

WaveGAN: Frequency-aware GAN for High-Fidelity Few-shot Image Generation

Mengping Yang^{1,2}, Zhe Wang^{1,2, (✉)}, Ziqiu Chi^{1,2}, and Wenyi Feng^{1,2}

¹ Department of Computer Science and Engineering, East China University of Science and Technology, China

² Key Laboratory of Smart Manufacturing in Energy Chemical Process, East China University of Science and Technology, China

mengpingyang@mail.ecust.edu.cn wangzhe@ecust.edu.cn
{chiziqiu Y10200096}@mail.ecust.edu.cn

Abstract. Existing few-shot image generation approaches typically employ fusion-based strategies, either on the image or the feature level, to produce new images. However, previous approaches struggle to synthesize high-frequency signals with fine details, deteriorating the synthesis quality. To address this, we propose WaveGAN, a frequency-aware model for few-shot image generation. Concretely, we disentangle encoded features into multiple frequency components and perform low-frequency skip connections to preserve outline and structural information. Then we alleviate the generator’s struggles of synthesizing fine details by employing high-frequency skip connections, thus providing informative frequency information to the generator. Moreover, we utilize a frequency L_1 -loss on the generated and real images to further impede frequency information loss. Extensive experiments demonstrate the effectiveness and advancement of our method on three datasets. Noticeably, we achieve new state-of-the-art with FID 42.17, LPIPS 0.3868, FID 30.35, LPIPS 0.5076, and FID 4.96, LPIPS 0.3822 respectively on Flower, Animal Faces, and VGGFace. GitHub: https://github.com/kobeshegu/ECCV2022_WaveGAN

Keywords: GANs, Few-shot Learning, Image Generation, Wavelet Transformation

1 Introduction

Recent years have witnessed remarkable developments in visual generative tasks with the rapid progress of generative models, especially Generative Adversarial Networks (GANs) [19] [22] [23] [21]. Despite being applied to various domains, the success of GANs mainly comes from immense data, and GANs struggle to synthesize high-quality images given insufficient data. Few-shot learning [7], being proposed to improve the generalization ability in the limited-data scenarios, has gained extensive attention and focused research. However, most of existing few-shot algorithms are designed for classification [27] and segmentation [25]

problems, few studies address few-shot image generation. Therefore, exploring and facilitating the generation quality in the few-shot regime is necessary.

Few-shot image generation aims at generating novel images for a category when given a few images from the same category. Specifically, the model is first trained on an auxiliary dataset (seen classes) with sufficient data in an episode-based manner [37], *i.e.*, feeding a specific number of images (*e.g.*, 3, 5) into the model in each episode. The trained model is then expected to produce new images when given a few images of a category from a new dataset (unseen classes). There is no overlap between the auxiliary dataset and the new testing dataset. The model is encouraged to capture the transferable ability from seen classes to unseen classes to generate new images for unseen classes.

Previous methods try to 1) transform intra-class information [1], 2) design new optimization schemes by combining GANs with meta-learning [4] [26], and 3) fuse the given images [14] [15] [12] to address few-shot image generation. Among these methods, LoFGAN [12] achieves current state-of-the-art performance by fusing local representations based on the semantic similarity of features. However, existing approaches ignore the enormous impact of the frequency information throughout the generating process. F-principal [39] proves that neural networks preferentially fit frequency signals from low to high. Consequently, the model tends to generate frequencies with higher priority and more superficial complexity, *i.e.*, only generating low-frequency signals.

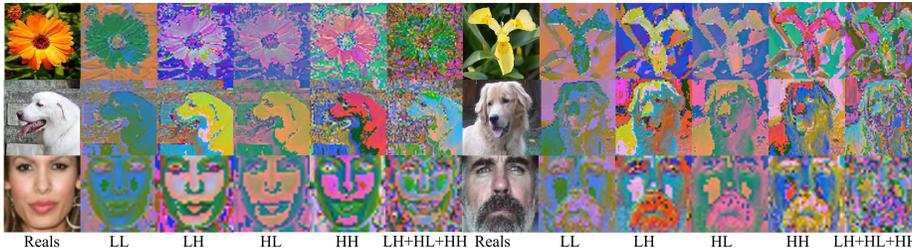


Fig. 1. Visualization of the transformed frequency components. *LL* denotes low-frequency component, and *LH*, *HL*, *HH* denote high-frequency components.

We visualize different frequency components of real images in Fig. 1. The low-frequency component (*i.e.*, *LL*) contains general information like the overall surface, outline, and structure. While rich details and perceptible information like the leaves of flowers, the tongue of dogs, and the hair of human, lie in the high-frequency components (*i.e.*, *LH*, *HL*, *HH*). Rich details can be obtained by adding all high-frequency components (*i.e.*, *LH+HL+HH*) together. Since high frequency components contain meticulous information, losing them may lead the generator to synthesize blurry images with more aliasing artifacts. This issue highlights the necessity of considering frequency signals in generating images, especially high-frequency signals as the generator usually eschews them [18] [39].

