

Automatic Correction of Internal Units in Generative Neural Networks

Ali Tousi^{1,*}, Haedong Jeong^{1,2,*}, Jiyeon Han¹, Hwanil Choi¹, Jaesik Choi^{1,3,†}

¹Korea Advanced Institute of Science and Technology (KAIST), South Korea

²Ulsan National Institute of Science and Technology (UNIST), South Korea

³INEEJI, South Korea

{ali.tousi, haedong.jeong, j.han, hwanil.choi, jaesik.choi}@kaist.ac.kr

Abstract

Generative Adversarial Networks (GANs) have shown satisfactory performance in synthetic image generation by devising complex network structure and adversarial training scheme. Even though GANs are able to synthesize realistic images, there exists a number of generated images with defective visual patterns which are known as artifacts. While most of the recent work tries to fix artifact generations by perturbing latent code, few investigate internal units of a generator to fix them. In this work, we devise a method that automatically identifies the internal units generating various types of artifact images. We further propose the sequential correction algorithm which adjusts the generation flow by modifying the detected artifact units to improve the quality of generation while preserving the original outline. Our method outperforms the baseline method in terms of FID-score and shows satisfactory results with human evaluation.

1. Introduction

In recent years, GANs have become more powerful in terms of producing photo-realistic images [13, 12] which are often hard to distinguish from real samples. In addition, they have become increasingly better at producing diverse sets of generated samples. These outstanding abilities pave way for GANs to be employed in various real-life domains [9, 18]. The main focus of existing work in the GAN domain has been on improving the quality of synthesis by changing the training scheme or devising more complex models. Despite considerable successes, GANs still suffer from producing outputs that contain unrealistic regions, so-called artifacts, which make them unsuitable for being employed in mission-critical applications. As a result, examining the root cause of such phenomena and pos-

sible solutions to enhance the overall quality has proved to be important.

Recent work removes artifact areas by perturbing the latent codes [22]. Defective units are removed based on understanding the representation of internal units and human annotation [3]. Unlike the previous methods, we identify defective units by learning a classifier from annotated samples. Furthermore, by utilizing an obtained explanation map, we devise our own global multi-layer artifact unit ablation scheme which enhances the quality of defective generations while preserving plausible generations.

In our approach, we first annotate randomly sampled generations into two categories, *Normal* and *Artifact*, based on predefined criteria. Then we train a classifier on all the annotated generations and some randomly sampled real images to classify images into their corresponding categories. Our trained classifier can generate an estimated mask for defective regions by employing an explanation method [21]. By measuring the alignment between an individual internal unit's activation and the defective regions' segmentation mask, we identify units inducing artifacts. In order to correct the artifact areas in the generations, we ablate units with the highest overlap score.

In summary, the contribution of our work is three-fold:

- We compile a large dataset of curated flawed generations and provide a comprehensive analysis on artifact generations.
- We identify defective units in a generative model by measuring the intersection-over-union (IoU) of the unit's activation map and pseudo artifact region masks obtained by training a simple classifier on our dataset.
- We propose an artifact removal method by globally ablating defective units which enhances the quality of artifact samples while maintaining normal samples from drastic change. We further improve the approach by sequentially ablating the defective units throughout consequent layers.

