

SSD-GAN: Measuring the Realness in the Spatial and Spectral Domains

Yuanqi Chen^{1,2}, Ge Li^{1*}, Cece Jin^{1,2}, Shan Liu⁴, Thomas Li^{1,3}

¹ School of Electronic and Computer Engineering, Peking University ² Peng Cheng Laboratory

³ Advanced Institute of Information Technology, Peking University ⁴ Tencent America
cyq373@pku.edu.cn, geli@ece.pku.edu.cn, fordacre@pku.edu.cn, shanli@tencent.com, tli@aiit.org.cn

Abstract

This paper observes that there is an issue of high frequencies missing in the discriminator of standard GAN, and we reveal it stems from downsampling layers employed in the network architecture. This issue makes the generator lack the incentive from the discriminator to learn high-frequency content of data, resulting in a significant spectrum discrepancy between generated images and real images. Since the Fourier transform is a bijective mapping, we argue that reducing this spectrum discrepancy would boost the performance of GANs. To this end, we introduce SSD-GAN, an enhancement of GANs to alleviate the spectral information loss in the discriminator. Specifically, we propose to embed a frequency-aware classifier into the discriminator to measure the realness of the input in both the spatial and spectral domains. With the enhanced discriminator, the generator of SSD-GAN is encouraged to learn high-frequency content of real data and generate exact details. The proposed method is general and can be easily integrated into most existing GANs framework without excessive cost. The effectiveness of SSD-GAN is validated on various network architectures, objective functions, and datasets. Code is available at <https://github.com/cyq373/SSD-GAN>.

Introduction

Generative Adversarial Networks (GANs) (Goodfellow et al. 2014) involve training a generator and discriminator network in an adversarial manner, such that the generator learns to reproduce the desired data distribution. Despite the remarkable achievements in image generation tasks (Isola et al. 2017; Karras, Laine, and Aila 2019), as shown in recent works (Zhang, Karaman, and Chang 2019; Durall, Keuper, and Keuper 2020; Dzanic and Witherden 2019; Frank et al. 2020), we can efficiently distinguish GAN-generated images from real images in the frequency domain, which indicates that existing GANs fail to learn the spectral distributions.

Recent studies (Dzanic and Witherden 2019; Frank et al. 2020) show that the frequency spectrum discrepancy mainly exists at high frequencies. Because high-frequency components of images influence the exactness of details, the discrepancy cannot be ignored for generative tasks where details matter. As shown in Fig. 1, when real data contains




Figure 1: Standard GAN (SGAN) fails to learn high frequencies of real data. In this toy example, the real data is a single 16×16 image with a checkerboard pattern, which contains significant high frequencies. Unlike SGAN, our proposed SSD-GAN alleviates the spectral information loss in the discriminator and can well reproduce the real data. The detailed experimental setup is described in the supplemental materials.

significant high frequencies, standard GAN might fail to reproduce the desired data distribution. Moreover, since the Fourier transform is a bijective mapping, the frequency spectrum discrepancy between real data and the generated samples also indicates that the data distribution in image space is not well captured. We believe that reducing the spectrum discrepancy would boost the performance of GANs.

In this paper, we first attempt to explore why there is a spectrum discrepancy between real data and the generated samples. By investigating the downsampling techniques widely used in the discriminator networks, we reveal that both of these downsampling strategies, downsampling with anti-aliasing and downsampling without anti-aliasing, would lead to high frequencies missing in the discriminator. Since the training of GANs is a two-player minimax game, the generator lacks incentives from the discriminator to learn the high-frequency information of the data.

To address the issue of high frequencies missing in the discriminator, we propose SSD-GAN, whose discriminator can measure the realness of the input in both the spatial and spectral domains. To instantiate the idea, we introduce an additional spectral classifier to detect frequency spectrum discrepancy between real and generated images and integrate it into the discriminator of GANs. With the enhanced discriminator, the generator of SSD-GAN is encouraged to

*Corresponding author.

