

Wavelet Diffusion Models are fast and scalable Image Generators

Hao Phung[†], Quan Dao[†], Anh Tran
VinAI Research

{v.haopt12, v.quandm7, v.anhtt152}@vinai.io

Abstract

Diffusion models are rising as a powerful solution for high-fidelity image generation, which exceeds GANs in quality in many circumstances. However, their slow training and inference speed is a huge bottleneck, blocking them from being used in real-time applications. A recent DiffusionGAN method significantly decreases the models' running time by reducing the number of sampling steps from thousands to several, but their speeds still largely lag behind the GAN counterparts. This paper aims to reduce the speed gap by proposing a novel wavelet-based diffusion scheme. We extract low-and-high frequency components from both image and feature levels via wavelet decomposition and adaptively handle these components for faster processing while maintaining good generation quality. Furthermore, we propose to use a reconstruction term, which effectively boosts the model training convergence. Experimental results on CelebA-HQ, CIFAR-10, LSUN-Church, and STL-10 datasets prove our solution is a stepping-stone to offering real-time and high-fidelity diffusion models. Our code and pre-trained checkpoints are available at <https://github.com/VinAIRResearch/WaveDiff.git>.

1. Introduction

Despite being introduced recently, diffusion models have grown tremendously and drawn many research interests. Such models revert the diffusion process to generate clean, high-quality outputs from random noise inputs. These techniques are applied in various data domains and applications but show the most remarkable success in image-generation tasks. Diffusion models can beat the state-of-the-art generative adversarial networks (GANs) in generation quality on various datasets [4, 37]. More notably, diffusion models provide a better mode coverage [14, 22, 40] and a flexible way to handle different types of conditional inputs such as semantic maps, text, representations, and images [36].

[†]Equal contribution.




Figure 1. Comparisons between our method and other GAN and diffusion techniques in terms of FID and sampling time on the CIFAR-10 (32×32) dataset. Our method is $2.5\times$ faster than DDGAN [49], the fastest up-to-date diffusion method, and approaches the real-time speed of StyleGAN methods [19, 20, 56] while still achieving comparable FID scores.

Thanks to this capability, they offer various applications such as text-to-image generation, image-to-image translation, image inpainting, image restoration, and more. Recent diffusion-based text-to-image generative models [1, 34, 37] allow users to generate unbelievably realistic images just by text inputs, opening a new era of AI-based digital art and promising applications to various other domains.

While showing great potential, diffusion models have a very slow running speed, a critical weakness blocking them from being widely adopted like GANs. The foundation work Denoising Diffusion Probabilistic Models (DDPMs) [13] requires a thousand sampling steps to produce the desired output quality, taking minutes to generate a single image. Many techniques have been proposed to reduce the inference time [25, 39], mainly via reducing the sampling steps. However, the fastest algorithm before DiffusionGAN still takes seconds to produce a 32×32 image, which is about 100 times slower than GAN. DiffusionGAN [49] made a break-through in fastening inference speed by combining Diffusion and GANs in a single system, which ul-

timately reduces the sampling steps to 4 and the inference time to generate a 32×32 image as a fraction of a second. It makes DiffusionGAN the fastest existing diffusion model. Still, it is at least 4 times slower than the StyleGAN counterpart, and the speed gap consistently grows when increasing the output resolution. Moreover, DiffusionGAN still requires a long training time and a slow convergence, confirming that diffusion models are not yet ready for large-scale or real-time applications.

This paper aims to bridge the speed gap by introducing a novel wavelet-based diffusion scheme. Our solution relies on discrete wavelet transform, which decomposes each input into four sub-bands for low- (LL) and high-frequency (LH, HL, HH) components. We apply that transform on both image and feature levels. This allows us to significantly reduce both training and inference times while keeping the output quality relatively unchanged. On the image level, we obtain a high speed boost by reducing the spatial resolution four times. On the feature level, we stress the importance of wavelet information on different blocks of the generator. With such a design, we can obtain considerable performance improvement while inducing only a marginal computing overhead.

Our proposed Wavelet Diffusion provides state-of-the-art training and inference speed while maintaining high generative quality, thoroughly confirmed via experiments on standard benchmarks including CIFAR-10, STL-10, CelebA-HQ, and LSUN-Church. Our models significantly reduce the speed gap between diffusion models and GANs, targeting large-scale and real-time systems.

In summary, our contributions are as following:

- We propose a novel Wavelet Diffusion framework that takes advantage of the dimensional reduction of Wavelet subbands to accelerate Diffusion Models while maintaining good visual quality of generated results through high-frequency components.
- We employ wavelet decomposition in both image and feature space to improve generative models' robustness and execution speed.
- Our proposed Wavelet Diffusion provides state-of-the-art training and inference speed, which serves as a stepping-stone to facilitating real-time and high-fidelity diffusion models.