This paper proposes WaveGAN, an innovative and effective approach to ameliorate the few-shot synthesis quality from the perspective of frequency domains. We first perform wavelet decomposition to transform the encoded features from the spatial domain to multiple frequency domains, comprising low and high-frequency components. Then we feed the low-frequency component to the behind layers of the encoder via low-frequency skip connections, maintaining the overall outline and structural patterns. To mitigate the pressure of the generator to generate high-frequency signals and provide more details to the decoder, we directly feed the decomposed high-frequency signals to the decoder. Two strategies are designed to aggregate the high-frequency signals, namely WaveGAN-M and WaveGAN-B. Both of them are effective and can provide high-frequency information to the decoder. The high-frequency components are then precisely reconstructed back into the original features with our inverse frequency transformation operations, guaranteeing the minimal loss of high-frequency signals. In addition, we apply frequency L_1 -loss to the generated images and the real images, which is complementary to spatial losses and impedes losing frequency information. Our primary contributions can be summarized as follows:

- We propose WaveGAN, the first few-shot image generation method that exploits frequency components to promote synthesis quality. Adding low and high-frequency skip connections to the generator, our WaveGAN alleviates the generator’s struggles to encode high-frequency signals and provide more perceptible information, resulting in favorable generation quality.
- We design two techniques to aggregate the high-frequency information for reconstructing frequency signals back to the original features, *i.e.*, WaveGAN-M and WaveGAN-B, which preserve fine details and statistical properties. We also present frequency L_1 -loss to avoid losing frequency information.
- We conduct comprehensive experiments on three datasets. Both qualitative and quantitative results demonstrate the superiority and effectiveness of our method. Notably, our model outperforms the state-of-the-art approach with significant FID improvements (*e.g.*, from **102.07** to **30.35** on Animal Face).

2 Related Work

Generative Adversarial Network. Generative Adversarial Networks (GANs) have made significant progress since the pioneering work in [10]. Benefit from remarkable ability of capturing the data distribution, GANs have been successfully applied in various visual domains, including image generation [23] [21], video generation [38], image-to-image translation [32] [35], etc. Typically, a GAN model consists of a generator and a discriminator, and the two networks are updated alternatively in an adversarial manner. Training a GAN is notoriously formidable as it requires massive data and computation resources, and the adversarial training may make the model diverge. The discriminator is tend to overfitting when given limited data, resulting in poor generation quality. Several works have been proposed to mitigate the discriminator overfitting. Different data augmentation techniques, including differentiable [43], non-leaking [20] and adaptive pseudo

augmentation [17] are designed to expand the limited training data. Lecam [36] regularizes the output of the discriminator to avoid overfitting. Unlike these efforts made for unconditional image generation with limited data, in this paper, we seek to generate novel images for one specific category when given a few images from this category.

Wavelet Transformation in GANs. Decomposing given signals into different frequency components, wavelet transformation has made great success in various generative tasks such as style transfer [40], image reconstruction [18], image inpainting [41], image editing [9] and image super-resolution [6] [16]. These approaches try to narrow the information gaps in the frequency domain to boost the model’s performance. For example, Jiang *et al.* propose focal frequency loss to avoid the loss of important frequency information for image reconstruction tasks [18]. WaveFill [41] decomposes images into multiple frequency components and fills the corrupted image regions with decomposed signals, which achieves superior image inpainting. Different from these methods, we try to generate realistic and plausible images when given only a few data. We are interested in the influence of frequency information on the challenging few-shot image generation.

Few-shot Image Generation. Inspired by the human’s great generalization ability from a few observations, few-shot image generation models try to generate new images given a few images. Existing few-shot image generation approaches can be roughly divided into three categories: 1) Optimization-based, 2) Fusion-based, and 3) Transformation-base methods. DAGAN [1] transforms combined projected latent codes and encoded images to new images. The optimized-based methods FIGR [4] and DAWSON [26] combine generative models with optimization-based meta learning Reptile [29] and MAML [8], respectively. The fusion-based methods fuse the local feature [12] or the input images [15] [14] to synthesis novel images. GMN [2] combines VAE [24] with Matching Networks [37] to capture the few-shot distribution. MatchingGAN [14] matches random vectors with given real images and mapping the fused features to novel images. F2GAN [15] further improves MatchingGAN with a fusing-and-filling paradigm. By fusing local representations with semantic similarity, LoFGAN [12] promotes the generation quality. Notably, zero-shot or few-shot text-to-image generation methods [34] [11] [33] have made great progress recently. Differently, this paper focuses on the problem of few-shot image generation for generating new images for a given class as defined in Sec. 3.1.

However, existing methods ignore the influence of frequency components on the quality of generated images, leading the generator to synthesize unfavorable images with more artifacts and fewer details. In this paper, we present a frequency-aware model that can generate appealing and photorealistic images by adding low and high-frequency skip connections to the generator. Such design mitigates the generator’s pressure of synthesising high-frequency signals. Our work explores an effective solution for few-shot image generation from the frequency domain perspective, which complements previous fusion-based methods.

3 Methodology

3.1 Overview

Problem Definition. Given K images from a new class, our model’s goal is to synthesis diverse and plausible images for the given class. The number of images K defines a K -shot image generation task. Generally, this task is accomplished in two phases, *i.e.*, training and testing. The datasets is first split into seen classes \mathbb{C}_s and unseen classes \mathbb{C}_u , where \mathbb{C}_s and \mathbb{C}_u have no overlap. In the training phase, a substantial amount of K -shot image generation tasks sampled from \mathbb{C}_s are fed into the model, expecting the model to transfer the knowledge of generating new images learned from \mathbb{C}_s to \mathbb{C}_u . In the testing phase, the model takes images from \mathbb{C}_u as input to synthesis new images.