*Equal contribution

†Corresponding Author

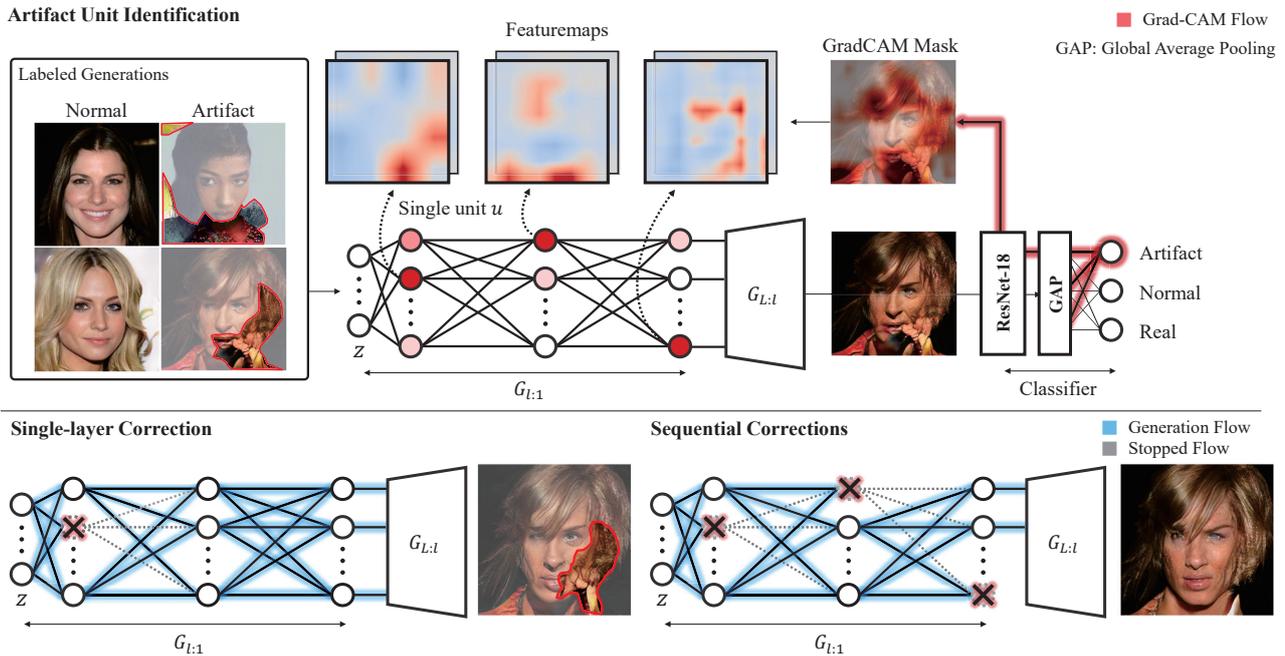


Figure 1: Identification of the artifact units for each layer (top) and the generation flow for two correction methods (bottom). Identification measures the IoU between defective regions (GradCAM mask) and a unit u for artifact generations. The average of IoUs over samples is used as the defective scores in the layer. The sequential correction method adjusts the generation flow of defective units and improves GANs without retraining.

2. Related Work

Generative Adversarial Networks. Since the introduction of GAN [6], the realism and diversity of outputs of the generator have increased steadily [1, 12, 13, 7]. Generally, GAN models take a sampled latent vector and output a synthesized image. While the samples from the original GAN model could be easily identified, recent models produce indistinguishable samples from real data. Despite recent advances, little research has been conducted in order to understand the inner mechanism of a GAN model.

Characterizing Deep Networks' Units. Various techniques have been proposed to examine and understand the internal representation of deep networks [23, 11]. Explanatory heatmaps can be used to explain individual network decisions [16, 21]. The heatmaps visualize which input regions contribute the most to the categorical prediction given by the networks. Recently, [2] introduces the *Network Dissection* framework for identifying the role of internal units of CNN models by gathering a dataset with semantic concepts and label each hidden unit based on the alignment of its activation map and the concept annotation.

Artifacts in Generative Models. Defective regions can be observed in synthesized images by deep neural networks [19, 24, 3, 22]. [19] finds the cause of checkerboard artifacts to be in the deconvolution operation and suggests

a simple resize-convolution upsampling to resolve this issue. [17] exploits the distinctive inherent artifacts that appear in the warping process of generating deep fake videos to be used for detection. [22] tries to first divide the latent space into normal and artifact regions by training a linear SVM model on manually labeled generations. By moving the latent codes toward the good quality direction obtained from the learned hyperplane, they gradually correct synthesized artifacts images. [3] finds artifact-inducing units with/without human supervision. They visualized the highest activation images for each unit and label them as normal or artifact units. Then, they ablate defective units in order to fix the generation.

3. Ablation of Artifact Units in GAN

The term artifact has been used in the previous work [3, 19] to describe the synthesized images which have unnatural (or undesired) visual patterns. Figure 2 shows illustrative examples of artifacts, where the church is distorted or some parts of the image are transparent. When we label the generations into *artifact* or *normal* generations, we observe one could get non-negligible amount of artifacts as listed in Table 1.