2. Related work

2.1. Diffusion models

Diffusion models [13, 38] are inspired by non-equilibrium thermodynamics where the diffusion and reverse processes are Markov chains. Unlike GANs [10, 20] and VAEs [24, 45], they sample using a large number of

denoising steps with a fixed procedure where each latent variable shares the same dimensionality as the original inputs. A line of methods shares the same motives but relies on score matching for the reverse process [41, 43, 47]. Later works focus on improving sample quality of diffusion models [4, 40, 42]. Some have shifted the process to latent space [36, 46], which enables the success of text-to-image generation [5, 34, 35, 37]. Similar to us, some approaches [30, 31] aim to improve the efficiency of sampling process for better convergence and faster sampling. As the Markov process is a sequential probability model with a finite step, some try to break it into non-Markov chains for faster sampling [39]. Despite several efforts to accelerate the sampling process, there is an inevitable trade-off between sampling speed and quality.

Diffusion GAN [49] was recently proposed as a novel approach to sidestep the unimodal Gaussian assumption of small step sizes via modeling the complex multimodal distribution of large step sizes with generative adversarial networks. This proposal greatly reduces the number of denoising steps to just a finger count (e.g., 2, 4), which leads the inference time to a fraction of a second. However, it is still largely slower than GAN competitors. Hence, we further maximize its full potential by introducing new wavelet components on top of this framework.

2.2. Wavelet-based approaches

Wavelet decomposition [2, 32] is a classical operator widely used in various computer vision tasks. It is the fundamental process behind a popular JPEG-2000 image compression format [44]. In recent years, wavelet transform has started being incorporated into many deep-learning-based systems as they can take advantage of spatial and frequency information to aid the training. Some have applied it to the design of networks to improve visual representation learning [16, 28, 29, 52]. Some have extended it to image restoration [3, 54], style transfer [53], and face problems [9, 15]. On the line of image generation, wavelet transform is plugged in recent generative adversarial networks [7, 50, 51, 55] to improve the visual quality of output images.

Regarding diffusion model families, several works start employing wavelet transform. Guth et al. [11] accelerate score-based models by modeling conditional probabilities of multiscale wavelet coefficients, but still suffer from low-quality output images. Li et al. [27] perform score matching on the wavelet domain for better image colorization.

However, none of the mentioned techniques can balance the diffusion models' quality and speed. Here, we introduce a novel wavelet diffusion approach to not only emphasize the importance of frequency information but also to reduce the sampling time by a large margin. We take advantage of frequency sparsity and dimensionality reduction of wavelet transform to reduce high-dimensional space to the

lower-dimensional manifolds, which are essential for faster and more efficient sampling. Unlike existing methods on diffusion families, we enrich the frequency awareness on both image and feature levels to facilitate high-fidelity image generation. This is greatly inspired by the success of wavelet-based GANs.

3. Background

3.1. Diffusion Models

The traditional diffusion process requires thousand timesteps to gradually diffuse each input x_0 , following a data distribution $p(x_0)$, into pure Gaussian noise. The posterior probability of a diffused image x_t at timestep t has a closed form:

$$q(x_t|x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_t)\mathbf{I}) \quad (1)$$

where $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$, and $\beta_t \in (0, 1)$ is defined to be small through a variance schedule which can be learnable or fixed at timestep t in the forward process.

Since the diffusion process adds relatively small noise per step, the reverse process $q(x_{t-1}|x_t)$ can be approximated by Gaussian process $q(x_{t-1}|x_t, x_0)$. Therefore, the trained denoising process $p_\theta(x_{t-1}|x_t)$ can be parameterized according to $q(x_{t-1}|x_t, x_0)$. The common parameterized form of $p_\theta(x_{t-1}|x_t)$ is:

$$p_\theta(x_{t-1} | x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \sigma_t^2 \mathbf{I}), \quad (2)$$

where $\mu_\theta(x_t, t)$ and σ_t^2 are the mean and variance of the parametric denoising model, respectively. The objective is to minimize the distance between a true denoising distribution $q(x_{t-1}|x_t)$ and the parameterized one $p_\theta(x_{t-1}|x_t)$ through Kullback-Leibler (KL) divergence.

Unlike the traditional diffusion methods, DiffusionGAN [49] enables large step sizes for faster sampling through generative adversarial networks. It introduces a discriminator D_ϕ and optimizes both the generator and the discriminator in an adversarial training manner:

$$\min_{\phi} \max_{\theta} \sum_{t \geq 1} \mathbb{E}_{q(\mathbf{x}_t)} [\mathbb{E}_{q(\mathbf{x}_{t-1}|\mathbf{x}_t)} [-\log(D_\phi(\mathbf{x}_{t-1}, \mathbf{x}_t, t))] + \mathbb{E}_{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)} [\log(D_\phi(\mathbf{x}_{t-1}, \mathbf{x}_t, t))], \quad (3)$$

where fake samples are sampled from a conditional generator $p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t)$. As large step sizes cause $q(x_{t-1}|x_t)$ to be no longer a Gaussian distribution, DiffusionGAN [49] aims to implicitly model this complex multimodal distribution with a generator $G_\theta(x_t, z, t)$ given a D -dimensional latent variable $z \sim \mathcal{N}(0, \mathbf{I})$. Specifically, DiffusionGAN first generates unperturbed sample x'_0 through the generator $G_\theta(x_t, z, t)$ and acquires the corresponding perturbed sample x'_{t-1} using $q(x_{t-1}|x_t, x_0)$. Meanwhile, the discriminator performs judgment on real pairs $D_\phi(x_{t-1}, x_t, t)$ and fake pairs $D_\phi(x'_{t-1}, x_t, t)$.