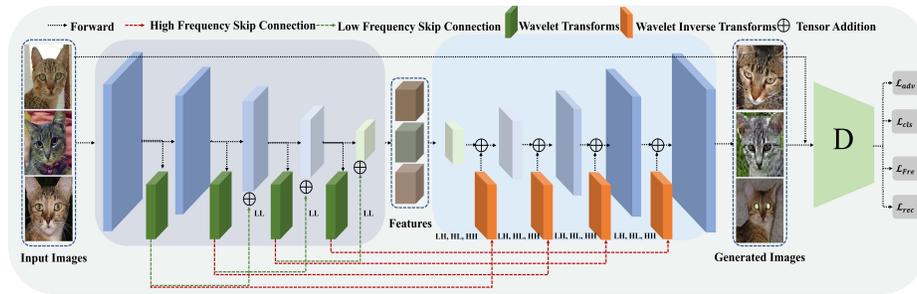


Fig. 2. The overall framework of our WaveGAN. We employ low frequency (LL) skip connections in the encoder and high frequency skip connection (LH, HL, HH) in the decoder to provide rich details to improve synthesis quality.

Overall Framework. As shown in Fig. 2, our model consists of a WaveEncoder, a WaveDecoder, and a Discriminator, the WaveEncoder and Wavedecoder constitute our generator. The WaveEncoder extracts feature representations of images, while the WaveDecoder maps the feature representation to new images. We perform wavelet transformation to the encoded features and obtain multiple frequency components. Then we employ low-frequency skip connections in the encoder to preserve the overall structure and outline. We exploit high-frequency skip connections to provide detailed information to the decoder. The wavelet inverse transformation module reconstructs these high-frequency signals to the original features. The high-frequency signals contain rich details and perceptible information, enabling the generator to synthesis high-quality images. The real and generated images are then fed into the discriminator to train the whole model. Next, we elaborate our WaveEncoder and WaveDecoder in detail.

3.2 WaveEncoder

Our WaveEncoder is composed of convolutional blocks and wavelet transformation blocks. The convolutional operations extract features for the decoder to

produce new images. To disentangle the extracted features into multiple frequency components, we adopt a simple yet effective wavelet transformation, *i.e.*, Haar wavelet [5]. Haar wavelet contains two operations: wavelet transform and inverse wavelet transformation, and four kernels, namely LL^T , LH^T , HL^T , and HH^T .

$$L^T = \frac{1}{\sqrt{2}} [1 \ 1], \quad H^T = \frac{1}{\sqrt{2}} [-1 \ 1] \quad (1)$$

where L and H denote the low and high pass filters, respectively. The low pass filter focuses on low-frequency signals containing the outline and structural information. In contrast, the high pass filter emphasizes high-frequency signals that capture fine-grained details like subtle edges and contour (See Fig. 1 and 5).

The wavelet transformation decomposes features into frequency components LL , LH , HL , HH . Among these frequency signals, LL captures the overall appearance and basic object structures of images (See Fig. 1). Thus we employ low-frequency skip connections in the encoder to obtain precise and faithful features throughout the feature extracting process. Specifically, for feature E_i obtained from the i -th convolutional block in the encoder, we adopt Haar wavelet transformation to extract the frequency components LL_i , LH_i , HL_i , HH_i . We then perform tensor addition on the low frequency signal LL_i and feature E_{i+1} obtained from the $(i + 1)$ -th convolutional block in the encoder.

$$E_{i+2} = LL_i + ConvBlock_{i+1}(E_{i+1}) \quad (2)$$

The obtained skip connected feature E_{i+2} is fed into the $(i + 2)$ -th convolutional block. The low-frequency skip connections contribute to the fidelity of the generated images. The experimental proofs are given in Sec. 4.4.

3.3 WaveDecoder

High-frequency components contain rich details of images. However, deep networks usually fit frequency signals from low to high, making it difficult for the generator to produce high-frequency information since it generates frequencies with higher priority. To alleviate the encoder’s pressure to synthesis rich details and provide fine-grained information to the decoder, we directly feed the decomposed high-frequency signals LH , HL , and HH into the WaveDecoder via high-frequency skip connections. Specifically, for the i -th layer of the encoder, we perform wavelet transformation on the features and obtain high frequency components LH_i , HL_i , HH_i , then we feed the inversed components to the $(n - i)$ -th layer of the decoder as exhibited in Fig. 2, where n is the number of all layers. This operation encourages the decoder to synthesis images with more details and fewer artifacts. We employ wavelet inverse transformation to reconstruct high-frequency signals back to original features. Our wavelet inverse transformation can be categorized into Mean and Base-index inverse transformation based on how these high-frequency components are aggregated.

Mean Inverse Transformation. As presented in Fig. 3, we calculate the frequency element-wise average of all high-frequency components of K features

from the same category, and take the averaged results as the input of our Mean inverse transformation module.

$$HF_M = \sum_{i=1}^K HF_i, HF_i \in \{LH_i, HL_i, HH_i\} \quad (3)$$

Although providing high-frequency signals to the decoder facilitates the generated images’ quality, the Mean inverse transformation may not be suited for one specific image as the averaged frequency information may shift the frequency signals. The averaged frequency information becomes more neutral with the number of training images K increases, leading to a decrease in generalization. This conjecture may go against our common sense that the generalization ability should improve as the number of images increases. We analyze this is because different images, even from the same category, have different frequency signals in the frequency domain. The experiments in Sec. 4 confirm our analysis that the averaged frequency transformation may fail to generalize when K is bigger. To improve the generalization ability of our inverse transformation in the frequency domain, we design shots-agnostic Base-index inverse transformation.