One may suggest using the output of the discriminator, which we call 'D-value' for the rest of this paper, as it is



Figure 2: Examples of artifact generations of PGGAN with LSUN-church (top), CelebA-HQ (middle), and LSUN-bedroom (bottom) datasets.

	CelebA-HQ	LSUN-church	LSUN-bedroom
Artifact ratio	38.71 %	16.79 %	40.91 %

Table 1: The ratio of artifact generations in PGGAN with CelebA-HQ, LSUN-church, and LSUN-bedroom datasets.

a natural metric to distinguish real images and generations. It is common to think the normal generations which look real would have D-values following the distribution of real training images, while artifact generations will be thought to have a distant distribution. However, it was not the case from our observations. Figure 3 shows the histogram of D-values for the normal and artifact generations.

It is non-trivial to classify the type of input (normal or artifacts) based on the D-value, because two histograms are overlapped. Also, the level of defectiveness is not proportional to the D-value. These observations show that we need to dig deeper into the network to detect and correct artifacts.

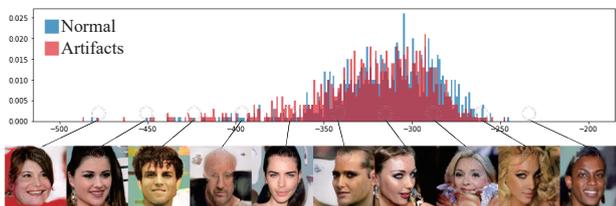


Figure 3: The D-value for each type of input in PGGAN with CelebA-HQ. The black line indicates the corresponding D-value for each generation. The second and the last generations have similar background problems, however, the corresponding D-values are different.

3.1. FID-Based Artifact Unit Identification

In the paper [3], artifacts are also studied as one type of object. To get the artifact units in an unsupervised manner, the authors use Fréchet Inception Distance (FID) which is a commonly used metric for measuring performance for generative models [8]. For each featuremap unit, they compute the FID score on 10K images which have the highest activation of the given featuremap unit among 200K generated images. Featuremap units which are highly activated for high-FID image sets are considered artifact units. The authors show that the FID score improves when the top 20 artifact units are ablated. While this automatic correction method has shown overall improvement, our observation shows there is still a possibility to improve the identification of artifact units. Figure 4 shows the unit with the 5th highest FID score among 512 units.



Figure 4: The representative generation and highly activated generations for unit 134 in layer 6 of PGGAN-LSUN church.

Although unit 134 has a high FID score, the related generation concepts seem to be natural. This example illustrates the reason why we need a more elaborate approach for the identification of artifact units.

3.2. Classifier-Based Artifact Unit Identification

In order to identify the internal units which cause high-level semantic artifacts, we hand-label 2k generated images to normal or artifact generations. Then, we build a model to classify our dataset into three categories, namely: artifact, normal, and randomly selected real samples. We employ an image classifier (e.g. ResNet-18) as our feature extraction module and introduce one fully connected layer on top of it for classification. During the training, we keep the parameters of the feature extraction module fixed and only optimize the classifier weights. In order to obtain a finer-level annotation compared to the image label, we apply GradCAM [21] to provide us with a mask for the defective regions. Such a mask highlights the regions that are effective for the model's decision. This operation is more efficient than hand labeling defective regions as the latter takes a lot of time.

Now that we have a pseudo segmentation ground truth mask for artifact areas, comparing feature activation maps with the mask can reveal the artifact-inducing units. For this purpose, we follow the principles in the prior work [2, 5]. For every unit u in the l -th layer of the generator, we compute $A_u(z_x) \in \mathbb{R}^{H_l \times W_l}$ which is the activation for



Figure 5: The explanation masks from the GradCAM with artifact class score. The masks focus on the defective regions which the normal generations are hard to contain.

a given generation x and the corresponding latent code z_x . We threshold the activation with the quantile T_u such that $P(A_u(z_x) > T_u) = \tau^1$. This is computed with regard to the feature map distribution of each unit for all images. After applying the threshold, we bilinearly upsample the result to the size of the pseudo artifact segmentation mask $L_a(x)$ and compute IoU with $A_u(z_x)$. To make the identification process global, we define the defective score which averages each IoU over artifact generations.