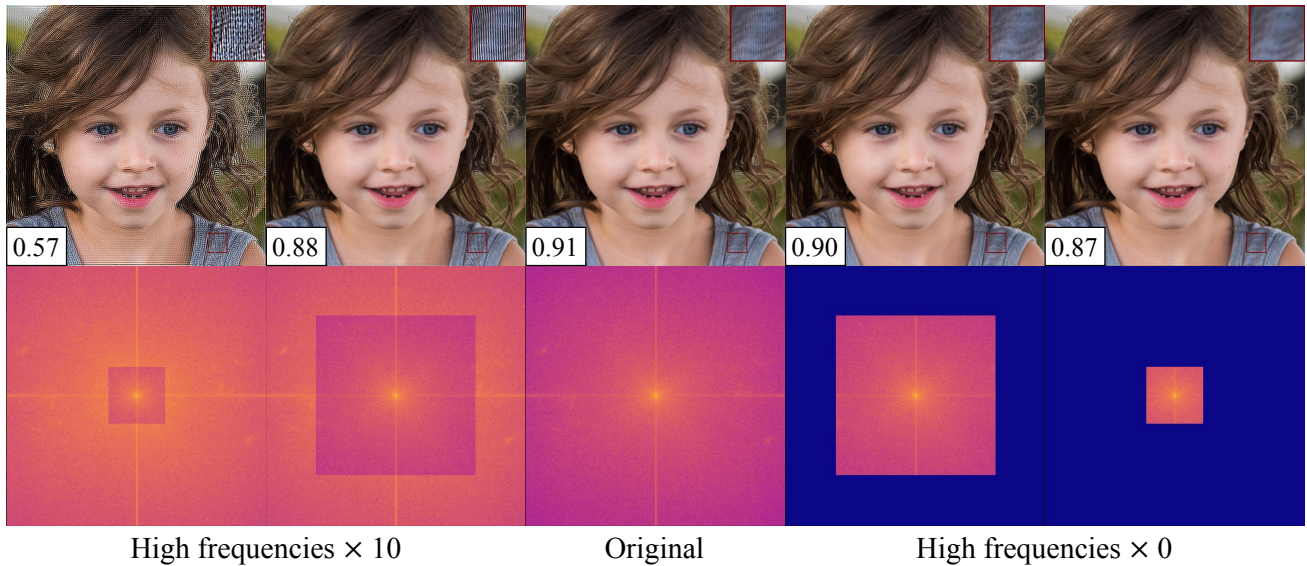


Figure 2: The StyleGAN’s discriminator cannot distinguish the difference of high frequencies. The numbers are estimates of $\mathbb{E}[D(x)]$ by averaging 1,000 samples. Unless we change the spectrum over a large bandwidth, the outputs of the discriminator are roughly the same, which makes StyleGAN fail to reproduce spectral distribution.

reduce frequency spectrum discrepancy and generate realistic images in both the spatial and spectral domains. Since a lightweight spectral classifier can be effective, the proposed method is general and can be easily integrated into most existing GANs framework without excessive cost. In the experiment, the effectiveness of the proposed method is validated on various network architectures, objective functions, and datasets.

The contributions of the paper can be summarized as follows:

- We observe there is an issue of high frequencies missing in the discriminator of GANs and reveal it stems from downsampling layers employed in the network architecture, which results in a significant spectrum discrepancy between generated images and real images.
- We introduce SSD-GAN, an enhancement of GANs to alleviate spectral information loss in the discriminator. With a discriminator that can measure both the spatial and spectral realism of an input sample, SSD-GAN can better capture the data distribution than standard GANs.
- We show experimentally that the quality of the generations can be improved by reducing the frequency spectrum discrepancy, which emphasizes the necessity of learning in the frequency domain.

Related Work

Generative Adversarial Networks Recent rapid advances in Generative Adversarial Networks (GANs) (Goodfellow et al. 2014) have greatly promoted the computer vision and image processing community, e.g., image inpainting (Yang et al. 2018; Ren et al. 2019), image colorization (Isola et al. 2017), image-to-image translation (Yu et al.

2019), etc. To enhance the quality of generated samples, PG-GAN (Karras et al. 2018) introduces a progressive growing manner for the training process to increase the resolution of synthesized images. StyleGAN (Karras, Laine, and Aila 2019) propose a style-based generator for finer control over the image synthesis. Other lines of work focus mainly on improving the discriminator of GANs. As details are important for generative tasks, PatchGAN discriminator (Isola et al. 2017) utilizes local discriminator feedback to capture better local structures. SNGAN (Miyato et al. 2018) limits the spectral norm of the weight matrices in the discriminator for Lipschitz constraint. For countering discriminator forgetting and stabilize the training process, SS-GAN (Chen et al. 2019) proposes to rotate the image and ask the discriminator to predict the rotation angle. To effectively balancing the performance of the generator and discriminator, variational discriminator bottleneck (Peng et al. 2019) constrains information flow in the discriminator. In this paper, after observing the issue of high frequencies missing in the discriminator, we aim to enhance the ability in the frequency domain of the discriminator. Thus the generator is encouraged to learn the spectral distribution of real data.