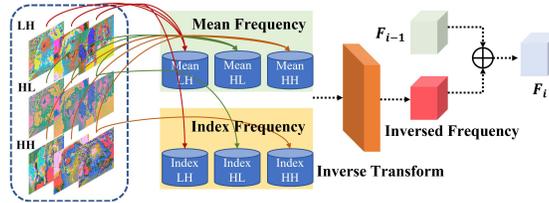


Fig. 3. Illustration of our Mean and Base-index inverse transformation.

Base-Index Inverse Transformation. Our Base-index inverse transformation is implemented based on the Local representation Fusion (LoF) strategy of LoFGAN [12]. We first give a brief introduction to LoFGAN. Given encoder features $\mathbb{F} = \mathbb{E}(X) \in \mathbb{R}^{k \times w \times h \times c}$, LoFGAN randomly selects one base feature $f_{base} \in \mathbb{R}^{w \times h \times c}$ and views the rest $(K - 1)$ features as reference, *i.e.*, $f_{ref} \in \mathbb{R}^{(k-1) \times w \times h \times c}$. LoFGAN fuses the local features based on the calculated semantic similarity map and replaces the closest base feature with the fused feature.

To provide customized high-frequency signals to the generated images, we recode the index of the selected base feature of the LoF module. As illustrated in Fig. 3, instead of calculating the averaged frequency signals, we explicitly feed the high-frequency component LH corresponding to the recorded index i .

$$HF_B = HF_i, HF_i \in \{LH_i, HL_i, HH_i\} \quad (4)$$

The high-frequency components are exact frequency signals of the selected feature, thus providing customized rich details and perceptible information to the

decoder. The generalization ability of our Base-index inverse transformation would not deteriorate with the training number increases.

After obtaining the aggregate high-frequency signals LH , LH , HH , we perform our inverse transformation to reconstruct these signals back to the original features. Concretely, we first perform transposed-convolution on each frequency component, then sum up all the output features. The summed result convert frequency signals back to original features precisely, theoretical analysis can be found in [40]. The inverse transformation can be formally expressed as:

$$F_{IF} = \sum TransConv(HF), HF \in \{HF_M, HF_B\} \quad (5)$$

HF is obtained from either the Mean or the Base-index inverse transformation. Element-wise tensor addition are employed to integrate the inversed high-frequency components with the former feature, *i.e.*, $F_i = F_{i-1} + F_{IF(n-i+1)}$. Such branches assuage the generator’s dilemmas in producing limited high-frequency containing fine details.

3.4 Optimization Objective

Our model has two networks to optimize, the generator (G) and the discriminator (D). The input of G is real images, and G tries to generate plausible and diverse new images $\hat{x} = G(X)$. Let $X = \sum_{i=1}^K x_i$ denotes the real images, and $c(x_i)$ denotes the labels for the image x_i (available for the \mathbb{C}_s only). The inputs of D are the real and generated images, and D tries to distinguish the real images from the generated ones. The generator and the discriminator are updated alternatively in an adversarial manner by optimizing the following losses.

Frequency L_1 -Loss. We employ the frequency L_1 -loss on the transformed frequency components of generated images and the real images, impeding losing the frequency information. Besides, the frequency loss complements existing spatial losses. We perform wavelet transformation on the generated and real images to compute the frequency L_1 loss.

$$\mathcal{L}_{Fre} = \sum \|Fre_x - Fre_{\hat{x}}\|_1, Fre \in \{LL, LH, HL, HH\} \quad (6)$$

Local Reconstruction Loss. We adopt local reconstruction loss to constrain the model to maintain the local features.

$$\mathcal{L}_{rec} = \|\hat{x} - LFM(X, \alpha)\|_1 \quad (7)$$

where α denotes the coefficient vector to fuse the features in the local fusion module (LoF).

Adversarial Loss. Following [12] and [31], we adopt the hinge version of adversarial loss to optimize the generator and the discriminator.

$$\begin{aligned} \mathcal{L}_{adv}^D &= \max(0, 1 - D(x)) + \max(0, 1 + D(\hat{x})) \\ \mathcal{L}_{adv}^G &= -D(\hat{x}) \end{aligned} \quad (8)$$

Classification Loss. Classification loss constrains the model to synthesis images that belong to one specific category. We add an auxiliary classifier to the generator and the discriminator following ACGAN [31]. The classification loss encourages the discriminator to identify which category an image belongs to while enabling the generator to synthesis images that belong to one specific category.

$$\begin{aligned}\mathcal{L}_{cls}^D &= -\log P(c(x) | x) \\ \mathcal{L}_{cls}^G &= -\log P(c(\hat{x}) | \hat{x})\end{aligned}\tag{9}$$

Our model is optimized with the following objective function with the linear combination of the above losses.

$$\begin{aligned}\mathcal{L}_G &= \mathcal{L}_{adv}^G + \lambda_{cls}^G \mathcal{L}_{cls}^G + \lambda_{Fre} \mathcal{L}_{Fre} + \lambda_{rec} \mathcal{L}_{rec}^G \\ \mathcal{L}_D &= \mathcal{L}_{adv}^D + \lambda_{cls}^D \mathcal{L}_{cls}^D\end{aligned}\tag{10}$$

4 Experiments

Datasets. We use three popular datasets in the few-shot image generation community to evaluate the performance of our model, namely Flower [30], Animal Faces [28] and VGGFace [3]. These datasets are split into seen classes \mathbb{C}_s and unseen classes \mathbb{C}_u . \mathbb{C}_s is used in the training stage, while \mathbb{C}_u is used in the testing stage. The adopted datasets are split in Tab. 1 following [12] and [15].