Definition 1 (Defective Score (DS)) Let the set of artifact generations X_a and artifact segmentation mask $L_a(x)$ and activation $A_u(z_x)$ of unit u in the l -th layer for $x \in X_a$. The defective score of unit u is defined as,

$$DS_{l,u,a} = \frac{1}{|X_a|} \sum_{x \in X_a} \frac{|A_u(z_x) \cap L_a(x)|}{|A_u(z_x) \cup L_a(x)|}$$

The defective scores for all units in the l -th layer are defined as $DS_{l,a} = \{DS_{l,1,a}, DS_{l,2,a}, \dots, DS_{l,D_l,a}\}$ where D_l is the number of the units in the l -th layer. Finally, we can sort the scores and choose units with higher scores as candidates for ablation.

Figure 6 denotes the results of ablation with the top 20 units in layer 6 comparing with the FID-based ablation. In the first row, we can identify that both methods barely harm the plausible regions in the generation. It implies that the selected units are related to the defective area and less correlated with the normal generations. In the second row with the shadow artifact, DS-based correction shows more reasonable performance than FID-based correction. However, both methods fail to correct in some cases as in the last row: (1) the hole on the church and (2) the texture error case.

3.3. Generation Concepts of Unit in GANs

From the previous observations, we can identify that a simple unit ablation cannot correct all types of artifacts. To

¹We empirically set $\tau = 0.005$ in the experiments.

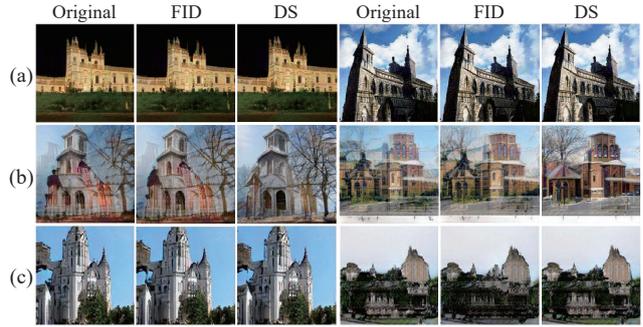


Figure 6: The single-layer ablation results (top 20) for the baseline method (FID) [3] and our method (DS). (a) Correction results on normal generations. The ablations from both methods do not affect plausible regions. (b) Our method can remove the shadow effect while maintaining the original outline. (c) Both methods fail to correct the hole on the wall (left) or the texture error (right).

analyze and explain this phenomenon, we investigate and reveal the generation concept of each featuremap unit. We generate 20k images and select the top 20 images which maximize the magnitude of activation for each unit u . Some featuremap units seem to have concrete concepts because the selected images share concrete semantic information (see more examples in Appendix 5). However, for some featuremap units, it is hard to define one clear generation concept. To better identify such concepts, we compute the mean featuremap amongst highly activated images to generate the representative image for each featuremap unit. Figure 7 denotes the examples of units with concrete/mixed generation concepts.

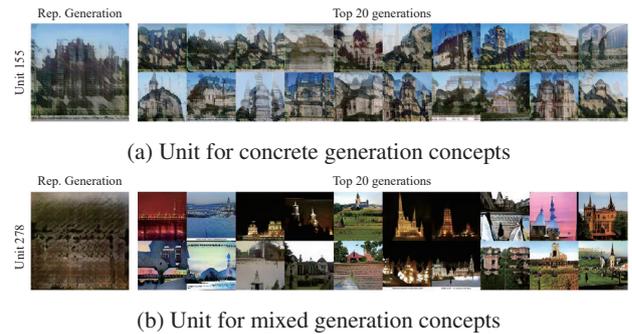


Figure 7: The examples of the representative image and highly activated generations for the given unit in layer 6 of PGGAN on LSUN-church. (a) A unit related to concrete generation concepts. (b) A unit that does not have concrete generation concepts. The representative image in this case seems to be blurred.

