective parts or unnatural objects, referred to as ‘artifacts’. As long as artifacts in generated images are not detected and fixed, the application of GANs will be limited.

Recently, there have been several attempts to identify the artifacts. GAN-Dissection [Bau *et al.*, 2019] analyzes internal units in a pre-trained GAN to identify which units contribute to generating a specific object in a generated image. In this work, the authors found several units that are responsible for typical artifacts. One method to identify the cause of an artifact is to find manually corresponding units creating artifacts. Another method for automatic identification of artifacts units is using unit-specific FID scores. They then turn off each identified unit, which means that activation values

*Corresponding Author

of all neurons in the selected unit are set to zero. Automatic Correction [Tousi *et al.*, 2021] exploits an auxiliary classifier that is trained to distinguish artifacts from normal images and identifies the locations of defective parts by applying Grad-CAM [Selvaraju *et al.*, 2019]. Handcrafted labels are required for fine-tuning this classifier. However, these studies lack analysis on the root cause of artifacts from the perspective of the generating process and they require manual sorting or additional training costs.

In this work, we focus on revealing which neurons make images distorted, how artifacts can be detected, and under which condition can make high quality natural images. In particular, we analyze PGGAN and StyleGAN2, which are known as high-performing models and have been widely adopted as the basic structures of generative models for high-quality high-resolution [Chen *et al.*, 2022; Kim *et al.*, 2021; Yao *et al.*, 2021]. We first analyze the frequency of activated neurons to reveal the following behavior of the generating process. Figure 1 illustrates the relation between defective parts of generated images and rarely activated neurons, which we will define as ‘low frequency rate’. We then empirically verify that defective parts in the artifact images are corresponding to neurons with low frequency rates. From this analysis, we show that artifacts can be censored without any additional training cost. Furthermore, we propose an ablation method to remove defective parts in an unsupervised manner. This method doesn’t require an additional training procedure. This experimental results support our hypothesis both qualitatively and quantitatively. The main contributions of our work are as follows:

- We analyze the frequency of activation in feature maps in order to reveal the properties of the generating procedure.
- We experimentally show that not only are the artifacts related to the rarely activated neurons, but also are the artifacts detected by this property.
- We provide qualitative and quantitative results. Furthermore, our analysis can be applied to detect and remove artifacts from generated images.

2 Related work

2.1 Generative Adversarial Networks

GANs [Goodfellow *et al.*, 2014] have been leading a new paradigm of generative models. A GAN consists of a discriminator and a generator. The discriminator classifies between two images distributions: real images and fake images. The generator creates fake images to deceive the discriminator. One improved version of conventional GAN structure is Progressive Growing GAN [Karras *et al.*, 2018], which is trained sequentially by increasing the resolution, and has emerged as a high-quality models. While recent models based on the GAN structure produce indistinguishable photo-realistic images, there have been few studies on the inner process of the GAN model.

2.2 Style-based GAN

StyleGAN [Karras *et al.*, 2019] was proposed as a novel generating mechanism for GAN to generate high quality and res-

olution images. This model has two parts: a mapping network and a style generator. The former transforms an input latent code z to a style code w . The latter generates an image with the trained constant tensor and the style code w . Although StyleGAN affords good performance, blob artifacts can sometimes be found in the images. To address this, they suggested an advanced version, StyleGAN2 [Karras *et al.*, 2020]. Despite their efforts to remove specific artifacts such as blobs, the general artifacts remain unresolved. In addition, they propose a method improving image quality, called ‘truncation’. This forces the style vector located near the average style vector that is calculated on training time. However, the authors state that this truncation method negatively effects the diversity of output images.

2.3 Semantic Hierarchy Emerges in GAN

A recent study [Yang *et al.*, 2021] revealed the hierarchy mechanism of generative models. The authors of this study noted that the generation process is the same as an artist drawing a picture by interpreting it from multiple abstraction levels. In the GAN models, the early layers create the layout of the output image. The middle layers create objects and attributes. Finally, the color scheme is rendered on the late layers. This implies that the artifacts are likely to occur in the early layers.