Frequency Analysis for CNNs Even though some GAN-generated images seem to be flawless for human perception, recent studies (Zhang, Karaman, and Chang 2019; Dzanic and Witherden 2019; Frank et al. 2020; Durall, Keuper, and Keuper 2020) find frequency analysis is effective for image forensics. They also show that existing GAN based models always fail to reproduce the spectral distribution of real data. AutoGAN (Zhang, Karaman, and Chang 2019) first identifies that spectral artifacts stem from upsampling modules included in the GANs pipeline. To compensate spectral distortions, a spectral regularization term (Durall, Keuper, and Ke-

uper 2020) is proposed to add to the generator loss. (Frank et al. 2020) examines StyleGAN instances using different upsampling techniques and finds bilinear sampling followed by anti-aliasing filters would help to alleviate the problem. In this paper, we investigate another source of spectral distortions, the issue of high frequencies missing in the discriminator.

Apart from frequency analysis for image forensics, researchers prove that neural networks exhibit a spectral bias (Rahaman et al. 2019; Xu, Zhang, and Xiao 2019); they learn filters with a strong bias towards lower frequencies. Based on this observation, to effectively control the resource usage, band-limited convolutional layer (Dziedzic et al. 2019) is introduced to constrain the frequency spectra of filters and data, while retaining high performance for classification tasks. However, high-frequency components cannot be ignored for generative tasks where details matter. To guarantee all the information can be kept in the model, MWCNN (Liu et al. 2018) utilizes discrete wavelet transform (DWT) as a downsampling module in the network architecture for image restoration. Different from it, we introduce a spectral classifier to compensate for high-frequency information loss of GANs’s discriminator.

High Frequencies Missing in the Discriminator

In the standard GANs (Goodfellow et al. 2014), the adversarial loss for the discriminator is defined as:

$$\mathcal{L}_D = -\mathbb{E}_{x \sim p_{data}(x)}[\log D(x)] - \mathbb{E}_{x \sim p_g(x)}[\log(1 - D(x))], \quad (1)$$

and $D(x)$ represents the probability that x comes from p_{data} rather than the generator’s distribution p_g . In other words, $D(x)$ measures the realism of the sample x . If x is realistic, then it is realistic in all aspects, such as in the spatial and frequency domains. However, as pointed out in recent works (Zhang, Karaman, and Chang 2019; Durall, Keuper, and Keuper 2020; Dzanic and Witherden 2019; Wang et al. 2020b; Frank et al. 2020), existing GAN based models usually fail to synthesize samples that are realistic in the frequency domain. It suggests that we cannot only measure the realism in the spatial domain.

Why do these GAN based models fail to reproduce the spectral distributions? We suspect that the generator lacks incentives from the discriminator to learn the high-frequency information of the data, since the training of GANs is a two-player minimax game. To validate the assumption, we first randomly sample 1,000 images from the real dataset. Then we modulate the amplitude of high frequencies of different bands. For the discriminator of StyleGAN (Karras, Laine, and Aila 2019), we compute the mean output of it for images after inverse Fourier transform of the modified spectra. As shown in Fig. 2, unless we change the spectrum over a large bandwidth, the discriminator cannot tell the difference of these spectra and the outputs are roughly the same. As a result, if the generated images contain some unusual high-frequency components, the discriminator may not distinguish them to be fake, which makes StyleGAN fail to reproduce spectral distribution.




Figure 3: Downsampling causes high frequencies missing. For these two downsampling strategies, the left column is the output of the raw image and its spectrum, while the right column is the output of the modulated image and its spectrum. After Blur + AvgPool, there is no significant difference between the spectra before and after the amplitude modulation. For AvgPool, the results exhibit aliasing (e.g., see shirt in the images), which indicates high-frequency details become invalid.

Why does the discriminator fail to distinguish the high-frequency contents of the images? We believe that there is an issue of high frequencies missing in the architecture of the discriminator. Specifically, this issue stems from the downsampling modules of the discriminator. When downsampling an input image, based on the classical sampling criterion (Nyquist 1928), a reasonable approach is to anti-alias by low-pass filtering the image. Some networks adopt this form of blurred-downsampling (LeCun et al. 1990; Karras, Laine, and Aila 2019). However, the low-pass filter removes the high frequencies of the input images. Since low-pass filtering before downsampling usually results in performance degradation, it is rarely used today. Another line of downsampling methods, such as max-pooling, strided-convolution, and average-pooling, abandons the use of the low-pass filter. However, these commonly used downsampling methods ignore the sampling theorem (Zhang 2019),

and high-frequency components are aliased and become invalid. To sum up, both of these downsampling strategies, downsampling with anti-aliasing and downsampling without anti-aliasing, lead to high frequencies missing in the discriminator.