Figure 8 denotes the examples which are clustered by the

type of artifacts with the corresponding unit indices. We can clearly identify that each artifact concept is related to multiple units in a hidden layer. It suggests that we need to ablate a set of units which includes most units sharing the defective concepts to ensure the correction performance. Figure 9 denotes the correction results comparing with manual unit selection. Because the top 20 units did not include all the texture error units in FID and DS-based correction, they cannot remove the texture error in the generation. However, when we manually ablate 15 units related to the texture error, the correction can be performed efficiently.

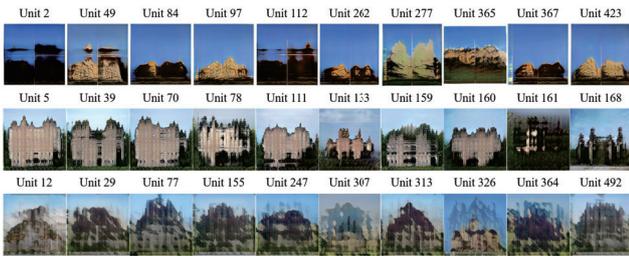


Figure 8: The clusters of the representative generations in layer 6 of PGGAN with LSUN-church. There are multiple units which share the same artifact concept.



Figure 9: An example of correction results for texture error. In the Manual case, we select and ablate the 15 units related to the texture error.

Although we can cover various types of artifacts by simply increasing the number of units for ablation, there exists a trade-off between removing artifacts and preserving the normal information in generations. Figure 10 illustrates such a trade-off. When we increase the number of units, it gradually removes the defective spot on the top-right corner in the generation, but heavy degradation also appears (e.g. the quality of trees and the building). In addition, we plot the FID scores according to the number of units ablated at layer 6 for the PGGAN model across various datasets in Figure 11. We can identify that the FID score increases exponentially when we increase the number of units for ablation.

4. Sequential Correction of Artifacts

We propose the sequential correction to improve the quality of generation performance by suppressing the side-



Figure 10: The change in generation over the number of ablation units. Although the size of the stain (top-right) is reduced when increasing the number of ablation units, the quality of trees and the church is degraded at the same time.

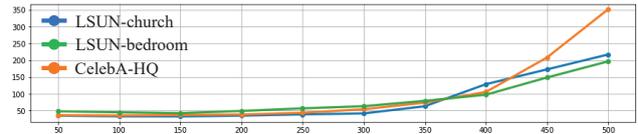


Figure 11: The FID scores on various numbers of ablated units on layer 6 of PGGAN with a various dataset.

effect of one layer ablations. Given generator with L layers, the function of generator G is decomposed into $G(z) = f_L(f_{L-1}(\dots(f_1(z)))) = f_{L:1}(z)$, where z is a vector in the latent space $\mathcal{Z} \subset \mathbb{R}^{D_z}$. $h_{l,u} = f_{l:1,u}(\cdot)$ denotes the values of the u -th unit in the l -th layer with $f_{l:1}(z) \in \mathbb{R}^{D_l \times H_l \times W_l}$. In general, the operation $f_l(\cdot)$ includes linear transformations and a non-linear activation function. For the given query z , we sequentially adjust the activation of units in the consequent layers. The detailed procedure is stated in Algorithm 1.

From the previous research [10] that shallow layers handle the abstract generation concepts and deeper layers handle localized information in GANs, we ablate the shallow layers from the first layer to the stopping layer $l < L$. To prevent the loss of semantic characteristics of a generation as pointed in Section 3.3, we adjust the magnitude of the original featuremaps instead of the simple zero ablation. Line 5 of Algorithm 1 states this soft ablation as,

$$h_{k+1,j} = \lambda(1 - DS_{k+1,j,a})h_{k+1,j}$$

where $\lambda \in [0, 1]$ is the scaling factor and $DS_{k+1,j,a} \in [0, 1]$ is normalized. $\lambda(1 - DS_{k+1,j,a})$ controls the relative generation flow of selected featuremap units. Note that if the scaling factor $\lambda = 0$, the algorithm performs simple zero ablations in consequent layers.