2.4 Detecting and removing Artifacts in Generative Models

In recent studies, GAN Dissection [Bau *et al.*, 2019] and Automatic Correction of Internal Units in Generative Neural Networks [Tousi *et al.*, 2021] have struggled to understand the inner mechanism of GANs. They suggest how the units where artifacts occur can be detected and fixing internal errors in the GAN so that the generator can provide realistic images without artifacts. GAN Dissection removes the internal units that are expected to cause defective elements. The latter trains an auxiliary neural network classifier. For the case of an artifact image, the auxiliary classifier is used to search a region that is considered to be strange by applying Grad-CAM [Selvaraju *et al.*, 2019].

3 Analysis for the Activation Frequency Rate

In this section, we attempt to shed light on the causes of the diversity of generated images by observing the frequency of activations in GAN, which may lead to the generation of artifacts. We first introduce notations for GAN structure to define the frequency rate of neurons. According to the frequency rate, we verify that rarely activated neurons are highly related to unusual features and defective parts in generated images by analyzing the effects of those neurons. Furthermore, we suggest an efficient detection and ablation techniques based on our analysis.

3.1 The Frequency Rate of Neurons

Notation. For convenience, we introduce some notations for GAN structure and define the frequency rate to measure the frequency of activations. GAN consists of two parts, a generator G and a discriminator D . The generator

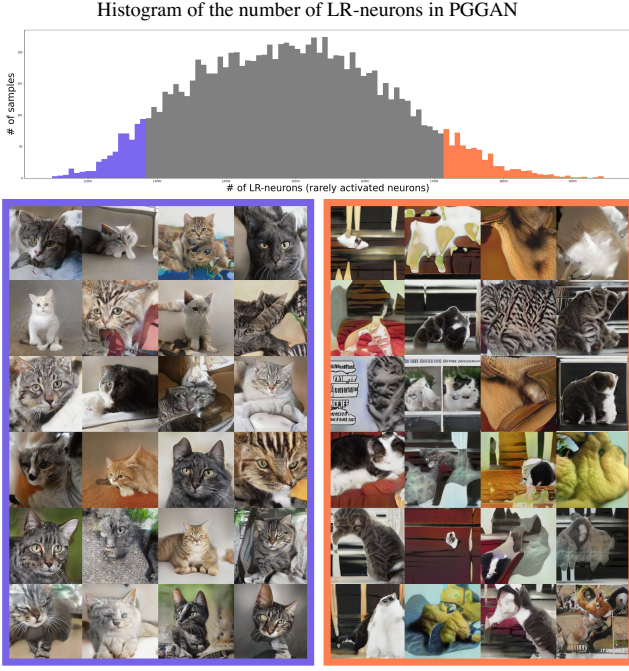


Figure 2: Generated images according to the number of rarely activated neurons in PGGAN. (Left) The samples in the blue box, of which featuremaps are composed mostly of frequently activated neurons, have general cat images. (Right) The samples in the red box have diverse components, but most of them fail to construct realistic cat images.

$G : \mathcal{Z} \rightarrow \mathcal{I}$ generates an output image where \mathcal{Z} is the latent vector space and \mathcal{I} is the space of output images. Since we analyze the process of image generation and internal neurons, we determine the generator G as our main target. A function from the 0 th layer to the i th layer of the generator is denoted by G^i . The output of the i th layer is denoted by $f^i(\mathbf{z}) = G^{(i-1):i}(f^{(i-1)}(\mathbf{z}))$, where $f^0(\mathbf{z}) = G^0(\mathbf{z})$.

Definition 1 (Activated Neuron). Given a generator G and its featuremap $f^i(\mathbf{z})$ for the i -th layer, the activation value of the n -th neuron is defined by $f_n^i(\mathbf{z})$. We call it ‘activated’ if $f_n^i(\mathbf{z}) > 0$. In particular, the set of the activated neurons in the i -th layer is defined by $[f^i]^+$.

Definition 2 (Neurons with Activation Rate R). Given a featuremap f^i , R_n^i is the probability that a neuron $f_n^i(\cdot)$ is activated. $f_{\cdot,R}^i$ is the set of neurons with an activation rate R such that $R_n^i = R$.