For convenience, we will abbreviate DiffusionGAN as DDGAN in later sections.

3.2. Wavelet Transform

Wavelet Transform is a classical technique widely used in image compression to separate the low-frequency approximation and the high-frequency details from the original image. While low subbands are similar to down-sampled versions of the original image, high subbands express the local statistics of vertical, horizontal, and diagonal edges. Notably, the Haar wavelet is widely adopted in real-world applications due to its simplicity. It involves two types of operations: discrete wavelet transform (DWT) and discrete inverse wavelet transform (IWT).

Let denote $L = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \end{bmatrix}$ and $H = \frac{1}{\sqrt{2}} \begin{bmatrix} -1 & 1 \end{bmatrix}$ be low-pass and high-pass filters. They are used for constructing four kernels with stride 2, namely LL^T, LH^T, HL^T , and HH^T to decompose the input $X \in \mathbb{R}^{H \times W}$ into four subbands X_{ll}, X_{lh}, X_{hl} , and X_{hh} with a size of $H/2 \times W/2$. As these filters are pairwise orthogonal, they can form a 4×4 invertible matrix to accurately reconstruct original signals X from frequency components by IWT.

In this paper, we use this kind of transform to decompose input images and feature maps to emphasize high-frequency components and reduce the spatial dimensions to four folds for more efficient sampling.

4. Method

This section describes our proposed Wavelet Diffusion framework. First, we present the core wavelet-based diffusion scheme for more efficient sampling (Sec. 4.1). We then depict the design of a new wavelet-embedded generator for better frequency-aware image generation (Sec. 4.2).

4.1. Wavelet-based diffusion scheme

First, we describe how to incorporate wavelet transform in the diffusion process. We decompose the input image into four wavelet subbands and concatenate them as a single target for the denoising process (illustrated in Fig. 2). Such a model does not perform on the original image space but on the wavelet spectrum. As a result, our model can leverage high-frequency information to increase the details of generated images further. Meanwhile, the spatial area of wavelet subbands is four times smaller than the original image, so the computational complexity of the sampling process is significantly reduced.

We build our method upon DDGAN model where the input is 4 wavelet subbands of wavelet transformation. Given input image $x \in \mathbb{R}^{3 \times H \times W}$, we decompose it into a set of low and high subbands which are further concatenated to form a matrix $y \in \mathbb{R}^{12 \times \frac{H}{2} \times \frac{W}{2}}$. This input is then **projected to the base channel D** via the first linear layer, keeping the

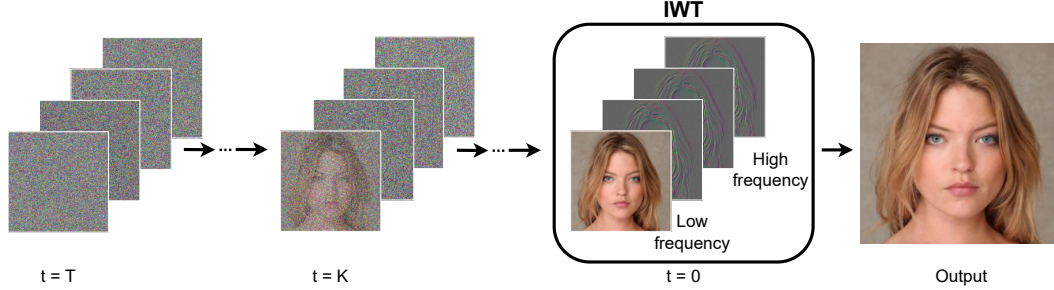


Figure 2. Illustration of Wavelet-based diffusion scheme. It performs denoising on wavelet space instead of pixel space. At each step t , a less-noisy sample y_{t-1} is generated by a denoiser $p_{\theta}(y_{t-1}|y_t)$ with parameters θ . After obtaining the clean sample y_0 through T steps, it is used to reconstruct the final image via Inverse wavelet transformation (IWT).

Algorithm 1 Wavelet-based sampling process

```

 $y_t \sim \mathcal{N}(0, \mathbf{I})$ 
for  $t = T, \dots, 1$  do
     $z \sim \mathcal{N}(0, \mathbf{I})$ 
     $y'_0 = G(y_t, z, t)$ 
     $y_{t-1} \sim q(y_{t-1}|y_t, y'_0)$ 
end for
 $x_0 = \text{IWT}(y_0)$ 
return  $x_0$ 

```

network width unchanged compared with DDGAN. Hence, most of the network benefits from $4\times$ reduction in spatial dimensions, significantly reducing its computation.