As shown in Fig. 3, we provide evidence for the above statement. For an input image, we first enhance the amplitude of high frequencies using a sharpening filter. Then we downsample the raw image and the modulated image and compare the results. For Gaussian blur followed by average-pooling, which belongs to downsampling with anti-aliasing, the high-frequency components are attenuated, and there is no significant difference between the spectra before and after the amplitude modulation. For average-pooling, since it does not provide the anti-aliasing capability, the results exhibit aliasing (e.g., see shirt in the images), which indicates high-frequency details become invalid. The more downsampling modules the deep network has, the wider the bandwidth of the lost high frequencies, which indicates that the high frequencies missing issue cannot be ignored, especially for generative tasks where details matter.

Methodology

In this section, we first introduce a spectral classifier C to detect frequency spectrum discrepancy between real and generated images. Then we integrate C into the discriminator of GANs to enhance its ability in the spectral domain, thereby reducing the spectrum discrepancy.

Detecting Frequency Spectrum Discrepancy

To address the issue of high frequencies missing in the discriminator, a straightforward approach is to discriminate in the frequency domain rather than the spatial domain. For a discrete two-dimensional signal $f(m, n)$ representing an image of size $M \times N$, we first compute the discrete Fourier transform \mathcal{F} of it,

$$\mathcal{F}(k, l) = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} f(m, n) e^{-2\pi i \frac{km}{M}} e^{-2\pi i \frac{ln}{N}}, \quad (2)$$

for the spectral coordinates $k = 0, \dots, M-1$ and $l = 0, \dots, N-1$. Then we convert it from Cartesian coordinates k and l to polar coordinates r and θ for better representing the frequencies of different bands,

$$\mathcal{F}(r, \theta) = \mathcal{F}(k, l) : r = \sqrt{k^2 + l^2}, \theta = \text{atan2}(l, k). \quad (3)$$

Recent works (Durall, Keuper, and Keuper 2020; Dzanic and Witherden 2019) have shown that a simple 1D representation of the Fourier power spectrum is effective to highlight the difference between the spectral characteristics of real and deep network generated images. Following these works, we get the reduced spectral representation v by azimuthally averaging over θ ,

$$v(r) = \frac{1}{2\pi} \int_0^{2\pi} |\mathcal{F}(r, \theta)| d\theta, \quad (4)$$

which represents the mean intensity of the signal with respect to the radial distance r . The reduced spectral representation smooths the fluctuations in the spectrum at high frequencies.

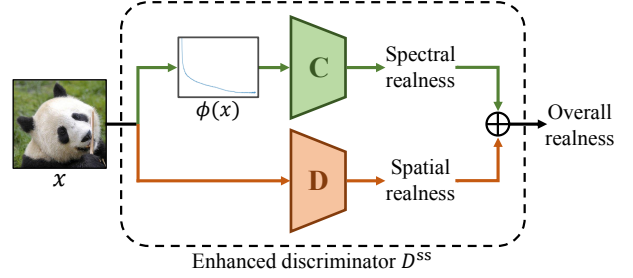


Figure 4: The enhanced discriminator D^{ss} measures both the spectral realism and the spatial realism.

For an input image x , we use the grayscale component of it to get its spectral vector v and denote the process as $v = \phi(x)$. The spectral classification loss is:

$$\begin{aligned} \mathcal{L}_{spectral} = & \mathbb{E}_{x \sim p_{data}(x)} [\log C(\phi(x))] \\ & + \mathbb{E}_{x \sim p_g(x)} [\log(1 - C(\phi(x)))] \end{aligned} \quad (5)$$

where $C(\phi(x))$ measures the spectral realism of x , and p_g is the generator G 's distribution.

Reducing Frequency Spectrum Discrepancy

Since a sample x is realistic if and only if it is realistic in both the spatial and frequency domains, we propose to measure the realism of x with the combination of spatial realism and spectral realism. We integrate the spectral classifier C into the discriminator of GANs to encourage the generator to learn the high-frequency content of the data. As shown in Fig. 4, our enhanced discriminator D^{ss} consists of two modules, a vanilla discriminator D which measures the spatial realism, and a spectral classifier C . Thus, D^{ss} is a discriminator measuring the realism of the input in both the spatial and spectral domains, and the overall realism of a sample x is represented as:

$$D^{ss}(x) = \lambda D(x) + (1 - \lambda) C(\phi(x)), \quad (6)$$

where λ is a hyperparameter that controls the relative importance of the spatial realism and the spectral realism. The adversarial loss of the framework can be written as:

$$\begin{aligned} \mathcal{L}_{adv} = & \mathbb{E}_{x \sim p_{data}(x)} [\log D^{ss}(x)] \\ & + \mathbb{E}_{x \sim p_g(x)} [\log(1 - D^{ss}(x))], \end{aligned} \quad (7)$$

where p_g represents the generator G 's distribution.