The activation rate R_n^i denotes how frequently the neuron n is indeed activated:

$$R_n^i = \mathbb{E}_{\mathbf{z} \sim \mathcal{Z}}[I(f_n^i(\mathbf{z}) > 0)] \quad (1)$$

where I is the indicator function. Since the true probability for each neuron cannot be computed exactly, 30K latent codes are sampled to estimate the relative frequency R_n^i . We analyze the effect of neurons with low frequency rates and that of neurons with high frequency rates. Then, with the constant rate R , we denote the set of neurons which have $R_n^i > R$ by $\bar{f}_{\cdot,R}^i$ and denote the set of neurons that have $R_n^i \leq R$ by $\underline{f}_{\cdot,R}^i$.

When there is no confusion in the context, we simply call $\bar{f}_{\cdot,R}^i$ ‘HR-neurons’ (high-rate neurons) and call $\underline{f}_{\cdot,R}^i$ ‘LR-neurons’ (low-rate neurons).

Hypothesis. In featuremaps of generated images, some neurons are rarely activated whereas other neurons are frequently activated. In other words, neurons have different activation frequency, which are precisely defined as the ‘frequency rate’ in Definition 2. We observe that featuremaps of normal images among generated results are composed mostly of frequently activated neurons. For instance, the rare cases shown in Figure 3 have many neurons with low frequency rates in their featuremaps unlike general cases. Considering this, we hypothesize that rarely activated neurons are used to create diverse components which may include failures, but frequently activated neurons are used to construct general components that are necessary for realistic images.

Figure 2 provides empirical evidences of our hypothesis. The histogram shows the number of generated images according to the number of activated neurons that are rarely activated during the generating procedure. The images in the red box, of which featuremaps have many rarely activated neurons, have diverse components, such as images with phrases. Most of them, however, fail to construct realistic images, while the images in the blue box show general cat images. It is reasonable to suspect that rarely activated neurons are highly related to artifact images generated from GANs.

3.2 Relationship between Frequency Rate and Generations

To observe the effects of neurons according to their frequency rates, we first introduce an ablation method, namely ‘Sequential Ablation’, in Algorithm 1. Given a latent code $\mathbf{z} \in \mathcal{Z}$, it iteratively computes its featuremap while turning off activated LR-neurons for each layer. In other words, it makes values of activated LR-neurons as zeros. The target layers to remove LR-neurons are limited to indices in \mathbf{M} , since neurons on different layers may have different effects on generated images. Previous research shows that each layer in

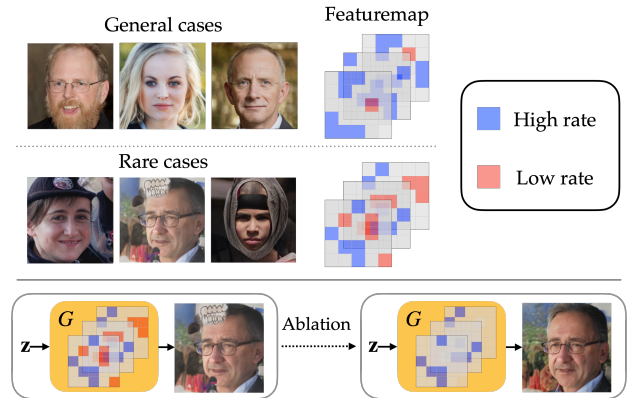


Figure 3: (Top) Comparison of feature maps between general samples and rare samples and (Bottom) our proposed ablation process. The red color denotes rarely activated neurons and the blue color denotes frequently activated neurons.

Algorithm 1 Sequential Ablation

Input: Generator G , A latent code $\mathbf{z} \in \mathcal{Z}$
Parameter: All Layer index list \mathbf{L} , Layer index list \mathbf{M} for ablation, Activation rate R
Output: Corrected Image I_{corr}

```

1: Let  $I \leftarrow \mathbf{z}$ .
2: for  $i$  in  $\mathbf{L}$  do
3:    $f^i \leftarrow G^{(i-1):i}(I)$ 
4:   if  $i$  in  $\mathbf{M}$  then
5:      $[f^i]_{\cdot, R}^+ \leftarrow 0$ 
6:   end if
7:    $I \leftarrow f^i$ 
8: end for
9:  $I_{corr} \leftarrow I$ 
10: return  $I_{corr}$ 
    
```

GANs plays a different role in the generating process in a hierarchical manner [Yang *et al.*, 2021]. The authors note that the early and middle layers determine the spatial layout and the objects, and the last layers render the attributes and color scheme. Therefore, we set \mathbf{M} to include the early or middle layers of the generator, since our work focuses on analyzing the root of unrealistic objects or defective parts.