Let denote y_0 be a clean sample and y_t be a corrupted sample at timestep t which is sampled from $q(y_t|y_0)$. In terms of the denoising process, a generator receives a tuple of variable y_t , a latent $z \sim \mathcal{N}(0, \mathbf{I})$, and a timestep t to generate an approximation of original signals y_0 : $y'_0 = G(y_t, z, t)$. The predicted noisy sample y'_{t-1} is then drawn from tractable posterior distribution $q(y_{t-1}|y_t, y'_0)$. The role of the discriminator is to distinguish the real pairs (y_{t-1}, y_t) and the fake pairs (y'_{t-1}, y_t) .

Adversarial objective Following [49], we optimize the generator and the discriminator through the adversarial loss:

$$\begin{aligned} \mathcal{L}_{adv}^D &= -\log(D(y_{t-1}, y_t, t)) + \log(D(y'_{t-1}, y_t, t)), \\ \mathcal{L}_{adv}^G &= -\log(D(y'_{t-1}, y_t, t)). \end{aligned} \quad (4)$$

Reconstruction term In addition to the adversarial objective in Eq. (4), we add a reconstruction term to not only impede the loss of frequency information but also preserve the consistency of wavelet subbands. It is formulated as an L1 loss between a generated image and its ground-truth:

$$\mathcal{L}_{rec} = \|y'_0 - y_0\|. \quad (5)$$

The overall objective of the generator is a linear combination of adversarial loss and reconstruction loss:

$$\mathcal{L}^G = \mathcal{L}_{adv}^G + \lambda \mathcal{L}_{rec}, \quad (6)$$

where λ is a weighting hyper-parameter (default value is 1).

After a few sampling steps as defined, we acquire the estimated denoised subbands y'_0 . The final image can be recovered via wavelet inverse transformation $x'_0 = \text{IWT}(y'_0)$. We depict the sampling process in Algorithm 1.

4.2. Wavelet-embedded networks

Next, we further incorporate wavelet information into feature space through the generator to strengthen the awareness of high-frequency components. This is beneficial to the sharpness and quality of final images.

Fig. 3 illustrates the structure of our proposed wavelet-embedded generator. It follows the UNet structure of [43] with M down-sampling and M up-sampling blocks plus skip connections between blocks of the same resolution, with M predefined. However, instead of using the normal downsampling and upsampling operators, we replace them with frequency-aware blocks. At the lowest resolution, we employ frequency-bottleneck blocks for better attention on low and high-frequency components. Finally, to incorporate original signals Y to different feature pyramids of the encoder, we introduce frequency residual connections using wavelet downsample layers. Let denote Y be the input image and F_i is the i -th intermediate feature map of Y . We will discuss below the newly introduced components:

Frequency-aware downsampling and upsampling blocks. Traditional approaches relied on a blurring kernel for the downsampling and upsampling process to mitigate the aliasing artifact. We instead utilize inherent properties of the wavelet transform for better upsampling and downsampling (depicted in Fig. 4). This, in fact, strengthens the awareness of high-frequency information on these operations. Particularly, the downsampling block receives a tuple of input features F_i , a latent z , and time embedding t , which are then processed through a sequence of layers to return downsampled features and high-frequency subbands. These returned subbands are served as an additional input to upsample features based on frequency cues in the upsampling block.

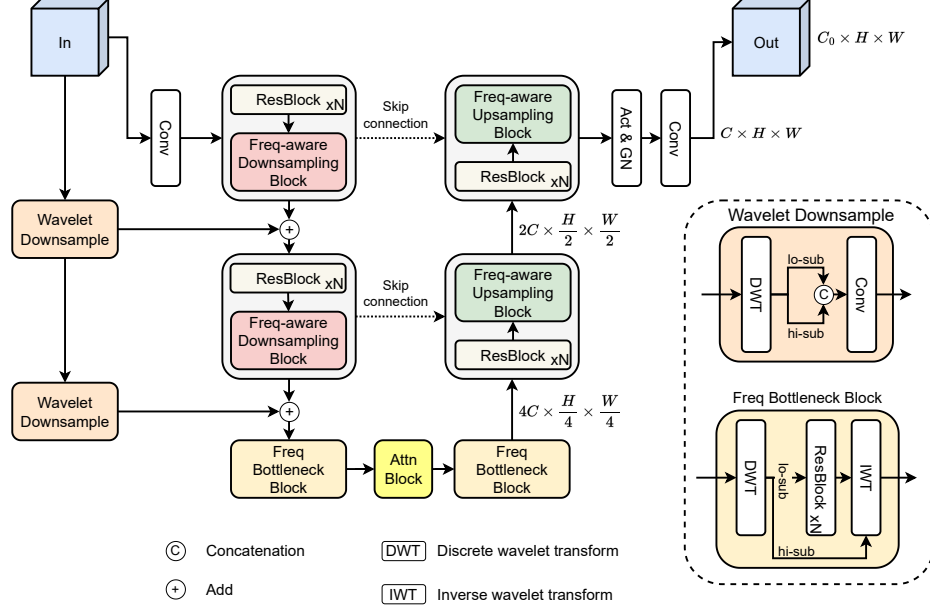


Figure 3. Illustration of Wavelet-embedded generator. For simplification, timestep embedding t and latent embedding z are ignored but they are injected in individual blocks of the denoising process. The inputs are noisy wavelet subbands of shape $[12 \times H \times W]$ at timestep t , which are processed by a sequence of our proposed components, including frequency-aware upsampling and downsampling blocks, frequency residual connections, and a brand new frequency bottleneck block. The outputs of the model are the approximation of unperturbed inputs.