Table 1. The split of experimental datasets.

Datasets	#Total classes	#Seen classes	#Unseen classes	#Images/class
Flower	102	85	17	40
Animal Faces	149	119	30	100
VGGFace	2354	1802	552	100

Evaluation and Baselines. We evaluate the quality of the generated images with two commonly used metrics: Fréchet Inception Distance (FID) [13] and Learned Perceptual Image Patch Similarity (LPIPS) [42]. The two metrics are calculated in the same setting with [12]. We compare our model with several few-shot image generation approaches, namely FIGR [4], GMN [2], DAWSON [26], DAGAN [1], MathingGAN [14], F2GAN [15] and LoFGAN [12]. We re-implement the current state-of-the-art LoFGAN for fair comparisons (denoted as “LoFGAN[‡]”), and all methods are evaluated under the same conditions. Implementation details are given in the appendix.

4.1 Quantitative Evaluation

We first train the model with \mathbb{C}_s and then use the data from \mathbb{C}_u to synthesis novel images for quantitative evaluation. Following LoFGAN [12] and [15], we split each unseen class into two parts, \mathbb{S}_{sup} and \mathbb{S}_{que} , the images in \mathbb{S}_{sup} are fed into

the model to generate images. We generate 128 images for each class (denoted as S_{gen}), S_{gen} and S_{que} are used to compute FID (lower is better) and LPIPS (higher is better) scores to evaluate the synthesis quality. The quantitative results of our model and the baselines are given in Tab. 2. All results in the table are conducted under 3-shot setting for both training and testing stages.

Table 2. Quantitative comparison results of our model and the baselines on FID and LPIPS. [†] results are quoted from LoFGAN [12]. [‡] results are re-implemented under the same condition with our model for fair comparison. The best and the second-ranked results are **bold** and underlined, respectively.

Method	Type	Flowers		Animal Faces		VGGFace	
		FID (↓)	LPIPS (↑)	FID (↓)	LPIPS (↑)	FID (↓)	LPIPS (↑)
FIGR [†] [4]	Optimization	190.12	0.0634	211.54	0.0756	139.83	0.0834
DAWSON [†] [26]	Optimization	188.96	0.0583	208.68	0.0642	137.82	0.0769
DAGAN [†] [1]	Transformation	151.21	0.0812	155.29	0.0892	128.34	0.0913
GMN [†] [2]	Fusion	200.11	0.0743	220.45	0.0868	136.21	0.0902
MatchingGAN [†] [14]	Fusion	143.35	0.1627	148.52	0.1514	118.62	0.1695
F2GAN [†] [15]	Fusion	120.48	0.2172	117.74	0.1831	109.16	0.2125
MatchingGAN+LoFGAN [†] [12]	Fusion	86.59	0.3704	112.99	<u>0.5024</u>	22.99	0.2687
LoFGAN [†] [12]	Fusion	79.33	<u>0.3862</u>	112.81	0.4964	20.31	0.2869
LoFGAN [‡]	Fusion	81.70	0.3768	102.07	0.5005	16.82	0.3041
WaveGAN-M (Ours)	Fusion	<u>63.79</u>	0.3709	<u>50.98</u>	0.5014	<u>8.62</u>	0.3822
WaveGAN-B (Ours)	Fusion	42.17	0.3868	30.35	0.5076	4.96	<u>0.3255</u>

As can be observed from Tab. 2, our waveGAN achieves the lowest FID and the highest LPIPS on all the datasets, and both WaveGAN-M and WaveGAN-B outperform the baseline models. Notably, Our model achieves much better FID results than the current state-of-the-art LoFGAN. Specifically, WaveGAN-B achieves the FID of less than 5 (**4.96**) on the challenging VGGFace dataset, while LoFGAN obtains 16.82. And WaveGAN-B lowers the FID from 102.07 (*resp.*, 81.70) to **30.35** (*resp.*, **42.17**) on Animal Faces (*resp.*, Flower). Such significant improvements on the quantitative metrics demonstrate that our model could generate plausible and vivid images. Since the upper bound of FID scores equals 0 and we achieve a single digit for the first time, demonstrating the efficacy of our method. As for the LPIPS metric, we calculate its upper bound by measuring the real images' LPIPS score and obtain 0.4393 for Flower, 0.5729 for Animal Faces, and 0.4389 for VGGFace. Our model yields favorable LPIPS scores that approach the upper bound, further substantiating the effectiveness of our model.

4.2 Qualitative Evaluation

We present the visualization results of LoFGAN [12] and our WaveGAN-B for qualitative comparison in Fig. 4. For each real image, we give two fake images generated by LoFGAN and our WaveGAN-B. As can be observed from the figure, images generated by our model are more plausible than that of LoFGAN.