Analysis for Layer

For a randomly selected artifact image generated from StyleGAN2, the first row in Figure 4 shows how LR-neurons are distributed in its featuremap for each layer. The generated image has a strange object on its chin and a grid-like background. The red part of the heatmaps indicates a high proportion of LR-neurons. The heatmaps on layers 5 and 7 have a large number of LR-neurons at the chin. The second row shows heatmaps and the third row shows generated images, when applying sequential ablation until the previous layers. In other words, the heatmap of layer 5 in the second row denotes the heatmap when applying sequential ablation at layers 0, 1, and 3. Through sequential ablation, LR-neurons at the defective parts of the chin are removed and the generated image is also successfully repaired. Therefore, we can identify that LR-neurons in the early layers are highly related to causing defective parts. However, when we apply sequential ablation for the higher layers, most of the background part disappears and the hairstyle become different. This implies that LR-neurons in layers 5 and 7 contributes to constructing detail features or the background. From this result, we conclude that applying sequential ablation until layer 3 is an appropriate choice to repair generated images without hindering the diversity of generations. In Section 4.2, the artifact correction results in Table 1 support our choice quantitatively.

Analysis for Rate

In this section, we provide qualitative analysis to verify our hypothesis and reveal the roles of neurons according to the frequency rate. Figure 5 shows the generation results of StyleGAN2 with sequential ablation applied in layers 0, 1, and 3 with different rates on FFHQ, LSUN-Cat, and LSUN-Church. In this analysis, we remove not only LR-neurons but also HR-neurons to identify the effects of HR-neurons




Figure 4: Heatmaps of neurons with low frequency rate and their effects on generations. The left most column is generated artifact image. The first row shows the masks of activations corresponding to $[f^i]_{\cdot, 0.3}^+(\mathbf{z})$. The second row shows the masks of activations that are applied sequential ablation upto a given layer. The third row describes the generated images after sequential ablation.

even though sequential ablation originally denotes ablating LR-neurons. For each generated image, the first row shows results from ablating HR-neurons with a given R and the second row shows results from ablating LR-neurons.

When ablating HR-neurons, the basic structures of the objects, such as the shapes of faces, totally collapse. Uncommon characteristics, such as the blue artifact in the first cat image and the human-like artifact in the third cat image, still remain after ablating HR-neurons. Even though the remaining neurons, which are rarely activated, may be necessary parts to generate images with diverse properties, they easily hinder visual fidelity.

Artifacts in the generated images are removed when LR-neurons are ablated. Distorted hands in the generated human images successfully disappear with $R = 0.3$ or 0.5 . The watermarks in the second cat image and the second church image are eliminated, and their shapes become more natural. However, ablating $[f^i]_{\cdot, 0.5}^+$ may remove too much information so that generated images have only fundamental components and lose their properties. In this case, $R = 0.3$ is a valid choice for sequential ablation to remove defective parts while maintaining the details.

3.3 Ablating Neurons to Repair Artifacts

From our previous analysis, we observe that LR-neurons in the early layers are highly related to generate defective parts in the results. Considering this, we propose an efficient ablation method to avoid generating artifact images. It follows the ‘Sequential Ablation’ method, which is described in Algorithm 1, and the early layers are set as a list of target layers \mathbf{M} with $R = 0.3$. Experimental results in Section 4.2 will show that this setting is valid.

Figure 3 describes the overall correction process. Otherwise, we denote applying ablation is only denoted for a single layer by ‘Single Ablation’. In Table 1, we use single ablation to compare the performance. Note that our ablation method does not require complex computations or manual analysis in the repair process. We can easily repair the image by setting activated LR-neurons as zero.