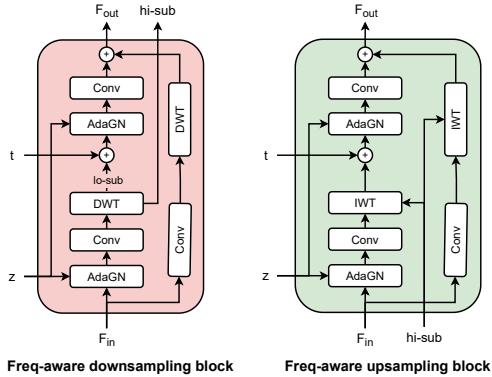


Figure 4. Details of frequency-aware downsampling and upsampling blocks in the wavelet-embedded generator.

Frequency bottleneck block locates at the middle stage, which includes two frequency bottleneck blocks and one attention block in-between. Each frequency bottleneck block first divides feature map F_i into the low-frequency subband $F_{i,l}$ and the concatenation of high-frequency subbands $F_{i,H}$. $F_{i,l}$ is then passed as input to resnet block(s) for deeper processing. The processed low-frequency feature map and the original high-frequency subbands $F_{i,H}$ are transformed back to the original space via IWT. With such a bottleneck, the model can focus on learning intermediate feature representations of low-frequency subbands while preserving the high-frequency details.

Frequency residual connection The original design of the network in [43] incorporates original signals Y to different feature pyramids of the encoder via a strided-convolution downsampling layer. We instead use a wavelet downsample layer to map residual shortcuts of input Y to the corresponding feature dimensions, which are then added to each feature pyramid. Specifically, the residual shortcuts of Y are decomposed into four subbands which are then concatenated and fed to a convolution layer for feature projection. This shortcut aims to enrich the perception of the frequency source of feature embeddings.

5. Experiments

In this section, we first provide details about the experimental setup and then present the empirical experiments on different datasets, including CIFAR-10, STL-10, CelebA-HQ, and LSUN. Finally, we ablate the important components of our proposed framework.

5.1. Experimental setup

Datasets We experiment on CIFAR-10 32×32 , STL-10 64×64 and two other higher-resolution datasets, including CelebA-HQ 256×256 and LSUN-Church 256×256 . We also examine our model training on high-resolution images with CelebA-HQ (512 & 1024).

Evaluation metrics We measure image fidelity by Fréchet

	DDGAN			Ours		
	Params	FLOPs	MEM	Params	FLOPs	MEM
CIFAR-10	48.43M	7.05G	0.31G	33.37M	1.67G	0.16G
STL-10	48.43M	27.26G	0.66G	55.58M	7.15G	0.34G
LSUN & CelebA HQ (256)	39.73M	70.82G	3.21G	31.48M	28.54G	1.07G
CelebA HQ (512)	39.73M	282.00G	12.30G	44.85M	74.35G	3.22G

Table 1. Model specifications of DDGAN [49] and our approach including a number of parameters (M), FLOPs (GB), and memory usage (GB) on a single GPU for one sample. Our method attains lower computing FLOPs and memory consumption while having a comparable number of parameters.

Model	FID↓	Recall↑	NFE↓	Time (s)↓
Ours	4.01	0.55	4	0.08
DDGAN [49]	3.75	0.57	4	0.21 (0.30*)
DDPM [13]	3.21	0.57	1000	80.5
NCSN [41]	25.3	-	1000	107.9
Score SDE (VE) [43]	2.20	0.59	2000	423.2
Score SDE (VP) [43]	2.41	0.59	2000	421.5
DDIM [39]	4.67	0.53	50	4.01
FastDDPM [25]	3.41	0.56	50	4.01
Recovery EBM [8]	9.58	-	180	-
DDPM Distillation [30]	9.36	0.51	1	-
StyleGAN2 w/o ADA [21]	8.32	0.41	1	0.04
StyleGAN2 w/ ADA [19]	2.92	0.49	1	0.04
StyleGAN2 w/ DiffAug [19]	5.79	0.42	1	0.04
Glow [23]	48.9	-	1	-
PixelCNN [33]	65.9	-	1024	-
NVAE [45]	23.5	0.51	1	0.36
VAEBM [48]	12.2	0.53	16	8.79

Table 2. Results on CIFAR-10. Among diffusion models, our method attains a state-of-the-art speed that is close to the speed of StyleGAN2 [19, 56] counterparts while preserving comparable image fidelity. Since the image resolution is too small (32), we do not use the wavelet-embedded generator in our network on this set. ‘**’ means the speed reproduced on our machine.

inception distance (FID) [12] and measure sample diversity by Recall metric [26]. Following [17], FID and Recall are computed over 50k generated samples. To further demonstrate our faster sampling, we measure the average inference time over 300 trials for a batch size of 100. Besides, the inference time of high-resolution images like CelebA-HQ 512×512 is computed from batches of 25 samples.

