

Focal Frequency Loss for Image Reconstruction and Synthesis

Liming Jiang¹ Bo Dai¹ Wayne Wu² Chen Change Loy^{1✉}

¹S-Lab, Nanyang Technological University ²SenseTime Research

{liming002, bo.dai, ccloy}@ntu.edu.sg wuwenyan@sensetime.com

Abstract

Image reconstruction and synthesis have witnessed remarkable progress thanks to the development of generative models. Nonetheless, gaps could still exist between the real and generated images, especially in the frequency domain. In this study, we show that narrowing gaps in the frequency domain can ameliorate image reconstruction and synthesis quality further. We propose a novel focal frequency loss, which allows a model to adaptively focus on frequency components that are hard to synthesize by down-weighting the easy ones. This objective function is complementary to existing spatial losses, offering great impedance against the loss of important frequency information due to the inherent bias of neural networks. We demonstrate the versatility and effectiveness of focal frequency loss to improve popular models, such as VAE, pix2pix, and SPADE, in both perceptual quality and quantitative performance. We further show its potential on StyleGAN2.^{1,2}

1. Introduction

We have seen remarkable progress in image reconstruction and synthesis along with the development of generative models [21, 38, 16, 36, 65], and the progress continues with the emergence of various powerful deep learning-based approaches [34, 50, 51, 64]. Despite their immense success, one could still observe the gaps between the real and generated images in certain cases.

These gaps are sometimes manifested in the form of artifacts that are discernible. For instance, upsampling layers using transposed convolutions tend to produce checkerboard artifacts [49]. The gaps, in some other cases, may only be revealed through the frequency spectrum analysis. Recent studies [68, 83, 24] in media forensics have shown some notable periodic patterns in the frequency spectra of manipulated images, which may be consistent with artifacts in the spatial domain. In Figure 1, we show some paired examples of real images and the fake ones generated by typical generative models for image reconstruction and synthesis. It is observed that the frequency domain gap between




Figure 1. Frequency domain gaps between the real and the generated images by typical generative models in image reconstruction and synthesis. Vanilla AE [21] loses important frequencies, leading to blurry images (Row 1 and 2). VAE [38] biases to a limited spectrum region (Row 3), losing high-frequency information (outer regions and corners). Unnatural periodic patterns can be spotted on the spectra of images generated by GAN (pix2pix) [26] (Row 4), consistent with the observable checkerboard artifacts (zoom in for view). In some cases, a frequency spectrum region shift occurs to GAN-generated images (Row 5).

the real and fake images may be a common issue for these methods, albeit in slightly different forms.

The observed gaps in the frequency domain may be imputed to some inherent bias of neural networks when applied to reconstruction and synthesis tasks. Fourier analysis highlights a phenomenon called *spectral bias* [55, 46, 62], a learning bias of neural networks towards low-frequency functions. Thus, generative models tend to eschew frequency components that are hard to synthesize, i.e., hard frequencies, and converge to an inferior point. *F-Principle* [77] shows that the priority of fitting certain fre-

¹ GitHub: <https://github.com/EndlessSora/focal-frequency-loss>.

² Project page: <https://www.mmlab-ntu.com/project/ffl/index.html>.

quencies in a network is also different throughout the training, usually from low to high. Consequently, it is difficult for a model to maintain important frequency information as it tends to generate frequencies with a higher priority.

In this paper, we carefully study the frequency domain gap between real and fake images and explore ways to ameliorate reconstruction and synthesis quality by narrowing this gap. Existing methods [38, 26, 50] usually adopt pixel losses in the spatial domain, while spatial domain losses hardly help a network find hard frequencies and synthesize them, in that every pixel shares the same significance for a certain frequency. In comparison, we transform both the real and generated samples to their frequency representations using the standard discrete Fourier transform (DFT). The images are decomposed into sines and cosines, exhibiting periodic properties. Each coordinate value on the frequency spectrum depends on all the image pixels in the spatial domain, representing a specific spatial frequency. Explicitly minimizing the distance of coordinate values on the real and fake spectra can help networks easily locate difficult regions on the spectrum, *i.e.*, hard frequencies.

To tackle these hard frequencies, inspired by hard example mining [12, 59] and focal loss [41], we propose a simple yet effective frequency-level objective function, named *focal frequency loss*. We map each spectrum coordinate value to a Euclidean vector in a two-dimensional space, with both the amplitude and phase information of the spatial frequency put under consideration. The proposed loss function is defined by the scaled Euclidean distance of these vectors by down-weighting easy frequencies using a dynamic spectrum weight matrix. Intuitively, the matrix is updated on the fly according to a non-uniform distribution on the current loss of each frequency during training. The model will then rapidly focus on hard frequencies and progressively refine the generated frequencies to improve image quality.

The main **contribution** of this work is a novel focal frequency loss that directly optimizes generative models in the frequency domain. We carefully motivate how a loss can be built on a space where frequencies of an image can be well represented and distinguished, facilitating optimization in the frequency dimension. We further explain the way that enables a model to focus on hard frequencies, which may be pivotal for quality improvement. Extensive experiments demonstrate the effectiveness of the proposed loss on representative baselines [21, 38, 26, 50], and the loss is complementary to existing spatial domain losses such as perceptual loss [30]. We further show the potential of focal frequency loss to improve state-of-the-art StyleGAN2 [34].

2. Related Work

Image reconstruction and synthesis. Autoencoders (AE) [21, 38] and generative adversarial networks (GAN) [16] are two popular models for image reconstruction and synthesis. The vanilla AE [21] aims at learning latent codes

while reconstructing images. It is typically used for dimensionality reduction and feature learning. Autoencoders have been widely used to generate images since the development of variational autoencoders (VAE) [38, 37]. Their applications have been extended to various tasks, *e.g.*, face manipulation [2, 1, 28, 27]. GAN [16, 47, 53], on the other hand, is extensively applied in face generation [32, 33, 34], image-to-image translation [26, 85, 8, 29], style transfer [42, 23], and semantic image synthesis [69, 50, 43]. Existing approaches usually apply spatial domain loss functions, *e.g.*, perceptual loss [30], to improve quality while seldom consider optimization in the frequency domain. Spectral regularization [11] presents a preliminary attempt. Different from [30, 11], the proposed focal frequency loss dynamically focuses the model on hard frequencies by down-weighting the easy ones and ameliorates image quality through the frequency domain directly. Some concurrent works on image reconstruction and synthesis via the frequency domain include [5, 15, 31].

Frequency domain analysis of neural networks. In addition to the studies [55, 46, 62, 77] we discussed in the introduction, we highlight some recent works that analyze neural networks through the frequency domain. Using coordinate-based MLPs, Fourier features [62, 56] and positional encoding [46, 66] are adopted to recover missing high frequencies in single image regression problems. Besides, several studies have incorporated frequency analysis with network compression [17, 76, 78, 7, 18] and feature reduction [39, 70] to accelerate the training and inference of networks. The application areas of the frequency domain analysis have been further extended, including media forensics [68, 83, 24, 13], super-resolution [14, 72], generalization analysis [67, 22], magnetic resonance imaging [61], image rescaling [73], *etc.* Despite the wide exploration of various problems, improving reconstruction and synthesis quality via the frequency domain remains much less explored.

Hard example processing. Hard example processing is widely explored in object detection and image classification to address the class imbalance problem. A common solution is to use a bootstrapping technique called hard example mining [59, 12], where a representative method is online hard example mining (OHEM) [59]. The training examples are sampled following the current loss of each example to modify the stochastic gradient descent. The model is encouraged to learn hard examples more to boost performance. An alternative solution is focal loss [41], which is a scaled cross-entropy loss. The scaling factor down-weights the contribution of easy examples during training so that a model can focus on learning hard examples. The proposed focal frequency loss is inspired by these techniques.

3. Focal Frequency Loss

To formulate our method, we explicitly exploit the frequency representation of images (Section 3.1), facilitating

the network to locate the hard frequencies. We then define a frequency distance (Section 3.2) to quantify the differences between images in the frequency domain. Finally, we adopt a dynamic spectrum weighting scheme (Section 3.3) that allows the model to focus on the on-the-fly hard frequencies.

3.1. Frequency Representation of Images

In this section, we revisit and highlight several key concepts of the discrete Fourier transform. We demonstrate the effect of missing frequencies in the image and the advantage of frequency representation for locating hard frequencies.