4.1. Analysis of Sequential Correction

To demonstrate the relation between the hyperparameters and correction performance, we first measure the FID after correction varying the stopping layers l and the portion of ablated units n in PGGAN with CelebA-HQ. In Figure 12a, ablating 20% of units in each layer shows the best performance in the FID comparison. The FID sharply increases

Algorithm 1 Sequential Correction

Input: z_0 : a query, $G(\cdot) = f_{L:1}(\cdot)$: a generator,
 l : a stopping layer, $DS_{l:1,a}$: normalized defective scores
for each layer, λ : a scaling factor, n : the number of ablated
units

Output: X : the corrected generation

```
1:  $h_0 = z_0$ 
2: for  $k \leftarrow 0$  to  $l$  do
3:    $h_{k+1} = f_{k+1:k}(h_k)$ 
4:   for  $j \leftarrow \text{Top } 1$  to  $\text{Top } n$  do
5:      $h_{k+1,j} = \lambda(1 - DS_{k+1,j,a})h_{k+1,j}$ 
6:   end for
7: end for
8:  $X = f_{L:l+1}(h_{l+1})$ 
9: return  $X$ 
```

in layers 9 and 12 whereas the differences are minimal for the lower layers. We also provide another metric called *Realism Score* (R) [15] which can measure the quality of generation for individual samples for various stopping layers with the fixed $n = 20\%$. Figure 12b indicates the average of R over the 1k artifact generations. We find that the stopping layer 6 shows the best R for the sequential correction.

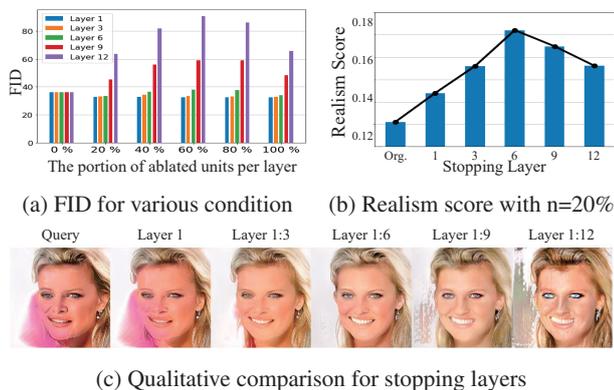


Figure 12: Quantitative and qualitative results over various hyper parameters (stopping layer l and the portion of ablated units n). We can identify that when the stopping layer l is set deeper layer, the quality of generation is degraded.

4.2. Local Region Correction

Although the proposed method defines the artifact units in a global sense, we can apply local region correction additionally, since we have the mask of defective regions for individual samples. We change the reducing weight $\lambda(1 - DS_{k+1,j,a})$ to $(1 - L_a(x))$ where $L_a(x)$ is the down-sampled GradCAM mask. In this scheme, we can perform the sample-specified sequential correction and the correc-

tion result can be local. As shown in Figure 13, the local region correction can minimize the change of unmasked contents per individual samples.



Figure 13: The effect of local region correction for the sequential correction (SeqC) in PGGAN-celebAHQ. We can identify that the area which was not focused on by the mask cannot be changed.

5. Experimental Evaluations

This section presents the analytical results of our algorithm and empirical comparisons of various correction methods. We perform correction on three PGGANs trained on LSUN-church, CelebA-HQ, and LSUN-bedroom datasets, respectively. We manually label generations from PGGANs and collect 1k artifact generations and its latent codes for each model. We obtain the featuremaps through the model and ablate the chosen featuremap units. Throughout the experiments in this paper, we use the stopping layer $l = 6$, the number of ablation units $n = 20\%$, and the scaling factor $\lambda = 0.9$. All the correction experiments are conducted on the same 1k latent codes of the original artifact generations. The qualitative and quantitative results are presented in the following subsections.

5.1. Qualitative Result

We first demonstrate how the generations can be corrected for each correction method. As shown in Figure 14, we can identify that the sequential correction can remove the defective region effectively. Especially, ablations in the shallow layer are helpful to correct for the regions in which the generation information is not clear (e.g. the constant green regions in row-5-LHS CelebA-HQ generation) and most information related to the normal characteristics can be maintained. It means that the proposed method minimizes the change of plausible regions and mainly focuses on the defective regions.