Implementation details Our implementation is mainly based on DDGAN [49]. We adopt the same training configurations as DDGAN [49] for all experiments. Training epochs are set as 500 for CelebA-HQ 256×256, 500 for LSUN 256×256, and 1800 for CIFAR-10 32×32. For CelebA-HQ (512 & 1024), we train our model and DDGAN for 400 epochs. We train our models on 1 - 8 NVIDIA A100 GPUs for each corresponding dataset. Notably, our models require fewer GPU resources and computations than DDGAN [49] thanks to the efficiency of our wavelet diffusion framework as demonstrated in Tab. 1. At the same

Model	FID↓	Recall↑	Time (s)↓
Ours + W-Generator	12.93	0.41	0.38
DDGAN [49]	21.79	0.40	0.58
StyleGAN2 w/o [19]	11.70	0.44	-
StyleGAN2 w/ ADA [56]	13.72	0.36	-
StyleGAN2 + DiffAug [56]	12.97	0.39	-

Table 3. Results on STL-10 64 × 64



Figure 5. Qualitative results of CIFAR-10 32 × 32



Figure 6. Qualitative results of CelebA-HQ 256 × 256

dataset, our model has a comparable number of parameters but requires less computing FLOPs and memory usage compared with DDGAN [49]. Following DDGAN [49], we use 2-sampling steps for CelebA-HQ (256, 512 & 1024) and 4-sampling steps for CIFAR-10 (32), STL-10 (64), and LSUN-Church (256) in both training and testing.

Our speed gain does not only come from the proposed framework but also from corresponding network configurations. As the wavelet input of our model is 4× smaller than the original dimensions, we need to provide suitable network configurations. On CIFAR-10, we use only 3 layers

Model	FID↓	Recall↑	Time (s)↓
Ours	6.55	0.35	0.60
Ours + W-Generator	5.94	0.37	0.79
DDGAN [49]	7.64	0.36	1.73
Score SDE [43]	7.23	-	-
NVAE [45]	29.7	-	-
VAEBM [48]	20.4	-	-
PGGAN [18]	8.03	-	-
VQ-GAN [6]	10.2	-	-

Table 4. Results on CelebA-HQ 256×256 . “Ours” denotes employing wavelet-based diffusion scheme without using the wavelet-embedded generator.

for both the generator and the discriminator instead of 4 as in DDGAN. On CelebA-HQ (256) and LSUN-Church, we use 5 instead of 6 layers. On other datasets, we use the same configurations as in DDGAN, with 6 layers for CelebA-HQ (512 & 1024) and 4 layers for STL-10.

5.2. Experimental results

CIFAR-10 As shown in Tab. 2, we have greatly improved inference time by requiring only $0.08(s)$ with our Wavelet based diffusion process, which is $2.5\times$ faster than DDGAN. This gives us a real-time performance along with StyleGAN2 methods while exceeding other diffusion models by a wide margin in terms of sampling speed.

STL-10 In Tab. 3, we not only achieve a better FID score at 12.93 but also gain faster sampling time at $0.38(s)$. We further analyze the convergence speed of our approach and DDGAN in Fig. 8a. Our approach offers a faster convergence than the baseline, especially in the early epochs. Generated samples of DDGAN can not recover objects’ overall shape and structure during the first 400 epochs. Besides, we provide sample images generated by our model in Fig. 8b.

CelebA-HQ At resolution 256×256 , we outperform some notable diffusion and GAN baselines with FID at 5.94 and Recall at 0.37 while achieving more than $2\times$ faster than DDGAN, as reported in Tab. 4. On high-resolution CelebA-HQ (512) in Tab. 5, our model is remarkably better than the DDGAN counterpart for both image quality (6.40 vs. 8.43) and sampling time (0.59 vs. 1.49). We also obtain a higher recall at 0.35 on high-resolution CelebA-HQ (512). Our further benchmarking on CelebA-HQ 1024 yielded a competitive FID score of 5.98, comparable to many high-res GAN baselines (StyleGAN - 5.06, PGGAN - 7.3). Notably, its inference time (0.59s for 25 samples) remains comparable to CelebA-HQ 512, thanks to our unchanged network configurations with only two modifications: removal of attention layers and use of patch size 2. For qualitative results, we provide the generated samples



Figure 7. Qualitative results of CelebA-HQ 512×512

Model	FID↓	Recall↑	Time (s)↓
Ours + W-Generator	6.40	0.35	0.59
DDGAN [49]	8.43	0.33	1.49

Table 5. Results on CelebA-HQ 512×512

Model	FID↓	Recall↑	Time (s)↓
Ours + W-Generator	5.06	0.40	1.54
DDGAN [49]	5.25	-	3.42
DDPM [13]	7.89	-	-
ImageBART [5]	7.32	-	-
PGGAN [18]	6.42	-	-
StyleGAN [20]	4.21	-	-
StyleGAN2 [19]	3.86	0.36	-

Table 6. Results on LSUN-Church 256×256 .