Discrete Fourier transform (DFT) is a complex-valued function that converts a discrete finite signal into its constituent frequencies, *i.e.*, complex exponential waves. An image³ can be treated as a two-dimensional discrete finite signal with only real numbers. Thus, to convert an image into its frequency representation, we perform the 2D discrete Fourier transform:

$$F(u, v) = \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} f(x, y) \cdot e^{-i2\pi(\frac{ux}{M} + \frac{vy}{N})}, \quad (1)$$

where the image size is $M \times N$; (x, y) denotes the coordinate of an image pixel in the spatial domain; $f(x, y)$ is the pixel value; (u, v) represents the coordinate of a spatial frequency on the frequency spectrum; $F(u, v)$ is the complex frequency value; e and i are Euler’s number and the imaginary unit, respectively. Following Euler’s formula:

$$e^{i\theta} = \cos \theta + i \sin \theta, \quad (2)$$

the natural exponential function in Eq. (1) can be written as:

$$e^{-i2\pi(\frac{ux}{M} + \frac{vy}{N})} = \cos 2\pi \left(\frac{ux}{M} + \frac{vy}{N} \right) - i \sin 2\pi \left(\frac{ux}{M} + \frac{vy}{N} \right). \quad (3)$$

According to Eq. (1) and Eq. (3), the image is decomposed into orthogonal sine and cosine functions, constituting the imaginary and the real part of the frequency value, respectively, after applied 2D DFT. Each sine or cosine can be regarded as a binary function of (x, y) , where its angular frequency is decided by the spectrum position (u, v) . The mixture of these sines and cosines provides both the horizontal and vertical frequencies of an image. Therefore, spatial frequency manifests as the 2D sinusoidal components in the image. The spectrum coordinate (u, v) also represents the angled direction of a spatial frequency (visualizations can be found in the *Appendix*), and $F(u, v)$ shows the “response” of the image to this frequency. Due to the periodicity of trigonometric functions, the frequency representation of an image also acquires periodic properties.

Note that in Eq. (1), $F(u, v)$ is the sum of a function that traverses every image pixel in the spatial domain, hence a

³ For simplicity, the formulas in this section are applied to gray-scale images, while the extension to color images is straightforward by processing each channel separately in the same way.




Figure 2. Standard bandlimiting operations on the frequency spectrum with the origin (low frequencies) center shifted and respective images in the spatial domain. These manual operations can be regarded as a simulation to show the effect of missing frequencies.

specific spatial frequency on the spectrum depends on all the image pixels. For an intuitive visualization, we suppress the *single* center point (the lowest frequency) of the spectrum (Column 2 of Figure 2), leading to *all* the image pixels being affected. To further ascertain the spatial frequency at the different regions on the spectrum, we perform some other standard bandlimiting operations and visualize their physical meanings in the spatial domain (Figure 2). A low-pass filter (Column 3), *i.e.*, missing high frequencies, causes blur and typical ringing artifacts. A high-pass filter (Column 4), *i.e.*, missing low frequencies, tends to retain edges and boundaries. Interestingly, a simple band-stop filter (Column 5), *i.e.*, missing certain frequencies, produces visible common checkerboard artifacts (zoom in for view).

Observably, the losses of different regions on the frequency spectrum correspond to different artifacts on the image. One may deduce that compensating for these losses may reduce artifacts and improve image reconstruction and synthesis quality. The analysis here shows the value of using the frequency representation of images for profiling and locating different frequencies, especially the hard ones.

3.2. Frequency Distance

To devise a loss function for the missing frequencies, we need a distance metric that quantifies the differences between real and fake images in the frequency domain. The distance has to be differentiable to support stochastic gradient descent. In the frequency domain, the data objects are different spatial frequencies on the frequency spectrum, appearing as different 2D sinusoidal components in an image. To design our frequency distance, we further study the real and imaginary part of the complex value $F(u, v)$ in Eq. (1).

Let $R(u, v) = a$ and $I(u, v) = b$ be the real and the imaginary part of $F(u, v)$, respectively. $F(u, v)$ can be rewritten as:

$$F(u, v) = R(u, v) + I(u, v)i = a + bi. \quad (4)$$

According to the definition of 2D discrete Fourier transform, there are two key elements in $F(u, v)$. The first el-




Figure 3. The necessity of both amplitude and phase information for a frequency distance verified by *single-image reconstruction*. “Amplitude/phase only” means solely applying Eq. (5)/(6) to calculate the distance between the real and reconstructed images.

ement is *amplitude*, which is defined as:

$$|F(u, v)| = \sqrt{R(u, v)^2 + I(u, v)^2} = \sqrt{a^2 + b^2}. \quad (5)$$

Amplitude manifests the energy, *i.e.*, how strongly an image responds to the 2D sinusoidal wave with a specific frequency. We typically show the amplitude as an informative visualization of the frequency spectrum (*e.g.*, Figure 1 and 2). The second element is *phase*, which is written as:

$$\angle F(u, v) = \arctan\left(\frac{I(u, v)}{R(u, v)}\right) = \arctan\frac{b}{a}. \quad (6)$$

Phase represents the shift of a 2D sinusoidal wave from the wave with the origin value (the beginning of a cycle).

A frequency distance should consider both the amplitude and the phase as they capture different information of an image. We show a single-image reconstruction experiment in Figure 3. Merely minimizing the amplitude difference returns a reconstructed image with irregular color patterns. Conversely, using only the phase information, the synthesized image resembles a noise. A faithful reconstruction can only be achieved by considering both amplitude and phase.

Our solution is to map each frequency value to a Euclidean vector in a two-dimensional space (*i.e.*, a plane). Following the standard definition of a complex number, the real and imaginary parts correspond to the x -axis and y -axis, respectively. Let $F_r(u, v) = a_r + b_r i$ be the spatial frequency value at the spectrum coordinate (u, v) of the real image, and the corresponding $F_f(u, v) = a_f + b_f i$ with the similar meaning w.r.t. the fake image. We denote \vec{r}_r and \vec{r}_f as two respective vectors mapped from $F_r(u, v)$ and $F_f(u, v)$ (see Figure 4). Based on the definition of amplitude and phase, we note that the vector magnitude $|\vec{r}_r|$ and $|\vec{r}_f|$ correspond to the amplitude, and the angle θ_r and θ_f correspond to the phase. Thus, the frequency distance corresponds to the distance between \vec{r}_r and \vec{r}_f , which considers both the vector magnitude and angle. We use the (squared) Euclidean distance for a single frequency:

$$d(\vec{r}_r, \vec{r}_f) = \|\vec{r}_r - \vec{r}_f\|_2^2 = |F_r(u, v) - F_f(u, v)|^2. \quad (7)$$

The frequency distance between the real and fake images can be written as the average value:

$$d(F_r, F_f) = \frac{1}{MN} \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} |F_r(u, v) - F_f(u, v)|^2. \quad (8)$$




Figure 4. Frequency distance between \vec{r}_r and \vec{r}_f mapped from two corresponding real and fake frequency values $F_r(u, v)$ and $F_f(u, v)$ at the spectrum position (u, v) . The Euclidean distance (purple line) is used, considering both the amplitude (magnitude $|\vec{r}_r|$ and $|\vec{r}_f|$) and phase (angle θ_r and θ_f) information.

3.3. Dynamic Spectrum Weighting

The frequency distance we defined in Eq. (8) quantitatively compares the real and fake images in the frequency domain. However, a direct use of Eq. (8) as a loss function is not helpful in coping with hard frequencies since the weight of each frequency is identical. A model would still bias to easy frequencies due to the inherent bias.

Inspired by hard example mining [12, 59] and focal loss [41], we formulate our method to focus the training on the hard frequencies. To implement this, we introduce a spectrum weight matrix to down-weight the easy frequencies. The spectrum weight matrix is dynamically determined by a non-uniform distribution on the current loss of each frequency during training. Each image has its own spectrum weight matrix. The shape of the matrix is the same as that of the spectrum. The matrix element $w(u, v)$, *i.e.*, the weight for the spatial frequency at (u, v) , is defined as:

$$w(u, v) = |F_r(u, v) - F_f(u, v)|^\alpha, \quad (9)$$

where α is the scaling factor for flexibility ($\alpha = 1$ in our experiments). We further normalize the matrix values into the range $[0, 1]$, where the weight 1 corresponds to the currently most lost frequency, and the easy frequencies are down-weighted. The gradient through the spectrum weight matrix is locked, so it only serves as the weight for each frequency.