Figure 5: Generated images from StyleGAN2 with ablating HR-neurons or LR-neurons. For each sample, the first row shows the results of ablating HR-neurons, and the second row shows the results of ablating LR-neurons.

| Layers | None | 0 | 1 | 3 | 0,1 | 0,1,3 | 0,1,3,5 |
|-------------------|------|-------------|-------------|-------------|-------------|-------------|-------------|
| $[f_{i,0.7}^i]^+$ | | 240.8 | 241.2 | 259.0 | 233.3 | 277.7 | 342.5 |
| $[f_{i,0.5}^i]^+$ | | 36.3 | 31.6 | 33.4 | 30.9 | 38.5 | 56.7 |
| $[f_{i,0.3}^i]^+$ | 48.9 | 44.9 | 37.0 | 31.9 | 35.6 | 29.6 | 39.2 |
| $[f_{i,0.1}^i]^+$ | | 48.7 | 46.9 | 41.5 | 46.8 | 40.3 | 38.3 |
| $[f_{i,0.7}^i]^+$ | | 74.1 | 124.1 | 118.6 | 147.2 | 155.1 | 163.0 |
| Random | | 43.5 | 44.5 | 55.4 | 52.4 | 178.0 | 226.7 |

Table 1: FID scores for repaired images with various target layer indices and rate R in StyleGAN2 with the FFHQ dataset.

4 Experiment

In this section, we carry out the following qualitative and quantitative experiments to evaluate our approach. The main results show the relationship between LR-neurons and low visual fidelity (or artifacts). Deep generative models were reproduced in PyTorch from genforce github [Shen *et al.*, 2020] which provides the official pre-trained weights; PGGANs trained on CelebA-HQ, LSUN-Cat, and LSUN-Church and StyleGAN2s trained on FFHQ, LSUN-Cat, and LSUN-Church.

4.1 Qualitative Results

50K latent codes were randomly sampled to generate images without the truncation method. The number of the total activated LR-neurons ($R = 0.3$) in the early layers was counted for each generated image, i.e., the sum of $|[f_{i,0.3}^i]^+|$. The target early layers were set to $i = 1, 3, 5$ for PGGAN and $i = 0, 1, 3$ for StyleGAN2. The dimensionality of each layer’s output corresponds to $\mathbb{R}^{512 \times 4 \times 4}$, $\mathbb{R}^{512 \times 8 \times 8}$, and $\mathbb{R}^{512 \times 16 \times 16}$, respectively.

Low Fidelity Detection in Image Generations

To detect images with low visual fidelity, the generated images were sorted according to the number of activated LR-neurons in decreasing order. the top 2K images and Bottom 2K images are then selected. As an example, Figure 2 shows the histogram of the number of activated LR-neurons. The blue region indicates that they have few activated LR-neurons and the red region indicates that the generated images have many activated LR-neurons. Consistent with our hypothesis, the generated images in the red box have more defective or distorted parts than the images in the blue box. The images in the blue box have a high visual fidelity, but lack diversity, whereas the images in the red box seem to have more uncommon properties. We provide abundant results of detection tasks in Appendix B-G.

Sequential Ablation

Figure 5 shows sequential ablation results of StyleGAN2 in FFHQ, LSUN-Cat, and LSUN-Church. The ablation method is conducted on layer 0, 1, and 3 with various rate R as described in Section 3.2. This supports that the LR-neurons induce the artifacts since removing them helps to generate realistic images. Furthermore, in Section 3.3, sequential ablation is conducted to correct artifacts generated from StyleGAN2 and PGGAN. As shown in Figure 6, we identify that our approach can repair the artifacts in various cases. In the case of StyleGAN2, a long crack on the face of the fifth human image was successfully eliminated. The buildings are not completely constructed and mixed with their backgrounds in the artifact church images generated from PGGAN. After sequential ablation, the buildings are repaired normally.