To train our model, we alternately update spectral classifier C , discriminator D , and generator G with the following gradients:

$$\begin{aligned} \theta_c & \leftarrow -\nabla_{\theta_c} \mathcal{L}_{spectral}, \\ \theta_d & \leftarrow -\nabla_{\theta_d} \mathcal{L}_{adv}, \\ \theta_g & \leftarrow \nabla_{\theta_g} \mathcal{L}_{adv}. \end{aligned} \quad (8)$$

Analyzing the Effect of the Spectral Classifier

Since much information of the image x is discarded in the spectral vector $\phi(x)$, we found it cannot provide an effective gradient for the adversarial training, which degrades the



Figure 5: Generations of SSD-StyleGAN trained on FFHQ at 1024×1024 .

performance of the model. To this end, we propose that the backpropagation process of Eq. 7 does not pass through the spectral classifier C , and $C(\phi(x))$ serves as a spectral modulating factor to the adversarial loss.

We compare the gradient of standard GAN (SGAN) and the proposed method for further insight. For a generated image $x = G(z)$, the gradients of the discriminator and generator in non-saturating SGAN are respectively:

$$\frac{1}{1 - D(x)} \nabla_{\theta_d} D(x), \quad (9)$$

$$-\frac{1}{D(x)} \nabla_x D(x) J_{\theta_g} G(z), \quad (10)$$

where θ_d and θ_g are the parameters of D and G , and J is the Jacobian. As for our method, the gradients are:

$$\frac{1}{1 - D(x) + \frac{1-\lambda}{\lambda}(1 - C(\phi(x)))} \nabla_{\theta_d} D(x), \quad (11)$$

$$-\frac{1}{D(x) + \frac{1-\lambda}{\lambda}C(\phi(x))} \nabla_x D(x) J_{\theta_g} G(z), \quad (12)$$

From these gradients, it can be observed that our method performs a hard example mining, where "hard" is defined in the frequency domain. For the discriminator, if $C(\phi(x)) \rightarrow 1$, the generated sample x has good spectral characteristics and is a hard example to be classified as fake. For the generator, x is a hard example when $C(\phi(x)) \rightarrow 0$. This means that x has poor spectral realism and needs more attention from the generator. In our model, when x is a hard example in the frequency domain, the gradients of the discriminator and generator are up-weighted, which induces the model to learn the spectral distribution of the real data.

Experiments

Since $\phi(x)$ is easy to compute and C can be a lightweight classifier (Durall, Keuper, and Keuper 2020), the proposed method is general and can be easily integrated into most existing GANs framework without excessive cost. In the experiment, we show that for various GANs frameworks of different objective functions, network architectures and datasets, our method can reduce the frequency spectrum discrepancy and improve the performance in the spatial domain.




Figure 6: The absolute difference of average spectra $|\mathbb{E}[\mathcal{F}(real)] - \mathbb{E}[\mathcal{F}(fake)]|$.

SSD-StyleGAN

Implementation Based on StyleGAN (Karras, Laine, and Aila 2019), We evaluate the effectiveness of our method on the FFHQ (Karras, Laine, and Aila 2019) dataset. It consists 70,000 high-quality images at 1024×1024 resolution. We use the same implementation as StyleGAN. In the discriminator, the activations are blurred before each downsampling layer for anti-aliasing. The training process is under a progressive growing manner (Karras et al. 2018) which starts from 8×8 to 1024×1024 . We apply the non-saturating loss (Goodfellow et al. 2014) as our adversarial loss with R_1 regularization (Mescheder, Geiger, and Nowozin 2018). We train all our models with Adam optimizer (Kingma and Ba 2015), setting $(\beta_1, \beta_2) = (0, 0.99)$. The total training time is 25M images. The hyperparameter λ is set to 0.5. The experiments are conducted on 4 Tesla V100 GPUs.

Evaluation in the Frequency Domain We first validate whether our proposed SSD-StyleGAN can reduce spectral distortions. We estimate the average spectrum $\mathbb{E}[\mathcal{F}(x)]$ by averaging over 5,000 images. Then we plot the absolute difference $|\mathbb{E}[\mathcal{F}(real)] - \mathbb{E}[\mathcal{F}(fake)]|$ between the two average spectra. As depicted in Fig. 6, compared to StyleGAN, the frequency spectrum discrepancy between images generated by SSD-StyleGAN and real data is significantly reduced. We also notice that both of the models have spectral distortions at the corners of the spectra that represent extremely high frequencies of images. In practice, due to image compression algorithms applied to real data, these high-frequency bands contain little information. Therefore, to discourage overfitting, GANs tend not to learn these extremely high frequencies, and these components of the generated images behave like white noise which has constant power density (Dzanic and Witherden 2019).