By performing the Hadamard product for the spectrum weight matrix and the frequency distance matrix, we have the *full form* of the focal frequency loss (FFL):

$$\text{FFL} = \frac{1}{MN} \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} w(u, v) |F_r(u, v) - F_f(u, v)|^2. \quad (10)$$

The focal frequency loss can be seen as a weighted average of the frequency distance between the real and fake images. It focuses the model on synthesizing hard frequencies by down-weighting easy frequencies. Besides, the focused region is updated on the fly to complement the immediate

hard frequencies, thus progressively refining the generated images and being adaptable to different methods.

In practice, to apply the proposed focal frequency loss to a model, we first transform both the real and fake images into their frequency presentations using the 2D DFT. We then perform the orthonormalization for each frequency value $F(u, v)$, *i.e.*, dividing it by \sqrt{MN} , so that the 2D DFT is unitary to ensure a smooth gradient. Finally, we employ Eq. (10) to calculate the focal frequency loss. We note that the exact form of focal frequency loss is not crucial. Some studies on the loss variants are provided in the *Appendix*.

4. Experiments

4.1. Settings

Baselines. We start from image reconstruction by vanilla AE [21] (*i.e.*, a simple 2-layer MLP) and VAE [38] (*i.e.*, CNN-based). We then study unconditional image synthesis using VAE, *i.e.*, generating images from the Gaussian noise. Besides, we also investigate conditional image synthesis using GAN-based methods. Specifically, we select two typical image-to-image translation approaches, *i.e.*, pix2pix [26] and SPADE [50]. We further explore the potential of focal frequency loss (FFL) on state-of-the-art StyleGAN2 [34]. In addition, we compare FFL with relevant losses [30, 11]. The implementation details are provided in the *Appendix*.

Datasets. We use a total of seven datasets. The datasets vary in types, sizes, and resolutions. For vanilla AE, we exploit the Describable Textures Dataset (DTD) [9] and CelebA [44]. For VAE, we use CelebA and CelebA-HQ [32] with different resolutions. For pix2pix, we utilize the officially prepared CMP Facades [54] and edges \rightarrow shoes [79] datasets. For SPADE, we select two challenging datasets, *i.e.*, Cityscapes [10] and ADE20K [84]. For StyleGAN2, we reuse CelebA-HQ. Please refer to the *Appendix* for the dataset details.

Evaluation metrics. To evaluate frequency domain difference, we introduce a frequency-level metric, named Log Frequency Distance (LFD), which is defined by a modified version of Eq. (8):

$$\text{LFD} = \log \left[\frac{1}{MN} \left(\sum_{u=0}^{M-1} \sum_{v=0}^{N-1} |F_r(u, v) - F_f(u, v)|^2 \right) + 1 \right], \quad (11)$$

where the logarithm is only used to scale the value into a reasonable range. A lower LFD is better. Note that LFD is a full reference metric (*i.e.*, requiring the ground truth image), so we use it in the reconstruction tasks.

Besides, we integrate the evaluation protocols from prior works [46, 4, 50, 29]. Specifically, we employ FID (lower is better) [20] for all tasks. For the reconstruction tasks of vanilla AE and VAE, we use PSNR (higher is better), SSIM (higher is better) [71], and LPIPS (lower is better) [82] in addition to LFD and FID. For the synthesis tasks of VAE,




Figure 5. **Vanilla AE image reconstruction** results on the DTD (64×64) and CelebA (64×64) datasets.

Table 1. The PSNR (higher is better), SSIM (higher is better), LPIPS (lower is better), FID (lower is better) and LFD (lower is better) scores for the **vanilla AE image reconstruction** trained with/without the focal frequency loss (FFL).

Dataset	FFL	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	LFD \downarrow
DTD	w/o	20.133	0.407	0.414	246.870	14.764
	w/	20.151	0.400	0.404	240.373	14.760
CelebA	w/o	20.044	0.568	0.237	97.035	14.785
	w/	21.703	0.642	0.199	83.801	14.403

pix2pix, and StyleGAN2, we apply IS (higher is better) [58] in addition to FID. For SPADE (task-specific method for semantic image synthesis), besides FID, we follow their paper [50] to use mIoU (higher is better) and pixel accuracy (accu, higher is better) for the segmentation performance of synthesized images. We use DRN-D-105 [80] for Cityscapes and UperNet101 [74] for ADE20K.

4.2. Results and Analysis

Vanilla AE. The results of vanilla AE [21] image reconstruction are shown in Figure 5. On DTD, without the focal frequency loss (FFL), the vanilla AE baseline synthesizes blurry images, which lack sufficient texture details and only contain some low-frequency information. With FFL, the reconstructed images become clearer and show more texture details. The results on CelebA show that FFL improves a series of quality problems, *e.g.*, face blur (Column 4), identity shift (Column 5), and expression loss (Column 6).

The quantitative evaluation results are presented in Table 1. Adding the proposed FFL to the vanilla AE baseline leads to a performance boost in most cases on the DTD and CelebA datasets w.r.t. five evaluation metrics. We note that the performance boost on CelebA is larger, indicating the effectiveness of FFL for the natural images.

VAE. The results of VAE [38] image reconstruction and unconditional image synthesis on CelebA are shown in Figure 6. For reconstruction, FFL helps the VAE model better retain the image clarity (Column 1), expression (Column 2), and skin color (Column 3). The unconditional synthesis results (Column 4, 5, 6) show that the quality of generated images is improved after applying FFL. The generated faces become clearer and gain more texture details. For a higher resolution, we present the VAE reconstruction and synthe-




Figure 6. VAE image reconstruction and unconditional image synthesis results on the CelebA (64×64) dataset.




Figure 7. VAE image reconstruction and unconditional image synthesis results on the CelebA-HQ (256×256) dataset.

Table 2. The PSNR (higher is better), SSIM (higher is better), LPIPS (lower is better), FID (lower is better) and LFD (lower is better) scores for the VAE image reconstruction trained with/without the focal frequency loss (FFL).

Dataset	FFL	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	LFD \downarrow
CelebA	w/o	19.961	0.606	0.217	69.900	14.804
	w/	22.954	0.723	0.143	49.689	14.115
CelebA-HQ	w/o	21.310	0.616	0.367	71.081	17.266
	w/	22.253	0.637	0.344	59.470	17.049

Table 3. The FID (lower is better) and IS (higher is better) scores for the VAE unconditional image synthesis trained with/without the focal frequency loss (FFL).

Dataset	FFL	FID \downarrow	IS \uparrow
CelebA	w/o	80.116	1.873
	w/	71.050	2.010
CelebA-HQ	w/o	93.778	2.057
	w/	84.472	2.060

sis results on CelebA-HQ in Figure 7. By adding FFL to the VAE baseline, the reconstructed images keep more original image information, *e.g.*, mouth color (Column 2) and opening angle (Column 1). Besides, high-frequency details on the hair are clearly enhanced (Column 1). For unconditional image synthesis, FFL helps reduce artifacts and ameliorates the perceptual quality of synthesized images.

The quantitative test results of VAE image reconstruction are shown in Table 2. Adding FFL to the VAE baseline



Figure 8. pix2pix image-to-image translation results on CMP Facades (256×256) and edges \rightarrow shoes (256×256) datasets.

Table 4. The FID (lower is better) and IS (higher is better) scores for the pix2pix image-to-image translation trained with/without the focal frequency loss (FFL).

Dataset	FFL	FID \downarrow	IS \uparrow
CMP Facades	w/o	128.492	1.571
	w/	123.773	1.738
edges \rightarrow shoes	w/o	80.279	2.674
	w/	74.359	2.804

achieves better performance w.r.t. all the metrics. Besides, both FID and IS are better in the unconditional image synthesis task (Table 3), indicating that the generated images are clearer and more photorealistic. The results suggests the effectiveness of the focal frequency loss in helping VAE to improve image reconstruction and synthesis quality.

pix2pix. For conditional image synthesis, the results of pix2pix [26] image-to-image translation (I2I) are shown in Figure 8. On CMP Facades, FFL improves the image synthesis quality of pix2pix by reducing unnatural colors (Column 1) or the black artifacts on the building (Column 2). Meanwhile, the semantic information alignment with the mask becomes better after applying FFL. For the edges \rightarrow shoes translation, pix2pix baseline sometimes introduces colored checkerboard artifacts to the white background (Column 3, zoom in for view). Besides, atypical colors appear in certain cases (Column 4). In comparison, the model trained with FFL yields fewer artifacts.