On the unaligned or noisy datasets such as LSUN, however, the HR-neurons have relatively less responsibility for creating a single representative object in the dataset. It is because the datasets that have sparse features can confuse the generative model to consider the objects to be mainly created



Figure 6: Sequential Ablation results for correcting artifacts in StyleGAN2s and PGGANs

| | Datasets | Bottom 2K | | | Top 2K | | |
|-----------|-------------|----------------------|--------------------|----------------------|----------------------|--------------------|---------------|
| | | Precision (fidelity) | Recall (diversity) | Realism Score | Precision (fidelity) | Recall (diversity) | Realism Score |
| PGGAN | CelebA-HQ | 0.9245 | 0.1056 | 1.0911±0.0712 | 0.4565 | 0.3194 | 0.9942±0.0663 |
| | LSUN-Cat | 0.8450 | 0.0893 | 1.0680±0.0686 | 0.4745 | 0.1089 | 1.0011±0.0571 |
| | LSUN-Church | 0.8985 | 0.1907 | 1.0754±0.0627 | 0.4165 | 0.3628 | 0.9902±0.0697 |
| StyleGAN2 | FFHQ | 0.9085 | 0.2153 | 1.1112±0.0913 | 0.762 | 0.4763 | 1.0370±0.0725 |
| | LSUN-Cat | 0.9195 | 0.1404 | 1.0875±0.0704 | 0.6385 | 0.4025 | 1.0206±0.0579 |
| | LSUN-Church | 0.8315 | 0.2226 | 1.0631±0.0640 | 0.4935 | 0.3646 | 1.0016±0.0729 |

Table 2: Precision, Recall, and Realism Score of Top/Bottom images with a lot/few of LR-neurons.

as part of diversity. Therefore, our approach may work better on the well-aligned datasets. Additional results are provided in Appendix H (including the results of StyleGAN3 [Karras *et al.*, 2021] with FFHQ).

4.2 Quantitative Results

Correction Study

Here, we evaluate our ablation method with various combinations of target layers and rate R in the FFHQ dataset with FID score [Heusel *et al.*, 2017]. Table 1 shows FID scores computed with 10K real images and 1.3K generated artifact images. Examples of hand labeled generated images are in Appendix A. ‘None’ denotes a case when we do not apply our ablation method and ‘Random’ denotes when we ablate 30% of randomly selected activated neurons. Note that we remove activated HR-neurons to identify the effects of HR-neurons on generations in the case of the 5th row. The ablation result when turning off $[f_{\cdot,0.3}^i]^+$ with $i = 0, 1, 3$ has the best FID score. Compared to the case $[f_{\cdot,0.3}^i]^+$, the case $[f_{\cdot,0.7}^i]^+$ has very high FID scores because HR-neurons play a significant role in generating essential parts of images. The most optimal setting is used ($[f_{\cdot,R}^i]^+$ with $R = 0.3, i = 0, 1, 3$) to conduct our qualitative experiments.

Diversity and Fidelity

We also quantify our detection method with the precision (fidelity), recall (diversity), and Realism Score [Kynkäänniemi *et al.*, 2019] with 10K real images, the Top 2K and Bottom 2K samples for each dataset, and the number of neighborhood $k = 3$. Table 2 shows that the Bottom 2K samples that contain few LR-neurons have a high score in fidelity and realism score, but a low score in recall. On the other hand, the fidelity and realism score for the Top 2K samples that contain a large number of LR-neurons are lower than those of the Bottom 2K samples, while diversity is higher in all cases. These results are linked to the evidence of our hypothesis that encouraging diversity leads to artifacts. In addition, it may imply a trade-off between fidelity and diversity in the generation procedure.

5 Conclusion

In this paper, we suggest a hypothesis that rarely activated neurons are highly related to generating artifacts, which might be induced as failures to create diversity. Our ablation study not only supports that low-rate neurons in the early layers induce defective parts but also shows that frequently activated neurons may affect essential parts. As a result, our approach can detect artifact images and correct defective parts by sequential ablation without any supervision.

Acknowledgements

This work was supported in part by the Institute of Information and Communications Technology Planning and Evaluation (IITP) grant funded by the Korea Government (MSIT) under Grant 2019-0-00075, Grant 2022-0-00184, and Grant 2022-0-00984 and partly supported by KAIST-NAVER Hypercreative AI Center.