Evaluation in the Spatial Domain Table 1 reports the performance in the spatial domain. We adopt Fréchet Inception Distance (FID) (Heusel et al. 2017) to evaluate the perceptual quality of generated images, and perceptual path length (PPL) (Karras, Laine, and Aila 2019) to measure the degree of disentanglement of representations. Because the feature extractors used in these metrics are neural networks that map from a high-dimensional input space to a low-dimensional space, they also suffer from some degree of



Figure 7: Interpolations in the \mathcal{W} space. For each comparison, the top row represents the reconstructions of the interpolations of StyleGAN, while the bottom row represents the results of SSD-StyleGAN.

high-frequency loss and mainly measure the characteristics in the spatial domain. It is evident that our method performs better than StyleGAN on both these metrics. We attribute the performance improvement to alleviating the high frequencies missing problem in the discriminator. By reducing spectral distortions, it helps to reproduce the spatial distribution of real data, since the Fourier transform is a bijective mapping.

For qualitative evaluation, we utilize the recent embedding algorithm (Abdal, Qin, and Wonka 2019) to map a given image into the \mathcal{W} space of a pre-trained StyleGAN and then reconstruct back for comparison. We note that the models may memory images during training and produce good reconstructions. To avoid overfitting, we propose to compare the interpolations of StyleGAN and SSD-StyleGAN. As shown in Fig. 7, compared to StyleGAN, the results of our proposed method show a smoother morphing and have better details, which is consistent with the quantitative evaluation. Fig. 5 shows a collection of generations obtained from SSD-StyleGAN.

Estimating the Spectral Quality of Samples In this section, we evaluate the performance of the spectral classifier C for estimating the spectral quality of samples. As shown in Fig. 8, for real data and generations of SSD-StyleGAN, we present two images with high and low spectral quality scores. In general, the samples with high spectral quality scores display a clear portrait. However, the images with low spectral quality scores are often overexposed and lose details, or have some unusual high-frequency components (e.g., see background and headwear in the right column). The above observation shows the effectiveness of the spectral classifier C , which is beneficial to the learning process of SSD-StyleGAN.

SSD-SNGAN

Implementation Based on SNGAN (Miyato et al. 2018), we evaluate the proposed method on a range

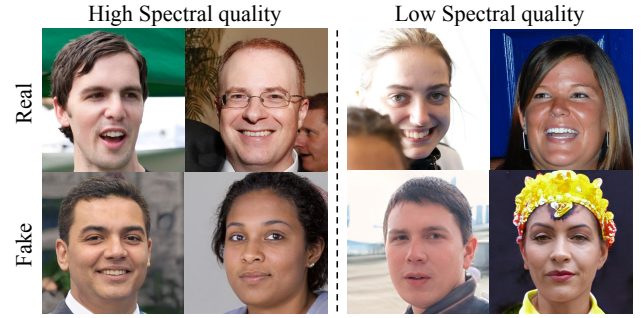


Figure 8: Spectral quality of real data and generations of SSD-StyleGAN.

Method	FID	Path length	
		full	end
StyleGAN	4.40	234.0	195.9
SSD-StyleGAN	4.06	229.1	189.7

Table 1: FID scores and perceptual path lengths (PPLs) on FFHQ (lower is better). PPLs are measured in the \mathcal{W} space.

of datasets including CIFAR100 (Krizhevsky and Hinton 2009), STL10 (Coates, Ng, and Lee 2011), and LSUN-bedroom (Yu et al. 2015). Since these datasets have various kinds of resolution, we mark them with the resolution: CIFAR100-32, STL10-48, and LSUN-128. We use the same training configurations as (Lee and Town 2020). We train all our models with Adam optimizer (Kingma and Ba 2015), setting $(\beta_1, \beta_2) = (0, 0.9)$. The learning rate is set to 0.0002, and the minibatch size is 64. The hyperparameter λ is set to 0.5. All models are trained on a single Tesla V100 GPU.

Baseline Models We compare our method against three baselines, including:

- **SNGAN** (Miyato et al. 2018) limits the spectral norm of the weight matrices in the discriminator for Lipschitz constraint. It adopts a ResNet (He et al. 2016) backbone and uses average-pooling as the downsampling layer in the discriminator. Different from StyleGAN, it utilizes the hinge version (Miyato et al. 2018) of the adversarial loss.
- **SNGAN+REG** (Durrall, Keuper, and Keuper 2020) adds a spectral regularization loss to the generator loss to penalize for synthesizing samples with abnormal spectra. The term can be written as $d(\phi(x^{real}), \phi(x^{fake}))$, where $d(\cdot, \cdot)$ measures the binary cross entropy.
- **SNGAN+DWT** (Liu et al. 2018) adopts discrete wavelet transform (DWT) as downsampling layer to avoid information loss. we use 2D Haar wavelet transform to decompose an input into an low-pass representation and three directions of high-frequency coefficients. Specifically, this DWT downsampling layer transforms the input raw images or a group of feature maps with height H , width W and channel C into a tensor of shape $(\frac{1}{2}H, \frac{1}{2}W, 4C)$.