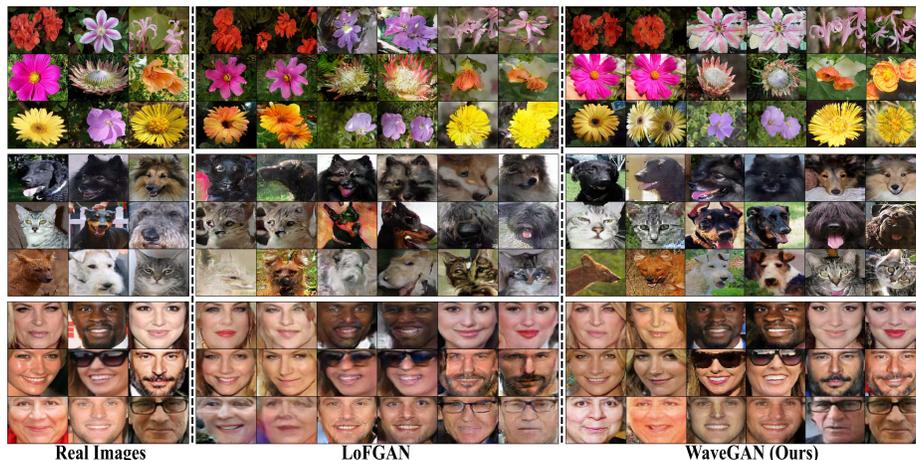


Fig. 4. Qualitative results of our WaveGAN and LoFGAN. The left-most three columns are real images, we give two generated images for each class of real image.

Moreover, images generated by our model contain rich details and perceptible information. Take the generated flowers as examples, the horizontal and vertical orientation of petals, the details of stamens, and the shapes of the leaves of images synthesised by our model are more reasonable and realistic. Moreover, animal and human face images generated by LoFGAN are distorted with blurry unfavorable artifacts, features like the eyes of cats, hair of dogs are even misplaced. By contrast, animal and human face images generated by our model have higher fidelity and even look indistinguishable from real images.

4.3 Visualization of the frequency components of generated images

We visualize the frequency components of images generated by our WaveGAN and LoFGAN [12] in Fig. 5. As observable in Fig. 5, in addition to the fact that our WaveGAN produces more realistic visual images, the decomposed high-frequency components of our WaveGAN contain more details and perceptible information than LoFGAN. Specifically, the frequency components of LoFGAN contain only the surface and texture information of images, indicating that the generator of LoFGAN fails to synthesize high-frequency information. In comparison, our WaveGAN is frequency-aware and can produce high-frequency signals that contain more fine details and statistical properties. Further, the frequency components of WaveGAN capture delicate information that is less noticeable (*e.g.*, the glasses frames and flower rhizomes in the second and fifth row of Fig. 5, respectively). Such observation further demonstrates the effectiveness and advancement of our method.

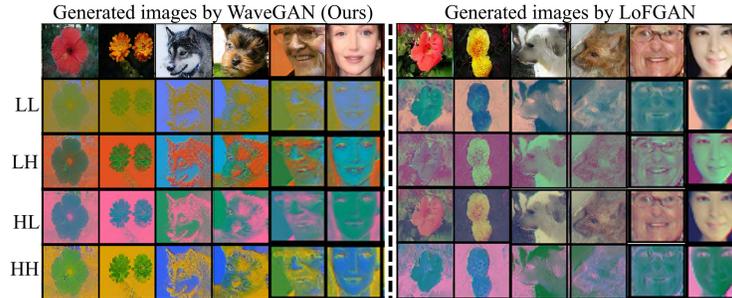


Fig. 5. Visualization results of the frequency components of images generated by our WaveGAN and LoFGAN [12].

4.4 Ablation Studies

We conduct ablation studies to evaluate the effectiveness of each component of the proposed WaveGAN. There are three main components of our WaveGAN, namely 1) the low-frequency skip connection, 2) the high-frequency skip connection, and 3) frequency L_1 -loss. We remove each component and keep other settings unchanged to validate their contributions. Besides, we remove the LoF module to investigate the influence of local fusion on our model. We also test the contribution of each component for our two transformation techniques (*i.e.*, WaveGAN-M and WaveGAN-B) quantitatively in the appendix.

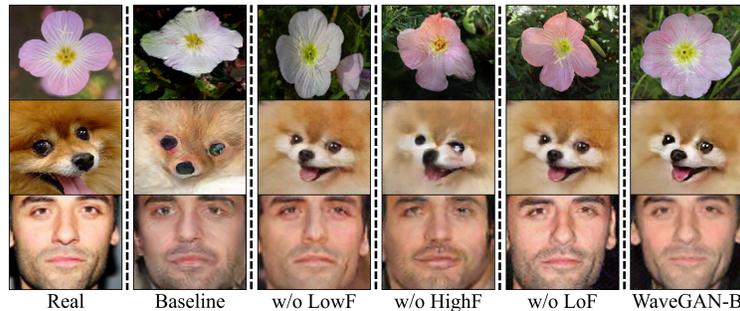


Fig. 6. Visualization comparison results of ablation studies.

We give visualization results of ablation studies in Fig. 6. The quantitative results of ablation studies are given in the appendix. Combining the qualitative and quantitative results, we can draw four conclusions as follows: 1) The skip connection of both low and high-frequency signals play an essential role in our model, and our high-frequency skip connections play a predominant role in our model. 2) High-frequency information provides detailed information to the generated images, and low-frequency information provides the overall outline of

images (compare baseline with other generated images in Fig. 6). 3) Our full model achieves the best results and can generate satisfactory images. 4) Our method complements to the local feature fusion approach.