References

- [Bau *et al.*, 2019] David Bau, Jun-Yan Zhu, Hendrik Strobelt, Bolei Zhou, Joshua B. Tenenbaum, William T. Freeman, and Antonio Torralba. Visualizing and understanding generative adversarial networks. In *Proceedings of ICLR*, 2019.
- [Brock *et al.*, 2019] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *Proceedings of ICLR*, 2019.
- [Chen and Yang, 2021] Zehao Chen and Hua Yang. Attentive semantic exploring for manipulated face detection. *Proceedings of ICASSP*, Jun 2021.
- [Chen *et al.*, 2022] Anpei Chen, Ruiyang Liu, Ling Xie, Zhang Chen, Hao Su, and Jingyi Yu. Sofgan: A portrait image generator with dynamic styling. *ACM Trans. Graph.*, 41(1), feb 2022.
- [Esser *et al.*, 2021] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of CVPR*, pages 12873–12883, June 2021.
- [Goodfellow *et al.*, 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in NeurIPS*, 27, 2014.
- [Heusel *et al.*, 2017] 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. *Advances in NeurIPS*, 30, 2017.
- [Hung *et al.*, 2021] Tun-Min Hung, Bo-Yu Chen, Yen-Tung Yeh, and Yi-Hsuan Yang. A benchmarking initiative for audio-domain music generation using the freesound loop dataset. In Jin Ha Lee, Alexander Lerch, Zhiyao Duan, Juhan Nam, Preeti Rao, Peter van Kranenburg, and Ajay Srinivasamurthy, editors, *Proceedings of ISMIR, 2021, Online, November 7-12, 2021*, pages 310–317, 2021.
- [Karras *et al.*, 2018] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *Proceedings of ICLR*, 2018.
- [Karras *et al.*, 2019] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of CVPR*, pages 4401–4410, 2019.
- [Karras *et al.*, 2020] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of CVPR*, pages 8110–8119, 2020.
- [Karras *et al.*, 2021] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in NeurIPS*, 34, 2021.
- [Kim *et al.*, 2021] Hyunsu Kim, Yunje Choi, Junho Kim, Sungjoo Yoo, and Youngjung Uh. Exploiting spatial dimensions of latent in gan for real-time image editing. In *Proceedings of CVPR*, pages 852–861, June 2021.
- [Kynkäänniemi *et al.*, 2019] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in NeurIPS*, volume 32. Curran Associates, Inc., 2019.
- [Liu *et al.*, 2021] Mingcong Liu, Qiang Li, Zekui Qin, Guoxin Zhang, Pengfei Wan, and Wen Zheng. Blendgan: Implicitly gan blending for arbitrary stylized face generation. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in NeurIPS*, volume 34, pages 29710–29722. Curran Associates, Inc., 2021.
- [Odena *et al.*, 2016] Augustus Odena, Vincent Dumoulin, and Chris Olah. Deconvolution and checkerboard artifacts. *Distill*, 2016.
- [Selvaraju *et al.*, 2019] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. *Advances in IJCV*, 128(2):336–359, Oct 2019.
- [Shen *et al.*, 2020] Yujun Shen, Yinghao Xu, Ceyuan Yang, Jiapeng Zhu, and Bolei Zhou. Genforce. <https://github.com/genforce/genforce>, 2020.
- [Tousi *et al.*, 2021] Ali Tousi, Haedong Jeong, Jiyeon Han, Hwanil Choi, and Jaesik Choi. Automatic correction of internal units in generative neural networks. In *Proceedings of CVPR*, pages 7932–7940, 2021.
- [Wang *et al.*, 2021] T. Y. Wang, D. Ceylan, K. Singh, and N. J. Mitra. Dance in the wild: Monocular human animation with neural dynamic appearance synthesis. In *Proceedings of 3DV*, pages 268–277. IEEE Computer Society, dec 2021.
- [Yang *et al.*, 2021] Ceyuan Yang, Yujun Shen, and Bolei Zhou. Semantic hierarchy emerges in deep generative representations for scene synthesis. *Advances in IJCV*, 129(5):1451–1466, 2021.
- [Yao *et al.*, 2021] Xu Yao, Alasdair Newson, Yann Gousseau, and Pierre Hellier. A latent transformer for disentangled face editing in images and videos. In *Proceedings of ICCV*, pages 13789–13798, October 2021.