The quantitative evaluation results of pix2pix image-to-image translation are shown in Table 4. FFL contributes to a performance boost on both of the two datasets. The results of the pix2pix baseline show the adaptability of the focal frequency loss for the image-to-image translation problem.

SPADE. We further explore semantic image synthesis (*i.e.*, synthesizing a photorealistic image from a semantic seg-



Figure 9. **StyleGAN2 unconditional image synthesis** results (without truncation) and the mini-batch average spectra (adjusted to better contrast) on the **CelebA-HQ** (256×256) dataset.

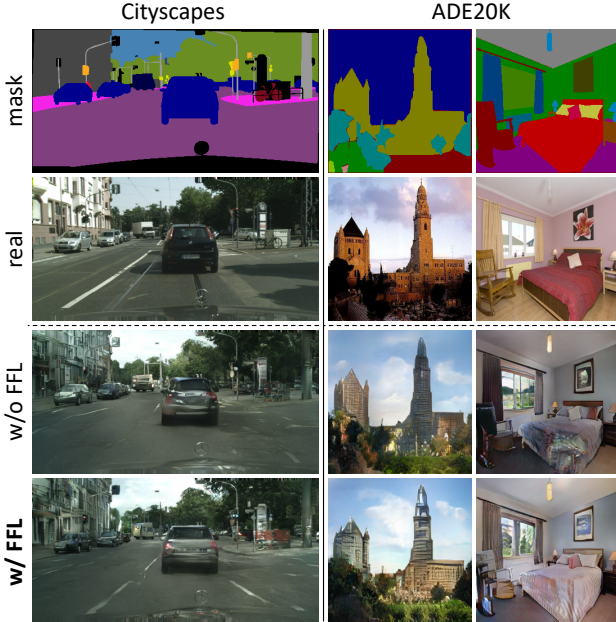


Figure 10. **SPADE semantic image synthesis** results on the **Cityscapes** (512×256) and **ADE20K** (256×256) datasets.

mentation mask) on more challenging datasets. The results of SPADE [50] are shown in Figure 10. In the street scene of Cityscapes (Column 1), SPADE baseline distorts the car and road, missing some important details (e.g., road line). The model trained with FFL demonstrates better perceptual quality for these details. In the outdoor scene of ADE20K (Column 2), applying FFL to SPADE boosts its ability to generate details on the buildings. Besides, for the ADE20K indoor images (Column 3), SPADE baseline produces some abnormal artifacts in certain cases. The model trained with the proposed FFL synthesizes more photorealistic images.

The quantitative test results are presented in Table 5 (the values used for comparison are taken from [50]). We compare SPADE trained with/without FFL against a series of open-source task-specific semantic image synthesis meth-

Table 5. The mIoU (higher is better), pixel accuracy (accu, higher is better) and FID (lower is better) scores for the **SPADE semantic image synthesis** trained with/without the focal frequency loss (FFL) compared to a series of task-specific methods.

Method	Cityscapes			ADE20K		
	mIoU↑	accu↑	FID↓	mIoU↑	accu↑	FID↓
CRN [6]	52.4	77.1	104.7	22.4	68.8	73.3
SIMS [52]	47.2	75.5	49.7	N/A	N/A	N/A
pix2pixHD [69]	58.3	81.4	95.0	20.3	69.2	81.8
SPADE [50]	62.3	81.9	71.8	38.5	79.9	33.9
SPADE + FFL	64.2	82.5	59.5	42.9	82.4	33.7

Table 6. The FID (lower is better) and IS (higher is better) scores for the **StyleGAN2 unconditional image synthesis** trained with/without the focal frequency loss (FFL).

Dataset	FFL	FID↓	IS↑
CelebA-HQ (256×256)	w/o	5.696	3.383
	w/	4.972	3.432

ods [6, 52, 69]. SIMS [52] obtains the best FID but poor segmentation scores on Cityscapes in that it directly stitches the training image patches from a memory bank while not keeping the exactly consistent positions. Without modifying the SPADE network structure, training with FFL contributes a further performance boost, greatly outperforming the benchmark methods, which suggests the effectiveness of FFL for semantic image synthesis.

StyleGAN2. We apply FFL to the mini-batch average spectra of the real images and the generated images by the state-of-the-art unconditional image synthesis method, i.e., StyleGAN2 [34], intending to narrow the frequency distribution gap and improve quality further. The results on CelebA-HQ (256×256) without truncation [33, 34] are shown in Figure 9. Although StyleGAN2 (w/o FFL) generates photorealistic images in most cases, some tiny artifacts can still be spotted on the background (Column 2 and 4) and face (Column 5). Applying FFL, such artifacts are reduced, ameliorating synthesis quality further. Observably, the frequency domain gaps between mini-batch average spectra are clearly mitigated by FFL (Column 8). Some higher-resolution re-

Table 7. **Comparison** of our focal frequency loss (FFL) **with relevant losses**, *i.e.*, perceptual loss (PL), spectral regularization (SpReg), and another transformation form for FFL, *i.e.*, discrete cosine transform (DCT), in different image reconstruction and synthesis tasks.

Method	VAE reconstruction (CelebA)					VAE synthesis (CelebA)		pix2pix I2I (edges \rightarrow shoes)	
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	LFD \downarrow	FID \downarrow	IS \uparrow	FID \downarrow	IS \uparrow
baseline	19.961	0.606	0.217	69.900	14.804	80.116	1.873	80.279	2.674
+ PL [30]	20.964	0.658	0.143	62.795	14.573	78.825	1.788	78.916	2.722
+ SpReg [11]	19.974	0.607	0.218	69.118	14.796	78.079	1.898	79.300	2.700
+ FFL (DCT)	22.677	0.711	0.150	51.536	14.179	71.827	1.932	79.045	2.754
+ FFL (Ours)	22.954	0.723	0.143	49.689	14.115	71.050	2.010	74.359	2.804

sults can be found in the *Appendix*.

The quantitative results are reported in Table 6. FFL improves both FID and IS, in line with the visual quality enhancement. The results on StyleGAN2 show the potential of FFL to boost state-of-the-art baseline performance.

Comparison with relevant losses. For completeness and fairness, we compare the proposed focal frequency loss (FFL) with relevant loss functions that aim at improving image reconstruction and synthesis quality. Specifically, we select the widely used spatial-based method, *i.e.*, perceptual loss (PL) [30], which depends on high-level features from a pre-trained VGG [60] network. We also study the frequency-based approach, *i.e.*, spectral regularization (SpReg) [11], which is derived based on the azimuthal integration of the Fourier power spectrum. Besides, we further compare with another transformation form for FFL, *i.e.*, discrete cosine transform (DCT).

The comparison results are reported in Table 7. FFL outperforms the relevant approaches (*i.e.*, PL and SpReg) when applied to our baselines in different image reconstruction and synthesis tasks. It is noteworthy that FFL and PL are complementary, as shown by our previous experiments on SPADE, which also uses PL. Even if we replace DFT with DCT as the transformation form of FFL, the results are still better than previous methods. The performance is only slightly inferior to that obtained by FFL with DFT (*i.e.*, Eq. (10)). We deduce that the transformation form for FFL may be flexible. At this stage, DFT may be a better choice.

Ablation studies. We present ablation studies of each key component for FFL in Figure 11 and corresponding quantitative results in Table 8. For intuitiveness and simplicity, we use vanilla AE image reconstruction on CelebA for the evaluation.

The full FFL achieves the best performance. If we do not use the frequency representation of images (Section 3.1) and focus the model on hard pixels in the spatial domain, the synthesized images become more blurry. The quantitative results degrade. Discarding either the phase or amplitude information (Section 3.2) harms the metric performance vastly. Visually, using no phase information (amplitude only), the contour of reconstructed faces is retained, but the color is completely shifted. Without amplitude (phase only), the model cannot reconstruct the faces at all, and the full identity information is lost. This further verifies the necessity of considering both amplitude and phase information. Without focusing the model on the hard frequen-




Figure 11. **Ablation studies** of each key component for the focal frequency loss (FFL), *i.e.*, frequency representation (freq), phase and amplitude (ampli) information, and dynamic spectrum weighting (focal) in the vanilla AE image reconstruction task on CelebA.

Table 8. The PSNR (higher is better), SSIM (higher is better), LPIPS (lower is better), FID (lower is better) and LFD (lower is better) scores for the **ablation studies** of each key component for the focal frequency loss (FFL).