4.5 Augmentation for Classification

To further investigate the quality of generated images, we augment the datasets with images generated by our model for downstream image classification tasks. Specifically, we first pre-train a ResNet18 network with seen classes following [12] and [15], we train the ResNet18 model for 100 epochs with batch size of 4. Then we split the unseen datasets into $\mathbb{D}_{\text{train}}$, \mathbb{D}_{test} and \mathbb{D}_{val} . For each category of the flower dataset, the number of train, test, and valid images are 10, 15, and 15, respectively. For each category of Animal Faces and VGGFace dataset, the number of train, test and valid images are 30, 35, and 35, respectively. We train a new classifier using the pre-trained model on seen classes with $\mathbb{D}_{\text{train}}$ without any augmentation, which is denoted as “Base”. Then we generate images to augment $\mathbb{D}_{\text{train}}$ with LoFGAN and our WaveGAN. The number of augmented images is 30 for Flower dataset and 50 for Animal Face and VGGFace datasets.

Table 3. Classification results of augmentation.

Datasets	Base	LoFGAN	WaveGAN-M (ours)	WaveGAN-B (ours)
Flower	64.71	80.78	70.20	84.71
Animals	20.00	26.10	31.81	32.19
VGGFace	50.76	64.74	62.96	77.36

The classification results are given in Tab. 3. Compared with the results without any augmentation, our model achieves significant improvements, and our WaveGAN-B outperforms LoFGAN and WaveGAN-M obviously. The effectiveness of using the generated images to augment the training dataset substantiates that our model can produce high-quality images, and the improvement on the classification accuracy provides a new data augmentation strategy for solving few-shot image classification problems.

4.6 Influence of the Number of Shots

All the experiments conducted before are 3-shot image generation tasks. We wonder to know the influence of different shots on our model. We perform different shots of experiments with $K \in \{2, 3, 5, 7, 9\}$. The number of training and testing images for different shots of the experiment are the same. Fig. 7 demonstrates the performance of our WaveGAN and LoFGAN. We can observe from the figure that when K is relatively small, WaveGAN-M is better than LoFGAN. However, the performance of WaveGAN-M degrades with the number of images increases, making WaveGAN-M inadequate for image generation tasks

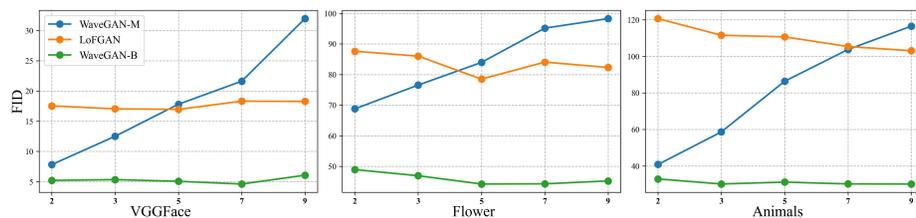


Fig. 7. Comparison results of our model under different shots of generation tasks. The ordinate denotes FID scores, and the abscissa denotes different shots K .

with larger K . Such phenomena corroborates our analysis in Sec. 3 that the averaged transformation may fail to generalize to one specific image. By contrast, our WaveGAN-B has relatively low sensitivity to K and is shots-agnostic, thus WaveGAN-B has stronger generalization ability for different shots generation tasks, which manifests the superiority of our WaveGAN-B.

5 Conclusion and Acknowledgments

Conclusion. In this paper, we propose WaveGAN, the first few-shot image generation model that ameliorates the synthesis quality from the frequency domain perspective. The key ingredients of our method are a WaveEncoder and a WaveDecoder. Our WaveEncoder performs wavelet transformation on the different levels of features to obtain frequency signals. We feed the decomposed low-frequency signals to the behind layers of the encoder, and we feed the high-frequency signals to our WaveDecoder. Our design mitigates the generator’s struggles of generating rich details for images, especially when limited data are available. We further perform frequency L_1 -loss to maintain the frequency information of real images, facilitating the fidelity of generated images. Experimentally, our WaveGAN yields significant improvements on three challenging datasets, and the visualization and downstream classification results demonstrate that our WaveGAN can produce realistic images. Besides, the ablation studies suggest the efficacy of each component of our method and substantiate that our approach complements the existing local fusion-based strategy. Hopefully, our WaveGAN may inspire researchers to explore few-shot image generation from the frequency domain perspective.

Acknowledgments. This work is supported by Shanghai Science and Technology Program “Distributed and generative few-shot algorithm and theory research” under Grant No. 20511100600 and “Federated based cross-domain and cross-task incremental learning” under Grant No. 21511100800, Natural Science Foundation of China under Grant No. 62076094, Chinese Defense Program of Science and Technology under Grant No.2021-JCJQ-JJ-0041, China Aerospace Science and Technology Corporation Industry-University-Research Cooperation Foundation of the Eighth Research Institute under Grant No.SAST2021-007.