5.2. Quantitative Result

We use the FID score to quantitatively measure the improvements of artifacts for each approach. In FID calculation, 1k of real training samples are used. For better comparison, we report the FID scores on 1k original *artifact* generations and 1k *normal* generations in Table 2. As a



Figure 14: The correction results for each ablation method of the PGGAN with the various dataset (LSUN-Church Outdoor, LSUN-Bedroom, and CelebA-HQ.). We confirm that the sequential correction removes the defection regions effectively.

baseline, we choose random ablation and FID based correction method [3] and summarize it under *random* and *FID* row in Table 2. Then we ablate the global artifact units obtained by classifier-based approach and summarize it under *DS* row. The results on the sequential correction approach are summarized under *Sequential* row. The results on each model are summarized under the LSUN-church, CelebA-HQ, and LSUN-bedroom columns, respectively. The proposed method shows the best performance in all three models.

5.3. Human Evaluation

To support the quality of the corrected results, we provide human evaluation results. The experiments are constructed by two evaluation procedures: (1) re-labeling the corrected generations from the sequential correction, and (2) assessing the improvement (improved/not improved) for each pre-defined artifact type. We set the criteria for each

Correction	LSUN-church	CelebA-HQ	LSUN-bedroom
Random	53.43	42.10	67.46
FID	40.66	44.37	48.48
DS	32.82	35.40	44.93
Sequential	23.96	34.71	40.71
Artifacts	46.95	36.16	61.17
Normals	22.37	29.80	29.15

Table 2: FID scores of corrected artifact generations for LSUN-church, CelebA-HQ, and LSUN-bedroom datasets.

artifact type to make consistent evaluations. A detailed description for each criterion can be found in Appendix 1.2.

The results on 500 artifact corrections for each dataset are summarized in Table 3. For CelebA-HQ dataset, we could find 53% out of 500 correction samples obtained by applying our method for artifact-labeled generations are re-labeled as *normal*. While 47% are still containing artifacts,

Dataset	Corrected (%)	Improved (%)
CelebA-HQ	53.00 (4.20)	96.00 (2.00)
LSUN-church	54.50 (0.90)	86.10 (6.30)
LSUN-bedroom	46.80 (8.60)	95.50 (1.10)

Table 3: Human evaluation results on the corrected artifact generations. The number in the parentheses indicates the standard deviation over raters.

97 % of the total artifact generations have significant improvements in the artifact regions or in the quality of generation. The results on the other two datasets show the similar pattern that almost half of the artifact generations are corrected and the visual quality is mostly improved.

5.4. Generalization

Although the proposed method is performed on the generator with a conventional structure, the same approach with minor modification can be generalized for the recent state-of-the-art generators such as StyleGAN2 [14] or U-net GAN [20] which is a variation of BigGAN [4]. Because of the distinct structure of each generator (e.g. each generation uses a different convolution kernel in the StyleGAN2), it is non-trivial to align the proposed framework to identify the defective units in the global sense. However, we can obtain the relative defective score for each unit by individually comparing it with the GradCAM mask for implementing the sequential correction. Figure 15b indicates the sequential correction results in the StyleGAN2. More correction examples for StyleGAN2 and U-net GAN are in Appendix 3-4.

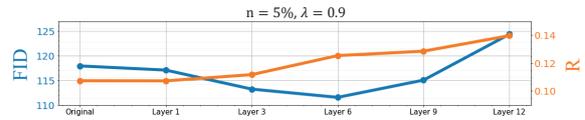
Model	Artifact	Corrected
StyleGAN2	117.91	113.08
U-net GAN	145.15	143.85

Table 4: FID scores of 100 artifact generations and the corrected generations for StyleGAN2 and U-net GAN on FFHQ.

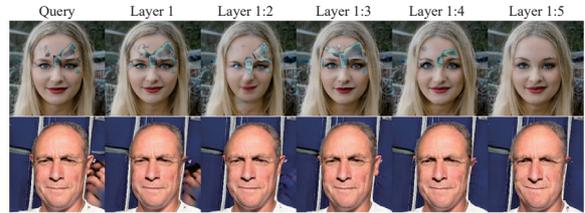
6. Discussion

In this paper, we propose the sequential correction method to improve the quality of generation without additional training of the generator. Especially, we define defective score which quantifies relations between each unit and artifact generations with a supervised approach. The sequential correction uses selected units by the DS-based unit identification and adjusts the generation flow in the consequent layers. We show proposed method achieves plausible correction performance and suggest the opportunity of generalization for the various structures of the generator.