	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	LFD \downarrow
baseline	20.044	0.568	0.237	97.035	14.785
full FFL	21.703	0.642	0.199	83.801	14.403
w/o freq	18.200	0.470	0.265	123.833	15.210
w/o phase	13.273	0.380	0.407	233.170	16.344
w/o ampli	15.640	0.359	0.539	323.528	15.799
w/o focal	20.163	0.574	0.234	94.497	14.758

cies by the dynamic spectrum weighting (*i.e.*, directly using Eq. (8)), the results are visually similar to baseline, in line with our discussion in Section 3.3. The metrics decrease, being close to but slightly better than baseline, which may benefit from the frequency representation.

5. Conclusion

The proposed focal frequency loss directly optimizes image reconstruction and synthesis methods in the frequency domain. The loss adaptively focuses the model on the frequency components that are hard to deal with to ameliorate quality. The loss is complementary to existing spatial losses of diverse baselines varying in categories, network structures, and tasks, outperforming relevant approaches. We further show the potential of focal frequency loss to improve synthesis results of StyleGAN2. Exploring other applications and devising better frequency domain optimization strategies can be interesting future works.

Acknowledgments. This study is supported under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

- [1] DeepFaceLab. <https://github.com/iperov/DeepFaceLab/>. Accessed: 2019-08-20. **2**
- [2] DeepFakes. <https://github.com/deepfakes/faceswap/>. Accessed: 2019-08-16. **2**
- [3] Gwern Branwen, Anonymous, and Danbooru Community. Danbooru2019 Portraits: A large-scale anime head illustration dataset. <https://www.gwern.net/Crops#danbooru2019-portraits>. Accessed: 2021-04-10. **16**
- [4] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *ICLR*, 2018. **5**
- [5] Mu Cai, Hong Zhang, Huijuan Huang, Qichuan Geng, and Gao Huang. Frequency domain image translation: More photo-realistic, better identity-preserving. *arXiv preprint*, arXiv:2011.13611, 2020. **2**
- [6] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *ICCV*, 2017. **7**
- [7] Wenlin Chen, James Wilson, Stephen Tyree, Kilian Q Weinberger, and Yixin Chen. Compressing convolutional neural networks in the frequency domain. In *KDD*, 2016. **2**
- [8] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation. In *CVPR*, 2018. **2**
- [9] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, , and A. Vedaldi. Describing textures in the wild. In *CVPR*, 2014. **5, 13, 14**
- [10] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The Cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016. **5, 13, 14**
- [11] Ricard Durall, Margret Keuper, and Janis Keuper. Watch your up-convolution: CNN based generative deep neural networks are failing to reproduce spectral distributions. In *CVPR*, 2020. **2, 5, 8, 14**
- [12] Pedro F Felzenszwalb, Ross B Girshick, David McAllester, and Deva Ramanan. Object detection with discriminatively trained part-based models. *TPAMI*, 32:1627–1645, 2009. **2, 4**
- [13] Joel Frank, Thorsten Eisenhofer, Lea Schönherr, Asja Fischer, Dorothea Kolossa, and Thorsten Holz. Leveraging frequency analysis for deep fake image recognition. In *ICML*, 2020. **2**
- [14] Manuel Fritsche, Shuhang Gu, and Radu Timofte. Frequency separation for real-world super-resolution. In *ICCVW*, 2019. **2**
- [15] Rinon Gal, Dana Cohen, Amit Bermano, and Daniel Cohen-Or. SWAGAN: A style-based wavelet-driven generative model. *arXiv preprint*, arXiv:2102.06108, 2021. **2**
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014. **1, 2, 13, 14**
- [17] Lionel Gueguen, Alex Sergeev, Ben Kadlec, Rosanne Liu, and Jason Yosinski. Faster neural networks straight from JPEG. In *NeurIPS*, 2018. **2**
- [18] Seungwook Han, Akash Srivastava, Cole Hurwitz, Prasanna Sattigeri, and David D Cox. not-so-biggan: Generating high-fidelity images on a small compute budget. *arXiv preprint*, arXiv:2009.04433, 2020. **2**
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. **13**
- [20] 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 *NeurIPS*, 2017. **5, 13**
- [21] Geoffrey E Hinton and Ruslan R Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313:504–507, 2006. **1, 2, 5, 13**
- [22] Jiaxing Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. RDA: Robust domain adaptation via fourier adversarial attacking. In *ICCV*, 2021. **2**
- [23] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, 2018. **2**
- [24] Yihao Huang, Felix Juefei-Xu, Qing Guo, Xiaofei Xie, Lei Ma, Weikai Miao, Yang Liu, and Geguang Pu. FakeRetouch: Evading deepfakes detection via the guidance of deliberate noise. *arXiv preprint*, arXiv:2009.09213, 2020. **1, 2**
- [25] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. **13**
- [26] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. **1, 2, 5, 6, 13, 16**
- [27] Liming Jiang, Zhengkui Guo, Wayne Wu, Zhaoyang Liu, Ziwei Liu, Chen Change Loy, et al. DeeperForensics Challenge 2020 on real-world face forgery detection: Methods and results. *arXiv preprint*, arXiv:2102.09471, 2021. **2**
- [28] Liming Jiang, Ren Li, Wayne Wu, Chen Qian, and Chen Change Loy. DeeperForensics-1.0: A large-scale dataset for real-world face forgery detection. In *CVPR*, 2020. **2**
- [29] Liming Jiang, Changxu Zhang, Mingyang Huang, Chunxiao Liu, Jianping Shi, and Chen Change Loy. TSIT: A simple and versatile framework for image-to-image translation. In *ECCV*, 2020. **2, 5**
- [30] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. **2, 5, 8, 13, 14**
- [31] Steffen Jung and Margret Keuper. Spectral distribution aware image generation. *arXiv preprint*, arXiv:2012.03110, 2020. **2**
- [32] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. *arXiv preprint*, arXiv:1710.10196, 2017. **2, 5, 13, 14**
- [33] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. **2, 7, 14, 16, 18**

- [34] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020. 1, 2, 5, 7, 13, 14, 16, 18
- [35] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*, arXiv:1412.6980, 2014. 13, 14
- [36] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *NeurIPS*, 2018. 1
- [37] Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In *NeurIPS*, 2014. 2
- [38] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint*, arXiv:1312.6114, 2013. 1, 2, 5, 13
- [39] A Levinskis. Convolutional neural network feature reduction using wavelet transform. *Elektronika ir Elektrotechnika*, 19:61–64, 2013. 2
- [40] Jae Hyun Lim and Jong Chul Ye. Geometric GAN. *arXiv preprint*, arXiv:1705.02894, 2017. 13
- [41] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017. 2, 4
- [42] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NeurIPS*, 2017. 2
- [43] Xihui Liu, Guojun Yin, Jing Shao, Xiaogang Wang, and Hongsheng Li. Learning to predict layout-to-image conditional convolutions for semantic image synthesis. In *NeurIPS*, 2019. 2
- [44] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, 2015. 5, 13, 14, 15
- [45] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *ICML*, 2018. 14
- [46] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 1, 2, 5
- [47] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint*, arXiv:1411.1784, 2014. 2, 13
- [48] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint*, arXiv:1802.05957, 2018. 13
- [49] Augustus Odena, Vincent Dumoulin, and Chris Olah. Deconvolution and checkerboard artifacts. *Distill*, 1:e3, 2016. 1
- [50] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, 2019. 1, 2, 5, 7, 13
- [51] Stanislav Pidhorskyi, Donald A Adjeroh, and Gianfranco Doretto. Adversarial latent autoencoders. In *CVPR*, 2020. 1
- [52] Xiaojuan Qi, Qifeng Chen, Jiaya Jia, and Vladlen Koltun. Semi-parametric image synthesis. In *CVPR*, 2018. 7
- [53] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint*, arXiv:1511.06434, 2015. 2
- [54] Radim Šára Radim Tyleček. Spatial pattern templates for recognition of objects with regular structure. In *GCPR*, 2013. 5, 13, 14
- [55] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *ICML*, 2019. 1, 2, 12
- [56] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *NeurIPS*, 2008. 2
- [57] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 13
- [58] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training GANs. In *NeurIPS*, 2016. 5
- [59] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *CVPR*, 2016. 2, 4
- [60] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint*, arXiv:1409.1556, 2014. 8, 13
- [61] Nalini M Singh, Juan Eugenio Iglesias, Elfar Adalsteinsson, Adrian V Dalca, and Polina Golland. Joint frequency-and image-space learning for fourier imaging. *arXiv preprint*, arXiv:2007.01441, 2020. 2
- [62] Matthew Tancik, Pratul P Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *NeurIPS*, 2020. 1, 2
- [63] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint*, arXiv:1607.08022, 2016. 13
- [64] Arash Vahdat and Jan Kautz. NVAE: A deep hierarchical variational autoencoder. In *NeurIPS*, 2020. 1
- [65] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with PixelCNN decoders. In *NeurIPS*, 2016. 1
- [66] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 2
- [67] Haohan Wang, Xindi Wu, Zeyi Huang, and Eric P Xing. High-frequency component helps explain the generalization of convolutional neural networks. In *CVPR*, 2020. 2
- [68] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. CNN-generated images are surprisingly easy to spot...for now. In *CVPR*, 2020. 1, 2
- [69] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional GANs. In *CVPR*, 2018. 2, 7, 13
- [70] Yunhe Wang, Chang Xu, Shan You, Dacheng Tao, and Chao Xu. CNNpack: Packing convolutional neural networks in the frequency domain. In *NeurIPS*, 2016. 2