References

1. Antoniou, A., Storkey, A., Edwards, H.: Data augmentation generative adversarial networks. arXiv preprint arXiv:1711.04340 (2017)
2. Bartunov, S., Vetrov, D.: Few-shot generative modelling with generative matching networks. In: International Conference on Artificial Intelligence and Statistics. pp. 670–678. PMLR (2018)
3. Cao, Q., Shen, L., Xie, W., Parkhi, O.M., Zisserman, A.: Vggface2: A dataset for recognising faces across pose and age. In: 2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018). pp. 67–74. IEEE (2018)
4. Clouâtre, L., Demers, M.: Figr: Few-shot image generation with reptile. arXiv preprint arXiv:1901.02199 (2019)
5. Daubechies, I.: The wavelet transform, time-frequency localization and signal analysis. IEEE transactions on information theory **36**(5), 961–1005 (1990)
6. Deng, X., Yang, R., Xu, M., Dragotti, P.L.: Wavelet domain style transfer for an effective perception-distortion tradeoff in single image super-resolution. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3076–3085 (2019)
7. Fei-Fei, L., Fergus, R., Perona, P.: One-shot learning of object categories. IEEE transactions on pattern analysis and machine intelligence **28**(4), 594–611 (2006)
8. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: International conference on machine learning. pp. 1126–1135. PMLR (2017)
9. Gao, Y., Wei, F., Bao, J., Gu, S., Chen, D., Wen, F., Lian, Z.: High-fidelity and arbitrary face editing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16115–16124 (2021)
10. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. Advances in neural information processing systems **27** (2014)
11. Gu, S., Chen, D., Bao, J., Wen, F., Zhang, B., Chen, D., Yuan, L., Guo, B.: Vector quantized diffusion model for text-to-image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10696–10706 (2022)
12. Gu, Z., Li, W., Huo, J., Wang, L., Gao, Y.: Lofgan: Fusing local representations for few-shot image generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8463–8471 (2021)
13. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems **30** (2017)
14. Hong, Y., Niu, L., Zhang, J., Zhang, L.: Matchinggan: Matching-based few-shot image generation. In: 2020 IEEE International Conference on Multimedia and Expo (ICME). pp. 1–6. IEEE (2020)
15. Hong, Y., Niu, L., Zhang, J., Zhao, W., Fu, C., Zhang, L.: F2gan: Fusing-and-filling gan for few-shot image generation. In: Proceedings of the 28th ACM International Conference on Multimedia. pp. 2535–2543 (2020)
16. Huang, H., He, R., Sun, Z., Tan, T.: Wavelet domain generative adversarial network for multi-scale face hallucination. International Journal of Computer Vision **127**(6), 763–784 (2019)
17. Jiang, L., Dai, B., Wu, W., Loy, C.C.: Deceive d: Adaptive pseudo augmentation for gan training with limited data. Advances in Neural Information Processing Systems **34** (2021)

18. Jiang, L., Dai, B., Wu, W., Loy, C.C.: Focal frequency loss for image reconstruction and synthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13919–13929 (2021)
19. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. In: International Conference on Learning Representations (2018)
20. Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., Aila, T.: Training generative adversarial networks with limited data. *Advances in Neural Information Processing Systems* **33**, 12104–12114 (2020)
21. Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., Aila, T.: Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems* **34** (2021)
22. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4401–4410 (2019)
23. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8110–8119 (2020)
24. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013)
25. Li, G., Jampani, V., Sevilla-Lara, L., Sun, D., Kim, J., Kim, J.: Adaptive prototype learning and allocation for few-shot segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8334–8343 (2021)
26. Liang, W., Liu, Z., Liu, C.: Dawson: A domain adaptive few shot generation framework. *arXiv preprint arXiv:2001.00576* (2020)
27. Liu, B., Cao, Y., Lin, Y., Li, Q., Zhang, Z., Long, M., Hu, H.: Negative margin matters: Understanding margin in few-shot classification. In: European Conference on Computer Vision. pp. 438–455. Springer (2020)
28. Liu, M.Y., Huang, X., Mallya, A., Karras, T., Aila, T., Lehtinen, J., Kautz, J.: Few-shot unsupervised image-to-image translation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10551–10560 (2019)
29. Nichol, A., Schulman, J.: Reptile: a scalable metalearning algorithm. *arXiv preprint arXiv:1803.02999* **2**(3), 4 (2018)
30. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: 2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing. pp. 722–729. IEEE (2008)
31. Odena, A., Olah, C., Shlens, J.: Conditional image synthesis with auxiliary classifier gans. In: International conference on machine learning. pp. 2642–2651. PMLR (2017)
32. Park, T., Efros, A.A., Zhang, R., Zhu, J.Y.: Contrastive learning for unpaired image-to-image translation. In: European Conference on Computer Vision. pp. 319–345. Springer (2020)
33. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125* (2022)
34. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: International Conference on Machine Learning. pp. 8821–8831. PMLR (2021)

35. Richardson, E., Alaluf, Y., Patashnik, O., Nitzan, Y., Azar, Y., Shapiro, S., Cohen-Or, D.: Encoding in style: a stylegan encoder for image-to-image translation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2287–2296 (2021)
36. Tseng, H.Y., Jiang, L., Liu, C., Yang, M.H., Yang, W.: Regularizing generative adversarial networks under limited data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7921–7931 (2021)
37. Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., Wierstra, D.: Matching networks for one shot learning. In: Proceedings of the 30th International Conference on Neural Information Processing Systems. pp. 3637–3645 (2016)
38. Wang, L., Ho, Y.S., Yoon, K.J., et al.: Event-based high dynamic range image and very high frame rate video generation using conditional generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10081–10090 (2019)
39. Xu, Z.Q.J., Zhang, Y., Luo, T., Xiao, Y., Ma, Z.: Frequency principle: Fourier analysis sheds light on deep neural networks. arXiv preprint arXiv:1901.06523 (2019)
40. Yoo, J., Uh, Y., Chun, S., Kang, B., Ha, J.W.: Photorealistic style transfer via wavelet transforms. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9036–9045 (2019)
41. Yu, Y., Zhan, F., Lu, S., Pan, J., Ma, F., Xie, X., Miao, C.: Wavefill: A wavelet-based generation network for image inpainting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14114–14123 (2021)
42. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)
43. Zhao, S., Liu, Z., Lin, J., Zhu, J.Y., Han, S.: Differentiable augmentation for data-efficient gan training. *Advances in Neural Information Processing Systems* **33**, 7559–7570 (2020)