Figure 9: Generations of SSD-SNGAN+REG trained on LSUN-bedroom at 128×128 .

Method	CIFAR100-32	STL10-48	LSUN-128
SNGAN	22.61	39.56	25.87
SNGAN+REG	21.39	38.16	16.95
SNGAN+DWT	25.95	40.46	148.05
SSD-SNGAN	19.28	36.41	15.17
SSD-SNGAN+REG	19.25	35.41	10.61

Table 2: FID scores on CIFAR100-32, STL10-48, and LSUN-128 (lower is better).

Results Table 2 reports the FID scores on CIFAR100-32, STL10-48, and LSUN-128. SNGAN+REG shows performance improvement over the baseline SNGAN, indicating that utilizing spectral information is effective. Compared with SNGAN, the scores of SNGAN+DWT are higher, especially for LSUN-128. We conjecture that because the DWT downsampling layer remains all the high-frequency information of the input, which contains both details and noises, it is difficult for the model to learn meaningful semantic representation. Moreover, as pointed out by (Wang et al. 2020a), learning high-frequency information may degrade the robustness and generalization of a model. Since LSUN-128 has higher resolution and contains more high-frequency information, performance degradation of SNGAN+DWT on this dataset is more dramatic. Our method gets lower FID scores than SNGAN+REG and SNGAN+DWT, indicating it is a smarter way to utilize high-frequency information of images. It is remarkable that SSD-SNGAN+REG achieves the best FID scores among all the datasets. These two techniques have complementary advantages, since they encourage the model to utilize high-frequency information from the aspects of generator and discriminator respectively. Fig. 9 shows a collection of generations obtained from SSD-

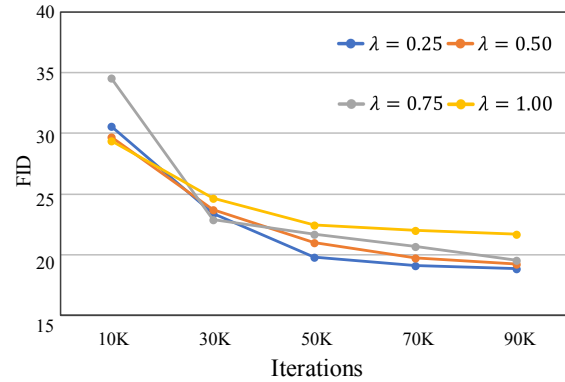


Figure 10: FID scores over the course of training with different hyperparameter λ .

SNGAN+REG on LSUN-128.

Robustness of the Hyperparameter λ

In Eq. 6, we introduce a hyperparameter λ to control the relative importance of spatial realness and spectral realness. Here, we evaluate different hyperparameter settings on CIFAR100-32 to investigate the robustness of λ . The baseline model is SNGAN, and other training settings remain the same as the previous section. Fig. 10 compares FID scores over the course of training when setting different values for λ . Note that the case of $\lambda = 1$ is the baseline model. We observe that the proposed approach yields consistent performance improvements and enjoys considerable tolerance for the selection of the hyperparameter λ .

Conclusion and Outlook

In this paper, we delve into why existing GANs fail to reproduce the spectral distribution of real data and reveal the issue of high frequencies missing in the discriminator. To alleviate the issue, we introduce SSD-GAN, whose discriminator is enhanced to measure the realness of samples in both the spatial and spectral domains. We provide empirical evidence that the proposed SSD-GAN can reduce frequency spectrum discrepancy, thus achieving performance improvement in the image domain.

Frequency analysis provides a novel perspective for analyzing and understanding GANs. It also opens some avenues for future research. First, although recent GAN based models achieve good performance under existing metrics, the generated samples of these models can easily be distinguished from real data in the frequency domain. A metric that quantitatively measure the performance of generative models in the frequency domain would promote the image synthesis community. Moreover, besides the discriminator in GANs, many machine learning tasks involve learning a mapping from a high-dimensional input space to a low-dimensional space. Constructing a general network architecture to learn semantic representations while taking high frequencies into consideration is interesting and challenging for future work.

Acknowledgements

We would like to thank Xiaoming Yu, Hengtao Zhang, Yurui Ren, Nannan Li, and our anonymous reviewers for their valuable feedback and helpful discussions. This work is supported by National Key R&D Program of China (No. 2020AAA0103500) and National Natural Science Foundation of China and Guangdong Province Scientific Research on Big Data (No. U1611461).