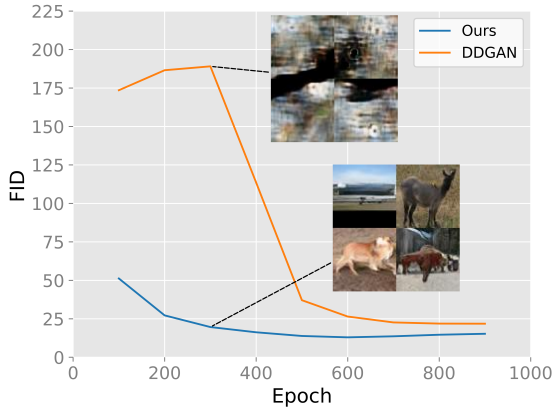
for the CelebA-HQ (256) and CelebA-HQ (512) datasets in Fig. 6 and Fig. 7, respectively.

LSUN-Church In Tab. 6, we obtain a superior image quality at 5.06 in comparison with other diffusion models while achieving comparable results with GAN counterparts. Notably, our model offers $2\times$ faster inference time compared with DDGAN while exceeding StyleGAN2 [19] in terms of sample diversity at 0.40. For qualitative results, randomly generated samples are shown at Fig. 9.

5.3. Ablation studies

Reconstruction term In this section, we verify the contribution of reconstruction loss Eq. (5) to the model performance on CelebA-HQ 256×256 . The FID score experiences a remarkable reduction of around 0.6 points to 5.94 when employing the term, confirming its usefulness in improving model’s quality.

Wavelet-embedded networks We validate the contribution of each individual component of our proposed wavelet-based generator on CelebA-HQ 256×256 in Tab. 7, where the full model includes residual connections, upsampling



(a) Training convergence comparisons



(b) Generated samples

Figure 8. Training convergence comparisons (a) and qualitative results (b) on STL-10 64×64 . Our method not only converges faster but also acquires better FID scores than DDGAN [49] across different epochs.

Model	FID↓	Time (s)↓
w/o residual	6.25	0.78
w/o up & down	6.23	0.61
w/o bottleneck	6.18	0.78
full model	5.94	0.79

Table 7. Ablation study of wavelet generator on CelebA-HQ 256×256 . Each setting is trained for 500 epochs.

	CIFAR-10	STL-10	CelebA(256)	CelebA(512)	CelebA(1024)	Church
Resolution	32	64	256	512	1024	256
#time steps	4	4	2	2	2	4
Time (s)	0.07	0.12	0.08	0.1	0.12	0.16

Table 8. Running time when using our full model to generate a single image on each benchmark set.

and downsampling blocks, and bottleneck blocks. As can be seen, each component has a positive effect on the model’s performance. By applying all three proposed components, our method achieves the best performance at 5.94, especially the bottleneck block is the least important component in the design of generator. Still, the performance gain comes along with a small cost in running speed.

5.4. Running time when generating a single image

We further demonstrate the superior speed of our models on single images, as expected in real-life applications. In Tab. 8, we present their time and key parameters. Our wavelet diffusion models can produce images up to 1024×1024 in a mere 0.1s, which is the first time for a diffusion model to achieve such almost real-time performance.

5.5. Why ours converges faster and more stably?

We believe that our model benefits from the frequency decomposition of wavelet transformation. Instead of learning from an entanglement of coarse and detailed informa-



Figure 9. Qualitative results of LSUN-Church 256×256 .

tion, our method separates them for efficient training and at multi scales in the feature space. First, it studies the low-frequency subbands easier due to lower spatial dimensions. Second, it quickly learns the sparse and repetitive high-frequency components, focusing on distinctive details.

6. Conclusions

This paper introduced a novel wavelet-based diffusion scheme that demonstrates superior performance on both image fidelity and sampling speed. By incorporating wavelet transformations to both image and feature space, our method can achieve the state-of-the-art running speed for a diffusion model, closing the gap with StyleGAN models [19, 20, 56] while obtaining a comparable image generation quality to StyleGAN2 and other diffusion models. Besides, our method offers a faster convergence than the baseline DDGAN [49], confirming the efficiency of our proposed framework. With these initial results, we hope our approach can facilitate future studies on real-time and high-fidelity diffusion models.