- [71] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 13:600–612, 2004. 5, 12
- [72] Yunxuan Wei, Shuhang Gu, Yawei Li, and Longcun Jin. Un-supervised real-world image super resolution via domain-distance aware training. *arXiv preprint*, arXiv:2004.01178, 2020. 2
- [73] Mingqing Xiao, Shuxin Zheng, Chang Liu, Yaolong Wang, Di He, Guolin Ke, Jiang Bian, Zhouchen Lin, and Tie-Yan Liu. Invertible image rescaling. In *ECCV*, 2020. 2
- [74] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *ECCV*, 2018. 5
- [75] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. In *ICCV*, 2015. 14
- [76] Kai Xu, Minghai Qin, Fei Sun, Yuhao Wang, Yen-Kuang Chen, and Fengbo Ren. Learning in the frequency domain. In *CVPR*, 2020. 2
- [77] Zhi-Qin John Xu, Yaoyu Zhang, Tao Luo, Yanyang Xiao, and Zheng Ma. Frequency principle: Fourier analysis sheds light on deep neural networks. *arXiv preprint*, arXiv:1901.06523, 2019. 1, 2, 12
- [78] Jaejun Yoo, Youngjung Uh, Sanghyuk Chun, Byeongkyu Kang, and Jung-Woo Ha. Photorealistic style transfer via wavelet transforms. In *ICCV*, 2019. 2
- [79] Aron Yu and Kristen Grauman. Fine-grained visual comparisons with local learning. In *CVPR*, 2014. 5, 13, 14
- [80] Fisher Yu, Vladlen Koltun, and Thomas Funkhouser. Dilated residual networks. In *CVPR*, 2017. 5
- [81] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. *arXiv preprint*, arXiv:1805.08318, 2018. 13
- [82] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 5
- [83] Xu Zhang, Svebor Karaman, and Shih-Fu Chang. Detecting and simulating artifacts in gan fake images. In *WIFS*, 2019. 1, 2
- [84] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ADE20K dataset. In *CVPR*, 2017. 5, 13, 14
- [85] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 2

Appendix

This appendix provides supplementary information that is not elaborated in our main paper: Section A shows some additional illustrations to explain our method further. Section B describes the implementation details in our experiments. Section C details our used datasets under diverse settings. Section D provides some studies on the variants of focal frequency loss. Section E presents additional results and analysis.

A. Additional Illustrations of Methodology

A.1. Spatial Frequency Visualization

After applying 2D discrete Fourier transform, an image is converted into its frequency representation and decomposed into orthogonal sine and cosine functions. The angular frequency of each sine and cosine function is decided by the frequency spectrum coordinate (u, v) . The spatial frequency manifests as the 2D sinusoidal components in the image. The spectrum coordinate also represents the angled direction of a specific spatial frequency. As an intuitive view, we show some examples of the 2D sinusoidal components with specific spatial frequencies in Figure 12. It is observed that the angled direction and density (angular frequency) of the waves depend on the spectrum coordinate (u, v) . Besides, the complex frequency value $F(u, v)$ can be regarded as the weight for each wave, and the weighted sum corresponds to the whole image in the spatial domain.

A.2. More Intuitive Illustration

To further explain the proposed focal frequency loss (FFL), we will provide a more intuitive illustration in this section. According to Figure 12, an image (gray-scale for simplicity) is the weighted sum of different spatial frequencies. We expand the accumulated frequencies into a new dimension, thus the image can be seen as a cube in a space. The length (L) and width (W) dimensions of the cube correspond to the pixel domain, and the height (H) dimension corresponds to the frequency domain, as shown in Figure 13. Therefore, a single pixel can be seen as the orange prism, and a specific frequency can be regarded as the green plane. It is observed that each frequency (*i.e.*, each coordinate value on the frequency spectrum) depends on all the image pixels. Due to the inherent bias of neural networks [55, 77], a model tends to eschew some frequency components that are hard to synthesize, *i.e.*, hard frequencies, in the H dimension. Optimizing in the spatial domain (*i.e.*, in the L and W dimensions) hardly help the model locate these hard frequencies in the H dimension. Similarly, focusing on certain pixels (*e.g.*, orange prism) hardly help the model tackle the hard frequencies (*e.g.*, green plane). Intuitively, when directly optimizing in the H dimension (*i.e.*, explicitly using the frequency representation of the image in




Figure 12. Two-dimensional sinusoidal components with specific spatial frequencies in an image. The angled direction and density (angular frequency) of the waves depend on the spectrum coordinate (u, v) , and $F(u, v)$ can be seen as the weight for each wave.




Figure 13. According to Figure 12, an image (gray-scale for simplicity) can be seen as a cube in a space, where its length (L) and width (W) dimensions correspond to the pixel domain, and the height (H) dimension corresponds to the frequency domain.

our method), the model can easily locate hard frequencies and in turn focus on them.

In Figure 13, it is noteworthy that each frequency also affects all the image pixels in the spatial domain. When FFL directly optimizes and adaptively focuses a model in the frequency domain, the frequency components in the H dimension will be reconstructed and synthesized better. Meanwhile, the general alignment and quality of all the image pixels in the L and W dimensions will be indirectly improved by FFL, thus boosting some pixel-based metrics (*e.g.*, PSNR and SSIM [71]) and ameliorating the image reconstruction and synthesis quality.

We wish to highlight that both the spatial-based loss and frequency-based loss are important since they consider different aspects and dimensions of an image, as illustrated in Figure 13. Hence, they are complementary and not replaceable. The proposed FFL is intending to complement existing spatial losses of different methods to improve reconstruction and synthesis quality further.

In fact, the actual situation of the frequency components in an image is much more complicated, which may be a higher-dimensional representation. The visualization in this section just provides a simple and intuitive illustration to help understand the proposed method in this paper.

B. Implementation Details

The code used for our experiments will be made publicly available. All the experiments are conducted on the NVIDIA Tesla V100 GPUs with 32 GB memory capacity.

B.1. Baseline Details

In this section, we will provide the implementation details of all the baselines in different image reconstruction and synthesis tasks. We select five representative methods from the two popular categories: autoencoder-based and GAN-based. Besides, we evaluate different network structures. Specifically, we explore the multilayer perceptron (MLP) network and the convolutional neural network (CNN). For CNN, the network details also vary, *e.g.*, with or without the skip connections. In addition, we consider various basic spatial domain losses, *e.g.*, MSE loss, L1 loss, GAN loss [16], perceptual loss [30], *etc.*, to test the ability of focal frequency loss to complement these losses.

Vanilla AE. Vanilla autoencoder [21] learns the image latent representation in an unsupervised manner, traditionally used for dimension reduction and feature learning. We employ vanilla AE in the image reconstruction task. The network is a simple 2-layer MLP with a hidden size of 256. We adopt ReLU activations (except the last layer using Tanh) and no norm layers. We use Adam [35] optimizer and set $\beta_1 = 0.9, \beta_2 = 0.999$. The learning rate is 0.001. Normal initialization (with mean 0.0 and standard deviation 0.02) is applied to all the networks of vanilla AE. The spatial loss is MSE loss. The models are trained on 1 GPU with a batch size of 128. We perform 200 epochs of training on DTD [9] and 20 epochs of training on CelebA [44].