References

- Abdal, R.; Qin, Y.; and Wonka, P. 2019. Image2stylegan: How to embed images into the stylegan latent space? In *ICCV*.
- Chen, T.; Zhai, X.; Ritter, M.; Lucic, M.; and Houthby, N. 2019. Self-supervised gans via auxiliary rotation loss. In *CVPR*.
- Coates, A.; Ng, A.; and Lee, H. 2011. An analysis of single-layer networks in unsupervised feature learning. In *AISTATS*.
- Durall, R.; Keuper, M.; and Keuper, J. 2020. Watch your Up-Convolution: CNN Based Generative Deep Neural Networks are Failing to Reproduce Spectral Distributions. In *CVPR*.
- Dzanic, T.; and Witherden, F. 2019. Fourier Spectrum Discrepancies in Deep Network Generated Images. *arXiv preprint arXiv:1911.06465*.
- Dziedzic, A.; Paparrizos, J.; Krishnan, S.; Elmore, A.; and Franklin, M. 2019. Band-limited Training and Inference for Convolutional Neural Networks. In *ICML*.
- Frank, J.; Eisenhofer, T.; Schönherr, L.; Fischer, A.; Kolossa, D.; and Holz, T. 2020. Leveraging Frequency Analysis for Deep Fake Image Recognition. In *ICML*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *NIPS*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *CVPR*.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *ICLR*.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *CVPR*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Krizhevsky, A.; and Hinton, G. 2009. Learning multiple layers of features from tiny images. *Tech Report*.
- LeCun, Y.; Boser, B. E.; Denker, J. S.; Henderson, D.; Howard, R. E.; Hubbard, W. E.; and Jackel, L. D. 1990. Handwritten digit recognition with a back-propagation network. In *NIPS*.
- Lee, K. S.; and Town, C. 2020. Mimicry: Towards the Reproducibility of GAN Research. In *CVPRW*.
- Liu, P.; Zhang, H.; Zhang, K.; Lin, L.; and Zuo, W. 2018. Multi-level wavelet-CNN for image restoration. In *CVPRW*.
- Mescheder, L.; Geiger, A.; and Nowozin, S. 2018. Which Training Methods for GANs do actually Converge? In *ICML*.
- Miyato, T.; Kataoka, T.; Koyama, M.; and Yoshida, Y. 2018. Spectral Normalization for Generative Adversarial Networks. In *ICLR*.
- Nyquist, H. 1928. Certain topics in telegraph transmission theory. *Transactions of the American Institute of Electrical Engineers*.
- Peng, X. B.; Kanazawa, A.; Toyer, S.; Abbeel, P.; and Levine, S. 2019. Variational Discriminator Bottleneck: Improving Imitation Learning, Inverse RL, and GANs by Constraining Information Flow. In *ICLR*.
- Rahaman, N.; Baratin, A.; Arpit, D.; Draxler, F.; Lin, M.; Hamprecht, F.; Bengio, Y.; and Courville, A. 2019. On the spectral bias of neural networks. In *ICML*.
- Ren, Y.; Yu, X.; Zhang, R.; Li, T. H.; Liu, S.; and Li, G. 2019. Structureflow: Image inpainting via structure-aware appearance flow. In *ICCV*.
- Wang, H.; Wu, X.; Huang, Z.; and Xing, E. P. 2020a. High-frequency Component Helps Explain the Generalization of Convolutional Neural Networks. In *CVPR*.
- Wang, S.-Y.; Wang, O.; Zhang, R.; Owens, A.; and Efros, A. A. 2020b. CNN-generated images are surprisingly easy to spot... for now. In *CVPR*.
- Xu, Z.-Q. J.; Zhang, Y.; and Xiao, Y. 2019. Training behavior of deep neural network in frequency domain. In *ICONIP*.
- Yang, D.; Hong, S.; Jang, Y.; Zhao, T.; and Lee, H. 2018. Diversity-Sensitive Conditional Generative Adversarial Networks. In *ICLR*.
- Yu, F.; Zhang, Y.; Song, S.; Seff, A.; and Xiao, J. 2015. LSUN: Construction of a Large-scale Image Dataset using Deep Learning with Humans in the Loop. *arXiv preprint arXiv:1506.03365*.
- Yu, X.; Chen, Y.; Liu, S.; Li, T.; and Li, G. 2019. Multi-mapping image-to-image translation via learning disentanglement. In *NIPS*.
- Zhang, R. 2019. Making Convolutional Networks Shift-Invariant Again. In *ICML*.
- Zhang, X.; Karaman, S.; and Chang, S.-F. 2019. Detecting and simulating artifacts in gan fake images. In *WIFS*.