While our method has shown to improve the artifact images in terms of both human evaluation and FID-score, there exist some cases that the model improves the image by sim-



(a) The quantitative results for various stopping layers.



(b) The qualitative results for various stopping layers.

Figure 15: Correction results on StyleGAN v2 with FFHQ. We select 100 artifact generations and calculate the FID and R for the sequential correction with various stopping layers l . We can identify that the scores are improved when the stopping layer l increases up to the middle layer ($l = 6$).



Figure 16: Examples of correction images with glasses or sunglasses. Although the artifact regions are expected background, the sequential correction removes the glasses or sunglasses at the same time.

plifying the generation. As shown in Figure 16, we can observe that glasses/sunglasses are removed instead of completing defective regions. We suspect this is caused by the undesirable features that are trained in the classifier which can be further explored in future work.

In addition, while the plausible regions are hardly changing dramatically, we could observe some cases which fail to maintain the original outline of generation. For example, the LSUN-church generation on row-2-RHS in Figure 14 shows that the original structure of the church is changed, although the unclear pattern on the door is repaired. In row-6-RHS celebA-HQ generation case, the stains are removed and the blond hair appears, while the angle of the face is changed. We consider this side-effect as a limitation of sequential correction since the identification of artifact units includes average over samples. For this purpose, the fusion of global and individual sample-based correction can be considered for future work.

Acknowledgement This work was supported by Institute for ITP grant funded by the Korea government (MSIT) (No.2017-0-01779, XAI and No.2019-0-00075, Artificial Intelligence Graduate School Program (KAIST)).

References

- [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein Generative Adversarial Networks. In *International Conference on Machine Learning*, pages 214–223. PMLR, Jul 2017.
- [2] David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. Network dissection: Quantifying interpretability of deep visual representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [3] David Bau, Jun-Yan Zhu, Hendrik Strobelt, Zhou Bolei, Joshua B. Tenenbaum, William T. Freeman, and Antonio Torralba. Gan dissection: Visualizing and understanding generative adversarial networks. In *International Conference on Learning Representations*, 2019.
- [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.
- [5] Ruth Fong and Andrea Vedaldi. Net2vec: Quantifying and explaining how concepts are encoded by filters in deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- [7] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [8] 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.
- [9] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *CVPR*, 2017.
- [10] Giyoung Jeon, Haedong Jeong, and Jaesik Choi. An efficient explorative sampling considering the generative boundaries of deep generative neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 4288–4295, 2020.
- [11] Andrej Karpathy, Justin Johnson, and Li Fei-Fei. Visualizing and Understanding Recurrent Networks. *arXiv*, Jun 2015.
- [12] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018.
- [13] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [14] 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.
- [15] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. *CoRR*, abs/1904.06991, 2019.
- [16] Sebastian Lapuschkin, Alexander Binder, Grégoire Montavon, Klaus-Robert Müller, and Wojciech Samek. The lrp toolbox for artificial neural networks. *Journal of Machine Learning Research*, 17(114):1–5, 2016.
- [17] Yuezun Li and Siwei Lyu. Exposing deepfake videos by detecting face warping artifacts. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2019.
- [18] Michaël Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- [19] Augustus Odena, Vincent Dumoulin, and Chris Olah. Deconvolution and Checkerboard Artifacts. *Distill*, 1(10):e3, Oct 2016.
- [20] Edgar Schonfeld, Bernt Schiele, and Anna Khoreva. A u-net based discriminator for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8207–8216, 2020.
- [21] 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. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [22] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *CVPR*, 2020.
- [23] Matthew D. Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 818–833, Cham, 2014.
- [24] Xu Zhang, Svebor Karaman, and Shih-Fu Chang. Detecting and simulating artifacts in gan fake images. In *2019 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–6, 2019.