VAE. Exploiting a reparameterization trick, the variational autoencoder [38] generates images by learning the latent representation in a probability distribution manner. We use VAE for image reconstruction and unconditional image synthesis. We employ CNN for VAE, with typical convolution and transposed convolution layers. Batch normalization [25] and Leaky ReLU (with a negative slope of 0.2, except the last layer using Tanh) are applied. Each convolution layer has a kernel size 4×4 , stride 2, and zero-padding amount 1. In the encoder, the feature map resolution is halved after each convolution block. Images are down-sampled to 4×4 , so the number of blocks depends on the input size (*e.g.*, if the input size is 64×64 , there will be 4 blocks). After an input layer, the number of feature channels is 64. Then, the number of feature channels will double after each convolution block, while we set a maximum channel number to 512 to avoid using redundant parameters. We apply two linear layers to the encoded feature to learn μ and σ for the reparameterization. The latent size is 256. After reparameterization, an additional linear layer is used to adjust the feature to the original shape. In the decoder, the network structure is completely inverse to the

encoder by replacing convolution layers with the transposed convolution layers. We use Adam [35] optimizer and set $\beta_1 = 0.9, \beta_2 = 0.999$. The learning rate is 0.001. Normal initialization (with mean 0.0 and standard deviation 0.02) is applied to all the networks of VAE. The spatial losses are MSE loss and KL divergence loss [38]. The models are trained on 1 GPU with a batch size of 128. We train our models for 20 epochs on CelebA [44] and 400 epochs on CelebA-HQ [32].

pix2pix. pix2pix [26] adopts conditional GAN [47] as a general-purpose solution to image-to-image translation with training pairs. We employ pix2pix for conditional image synthesis. The U-Net [57] generator is applied, which is an encoder-decoder with skip connections between mirrored layers in the encoder and decoder stacks. There are 8 skip connection blocks in the generator. The patch-based discriminator is used. Adam [35] optimizer is used with $\beta_1 = 0.5, \beta_2 = 0.999$. The learning rate is 0.0002. Normal initialization (with mean 0.0 and standard deviation 0.02) is applied to all the networks of pix2pix. The spatial losses are vanilla GAN loss [16] and L1 loss. The models are trained on 1 GPU. We conduct 200 epochs of training on CMP Facades [54] with a batch size of 1. We train the models for 15 epochs on edges \rightarrow shoes [79] with a batch size of 4. For other detailed network structures and parameters, we follow the original paper [26] and their released code.

SPADE. As a task-specific GAN-based method for semantic image synthesis (*i.e.*, synthesizing a photorealistic image from a semantic segmentation mask), SPADE [50] resizes the segmentation mask for modulating the activations in normalization layers by a learned affine transformation. The generator is built on a series of residual blocks [19] with the synchronized version of batch normalization. The multi-scale patch-based discriminator [69] with the instance normalization [63] is exploited. Besides, spectral normalization [48] is applied to all the convolutional layers in the generator and discriminator. Adam [35] optimizer is exploited with $\beta_1 = 0, \beta_2 = 0.9$. Two time-scale update rule [20] is applied, where the learning rates for the generator and the discriminator are 0.0001 and 0.0004, respectively. The spatial losses are hinge-based GAN loss [40, 48, 81], perceptual loss [30] calculated by VGG-19 [60] model, and feature matching loss [69]. The models are trained for 200 epochs on Cityscapes [10] and ADE20K [84] using 4 GPUs. The batch size is 32. For other detailed network structures and parameters, we follow the original paper [50] and their released code.

StyleGAN2. We further explore the potential of focal frequency loss on the state-of-the-art unconditional image synthesis method, StyleGAN2 [34]. We construct the StyleGAN2 baseline on top of its open-source official implementation. The mapping network consists of 8 fully connected layers. The dimensionality of both the input latent space and intermediate latent space is 512. The activation func-

tion is Leaky ReLU with a negative slope of 0.2. Several standard techniques in [32, 33] are applied, such as the exponential moving average of generator weights, mini-batch standard deviation layer at the end of the discriminator, equalized learning rate for all the trainable parameters, *etc.* Adam [35] optimizer is used with $\beta_1 = 0, \beta_2 = 0.99$. The spatial loss is non-saturating logistic loss [16, 34] with R_1 regularization [45]. All the models are trained with 8 V100 GPUs. The batch size is 64 for CelebA-HQ [32] (256×256) and 32 for the resolution of 1024×1024 . For other detailed network structures and parameters, we follow the original paper [34] and their released code.

As for the relevant losses used for comparison, *i.e.*, perceptual loss [30] and spectral regularization [11], we follow all the details in their papers and released code.

B.2. Computational Cost

The computational cost of the proposed focal frequency loss (FFL) is negligible. Take pix2pix image-to-image translation on the CMP Facades dataset as an example. The average computational training time only increases from 0.064 to 0.067 seconds per iteration after applying FFL. The memory consumption increases from 3513 to 3515 MB. This cost test is conducted on 1 NVIDIA Tesla V100 GPU.

C. Dataset Details

In this section, we will provide detailed information about the seven datasets we explored. The datasets vary in types, sizes, and resolutions.

- **Describable Textures Dataset (DTD).** We use DTD provided by [9], which is an evolving collection of textural images in the wild, annotated with human-centric attributes. DTD contains texture images with special frequency patterns. We perform vanilla AE image reconstruction using this dataset, with 4,512 images for training and 1,128 images for testing. The original images are scaled and center cropped to 64×64 .
- **CelebA.** CelebA [44] is a large-scale face attributes dataset covering large pose variations and background clutter. We conduct image reconstruction with vanilla AE and VAE on CelebA. Besides, we perform VAE unconditional image synthesis on CelebA. We use the cropped and aligned faces, which are more natural images. The training set contains 199,599 images, and the test set has 3,000 images. The images are resized and center cropped to 64×64 .
- **CelebA-HQ.** CelebA-HQ is a higher-quality version of the CelebA dataset provided by [32]. The original resolution is 1024×1024 . We perform VAE image reconstruction on this dataset. Besides, we study the unconditional image synthesis by VAE and StyleGAN2 using CelebA-HQ. The dataset is randomly

split, yielding 27,000 images for training and 3,000 images for evaluation. All the cropped and aligned face images are uniformly resized to 256×256 . For StyleGAN2, we also tried to synthesize images with a resolution of 1024×1024 besides 256×256 .

- **CMP Facades.** For pix2pix image-to-image translation, we utilize the officially prepared CMP Facades [54] dataset. The facades are collected from different cities around the world with diverse architectural styles. CMP Facades contains architectural labels and photos, which is suitable for mask \rightarrow image translation. The sizes of training and test sets are 400 and 106, respectively. The resolution is 256×256 .
- **Edges \rightarrow shoes.** We also exploit the officially prepared edges \rightarrow shoes dataset for pix2pix image-to-image translation. The shoe images are from UT Zappos50K [79]. The shoes are centered on a white background. The edge maps are detected by HED [75]. The numbers of images for training and testing are 49,825 and 200, respectively. The image size is 256×256 .
- **Cityscapes.** We use the Cityscapes [10] dataset for SPADE semantic image synthesis. Cityscapes dataset consists of street scene images that are mostly collected in Germany. The dataset provides instance-wise, dense pixel annotations of 30 classes. The training set has 2,975 images, and the test set contains 500 images. The images are scaled to 512×256 .
- **ADE20K.** ADE20K [84] dataset contains challenging in-the-wild images with fine annotations of 150 semantic classes. We also use ADE20K for SPADE semantic image synthesis, with 20,210 images for training and 2,000 images for evaluation. All the images are resized to 256×256 .

D. Variant Studies

In our main paper, we mentioned that the exact form of the proposed focal frequency loss (FFL) is not crucial. In this section, we will provide some variants to extend and modify FFL. We will show some studies on these variants. For simplicity and intuitiveness, we revisit the vanilla AE image reconstruction task on CelebA. We report quantitative evaluation results for the variant studies. The visual results of variants are similar.

Several simple variants can be derived by adjusting the spectrum weight matrix parameter α . The parameter α controls how close the weight matrix values are, *i.e.*, how focused the model is. The larger α is, the model will be more focused on the hard frequencies, *i.e.*, the weight difference for easy and hard frequencies will be larger. For the experiments we present in our main paper, we set $\alpha = 1$ (we call the main version). The results are shown in Table 9.

Table 9. The PSNR (higher is better), SSIM (higher is better), LPIPS (lower is better), FID (lower is better) and LFD (lower is better) scores for the **variant studies** on the spectrum weight matrix **parameter** α for the focal frequency loss.

	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	LFD \downarrow
baseline	20.044	0.568	0.237	97.035	14.785
$\alpha = 1$ (main)	21.703	0.642	0.199	83.801	14.403
$\alpha = 2$	21.376	0.621	0.203	102.329	14.478
$\alpha = 0.5$	21.521	0.635	0.197	82.561	14.445
$\alpha = 0.1$	20.497	0.591	0.225	89.792	14.681

Table 10. The PSNR (higher is better), SSIM (higher is better), LPIPS (lower is better), FID (lower is better) and LFD (lower is better) scores for the **variant studies** on **patch-based** focal frequency loss. Patch factor p is the number of patches on each edge.

	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	LFD \downarrow
baseline	20.044	0.568	0.237	97.035	14.785
$p = 1$ (main)	21.703	0.642	0.199	83.801	14.403
$p = 2$	21.836	0.648	0.185	88.475	14.372
$p = 4$	21.752	0.643	0.170	90.612	14.392
$p = 8$	21.414	0.627	0.176	102.334	14.470

Applying the main version of FFL ($\alpha = 1$) shows better performance than the baseline without FFL in all the five metrics. If we set $\alpha = 2$, the quantitative results degrade from the main version, especially FID. This suggests that the model may be too focused on the hard frequencies while ignoring some important easy frequency information, albeit the results are still better than the baseline in most cases. When setting $\alpha = 0.5$, all the metric results are better than the baseline. The LPIPS and FID scores become better than the main version. The results of this variant are close to the main version of FFL. If we set $\alpha = 0.1$, the quantitative results degrade from the main version despite still better than the baseline. This indicates that the model may be too unfocused. For a trade-off, we select $\alpha = 1$ as the main version of FFL, while one may consider choosing other variants regarding the parameter α in certain tasks for the flexibility.

Besides, we study another category of variants, the patch-based focal frequency loss, where we crop an image into small patches so that the focused frequencies are at the patch level. We define the patch factor p as the number of patches on each edge. For instance, if $p = 2$, the image will be cropped into $2 \times 2 = 4$ patches. Obviously, using the original image without cropping it into patches, *i.e.*, the main version of FFL we defined before, corresponds to $p = 1$. The results are shown in Table 10. We note that $p = 1, 2, 4$ achieve close performance regarding the five evaluation metrics, all of which are much better than the baseline. However, if we set $p = 8$, the quantitative performance will degrade from the previous versions, especially FID. Although the results are still better than the baseline in most cases, this indicates that the patch size should not be too small. We simply choose $p = 1$ as the main version of FFL for our experiments in the main paper. However, the variant studies show that the patch-based focal frequency loss may contribute to an additional performance boost in




Figure 14. The spatial losses (MSE) with the same weight and random seed of the two training processes with/without focal frequency loss (FFL) for vanilla AE image reconstruction on CelebA. The spatial loss converges to a lower point with the help of FFL.

certain cases. Thus, this may be another direction to extend and modify FFL.

E. Additional Results and Analysis

E.1. Training Loss

In the main paper, we have mentioned that the proposed focal frequency loss (FFL) is complementary to existing spatial losses, *e.g.*, MSE loss, to improve image reconstruction and synthesis quality. We further analyze the training loss in this section. We choose the vanilla AE image reconstruction task on CelebA [44] for simplicity. We plot the spatial losses with the same weight and random seed of the two training processes with/without FFL in Figure 14. It is readily observed that the spatial loss (MSE) converges to a lower point after applying FFL. This indicates that the model may converge to a better point with the help of FFL, in line with the better perceptual quality and quantitative performance we presented in our main paper.

E.2. Frequency Domain Gap

As mentioned in the main paper, we wish to improve the image reconstruction and synthesis quality by narrowing the frequency domain gap between the real and generated images using the proposed focal frequency loss (FFL). We have shown that the gaps between mini-batch average spectra of state-of-the-art StyleGAN2 are clearly mitigated by FFL. We will show some more examples of VAE image reconstruction on the CelebA [44] dataset and provide more analysis about the frequency domain gap in this section.

The results are shown in Figure 15. In the spatial domain, without applying FFL, the reconstructed faces are blurry. This may be attributed to the reparameterization operation in the latent space between the encoder and decoder, which increases the difficulty for reconstruction. Trained with FFL, the VAE model can synthesize much clearer results, being closer to the ground truth real images. The perceptual quality is better after applying FFL. In the fre-




Figure 15. Frequency domain gaps are narrowed by the focal frequency loss (FFL) for VAE image reconstruction on CelebA.

edge	baseline	full FFL	w/o freq	w/o phase	w/o ampli	w/o focal
FID ↓	80.279	74.359	86.674	98.778	89.255	77.864
IS ↑	2.674	2.804	2.713	2.667	2.527	2.705

Figure 16. **Additional ablation studies** of each key component for the focal frequency loss (FFL), *i.e.*, frequency representation (freq), phase and amplitude (ampli) information, and dynamic spectrum weighting (focal) in the pix2pix image-to-image translation task on edges \rightarrow shoes (256×256). The corresponding FID (lower is better) and IS (higher is better) scores are reported below the images.

quency domain, in line with our visualizations in the main paper, the VAE baseline without FFL bias to a limited spectrum region, losing high-frequency information (outer regions and corners). The frequency domain gaps are clearly narrowed after adopting FFL. The spectrum distribution becomes closer to the ground truth. Besides, some essential special spectrum patterns can be generated by applying FFL. This suggests the effectiveness of focal frequency loss to narrow the frequency domain gaps and ameliorate image quality further.

E.3. Additional Ablation Studies

In the main paper, we provided the ablation studies of vanilla AE image reconstruction on CelebA for intuitiveness and simplicity, intending to study the importance of each key component for the proposed focal frequency loss (FFL) while reducing the influence of other factors, such as the adversarial loss. In this section, we provide the additional ablation studies on higher-resolution images with GAN. We show the studies of pix2pix [26] (*i.e.*, GAN-based method) image-to-image translation on edges \rightarrow shoes (256×256) in Figure 16. The results are in line with the ablation studies in our main paper, further suggesting the importance of each key component for FFL.

E.4. Results on Non-Photorealistic Images

We further study the benefit of the proposed focal frequency loss (FFL) on non-photorealistic images. As an ex-

Table 11. The PSNR (higher is better), SSIM (higher is better), LPIPS (lower is better), FID (lower is better) and LFD (lower is better) scores for the **vanilla AE image reconstruction on Danbooru2019 Portraits (Anime)** trained with/without the focal frequency loss (FFL).

Dataset	FFL	PSNR↑	SSIM↑	LPIPS↓	FID↓	LFD↓
Anime	w/o	19.885	0.575	0.294	193.342	14.822
(64×64)	w/	20.657	0.628	0.267	184.443	14.644

ample, we provide the vanilla AE image reconstruction results on Danbooru2019 Portraits [3] (Anime) in Table 11. Empirically, we observe that all the metrics can still be boosted by FFL. Our intuition is that FFL can also help generate non-photorealistic images since they still possess special frequency patterns that may be hard for a network to learn. FFL is adaptive for dealing with these frequencies.

E.5. Higher-Resolution Results on StyleGAN2

In Figure 17, we show some higher-resolution images synthesized by StyleGAN2 [34] trained with or without the proposed focal frequency loss (FFL) on CelebA-HQ (1024×1024). The truncation trick [33, 34] is not applied. Although the original StyleGAN2 (w/o FFL) generates plausible images in most cases, it sometimes produces tiny artifacts on the face (Row 2) and eyes (Row 3). The details on the teeth are missing in certain cases (Row 1). The synthesized images by StyleGAN2 with FFL (w/ FFL) are very photorealistic. Besides, StyleGAN2 achieves a better FID score after applying FFL, indicating that the quality of generated images becomes better with the help of FFL. More random sampled synthesized images without truncation are shown in Figure 18, and the examples with truncation using $\psi = 0.5$ [33, 34] are presented in Figure 19. It is observed that all the images generated by StyleGAN2 with FFL are with very high fidelity.

w/o FFL

w/ FFL



Figure 17. Synthesis results (without truncation) of StyleGAN2 trained with/without the proposed FFL on CelebA-HQ (1024×1024). The model with FFL achieves the FID score of **3.374**, outperforming the original StyleGAN2 without FFL of 3.733.



Figure 18. More random sampled images (without truncation) synthesized by StyleGAN2 trained with the proposed FFL on CelebA-HQ (1024×1024).



Figure 19. More random sampled images (with truncation applied using $\psi = 0.5$ [33, 34]) synthesized by StyleGAN2 trained with the proposed FFL on CelebA-HQ (1024×1024).