References

- [1] CompVis/stable-diffusion · Hugging Face. [Online]. [1](#)
- [2] Charles K Chui. *An introduction to wavelets*, volume 1. Academic press, 1992. [2](#)
- [3] Xin Deng, Ren Yang, Mai Xu, and Pier Luigi Dragotti. 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*, pages 3076–3085, 2019. [2](#)
- [4] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. [1](#), [2](#)
- [5] Patrick Esser, Robin Rombach, Andreas Blattmann, and Bjorn Ommer. Imagebart: Bidirectional context with multinomial diffusion for autoregressive image synthesis. *Advances in Neural Information Processing Systems*, 34:3518–3532, 2021. [2](#), [7](#)
- [6] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021. [7](#)
- [7] Rinon Gal, Dana Cohen Hochberg, Amit Bermano, and Daniel Cohen-Or. Swagan: A style-based wavelet-driven generative model. *ACM Transactions on Graphics (TOG)*, 40(4):1–11, 2021. [2](#)
- [8] Ruiqi Gao, Yang Song, Ben Poole, Ying Nian Wu, and Diederik P Kingma. Learning energy-based models by diffusion recovery likelihood. In *International Conference on Learning Representations*, 2021. [6](#)
- [9] Yue Gao, Fangyun Wei, Jianmin Bao, Shuyang Gu, Dong Chen, Fang Wen, and Zhouhui Lian. High-fidelity and arbitrary face editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16115–16124, 2021. [2](#)
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, 2014. [2](#)
- [11] Florentin Guth, Simon Coste, Valentin De Bortoli, and Stephane Mallat. Wavelet score-based generative modeling. *arXiv preprint arXiv:2208.05003*, 2022. [2](#)
- [12] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 2017. [6](#)
- [13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in neural information processing systems*, 2020. [1](#), [2](#), [6](#), [7](#)
- [14] Chin-Wei Huang, Jae Hyun Lim, and Aaron C Courville. A variational perspective on diffusion-based generative models and score matching. *Advances in Neural Information Processing Systems*, 34:22863–22876, 2021. [1](#)
- [15] Huaibo Huang, Ran He, Zhenan Sun, and Tieniu Tan. Wavelet domain generative adversarial network for multi-scale face hallucination. *International Journal of Computer Vision*, 127(6):763–784, 2019. [2](#)
- [16] Liming Jiang, Bo Dai, Wayne Wu, and Chen Change Loy. Focal frequency loss for image reconstruction and synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13919–13929, 2021. [2](#)
- [17] Alexia Jolicoeur-Martineau, Ke Li, Rémi Piché-Taillefer, Tal Kachman, and Ioannis Mitliagkas. Gotta go fast when generating data with score-based models. *arXiv preprint arXiv:2105.14080*, 2021. [6](#)
- [18] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018. [7](#)
- [19] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *Advances in neural information processing systems*, 2020. [1](#), [6](#), [7](#), [8](#)
- [20] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. [1](#), [2](#), [7](#), [8](#)
- [21] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2020. [6](#)
- [22] Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *Advances in neural information processing systems*, 34:21696–21707, 2021. [1](#)
- [23] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in neural information processing systems*, 2018. [6](#)
- [24] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *The International Conference on Learning Representations*, 2014. [2](#)
- [25] Zhifeng Kong and Wei Ping. On fast sampling of diffusion probabilistic models. In *ICML Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models*, 2021. [1](#), [6](#)
- [26] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. *Advances in Neural Information Processing Systems*, 32, 2019. [6](#)
- [27] Jin Li, Wanyun Li, Zichen Xu, Yuhao Wang, and Qiegen Liu. Wavelet transform-assisted adaptive generative modeling for colorization. *IEEE Transactions on Multimedia*, 2022. [2](#)
- [28] Qiufu Li, Linlin Shen, Sheng Guo, and Zhihui Lai. Wavelet integrated cnns for noise-robust image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7245–7254, 2020. [2](#)
- [29] Pengju Liu, Hongzhi Zhang, Wei Lian, and Wangmeng Zuo. Multi-level wavelet convolutional neural networks. *IEEE Access*, 7:74973–74985, 2019. [2](#)
- [30] Eric Luhman and Troy Luhman. Knowledge distillation in iterative generative models for improved sampling speed. *arXiv preprint arXiv:2101.02388*, 2021. [2](#), [6](#)
- [31] Hengyuan Ma, Li Zhang, Xi Tian Zhu, and Jianfeng Feng. Accelerating score-based generative models with precondi-

- tioned diffusion sampling. In *European Conference on Computer Vision*, pages 1–16. Springer, 2022. 2
- [32] Stephane G Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE transactions on pattern analysis and machine intelligence*, 11(7):674–693, 1989. 2
- [33] Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. *International Conference on Machine Learning*, 2016. 6
- [34] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 1, 2
- [35] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021. 2
- [36] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 1, 2
- [37] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. 1, 2
- [38] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, 2015. 2
- [39] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. 1, 2, 6
- [40] Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in Neural Information Processing Systems*, 34:1415–1428, 2021. 1, 2
- [41] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in neural information processing systems*, 2019. 2, 6
- [42] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020. 2
- [43] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. 2, 4, 5, 6, 7
- [44] David Taubman and Michael Marcellin. *JPEG2000 Image Compression Fundamentals, Standards and Practice: Image Compression Fundamentals, Standards and Practice*, volume 642. Springer Science & Business Media, 2012. 2
- [45] Arash Vahdat and Jan Kautz. NVAE: A deep hierarchical variational autoencoder. In *Advances in neural information processing systems*, 2020. 2, 6, 7
- [46] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. In *Advances in neural information processing systems*, 2021. 2
- [47] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011. 2
- [48] Zhisheng Xiao, Karsten Kreis, Jan Kautz, and Arash Vahdat. Vaebm: A symbiosis between variational autoencoders and energy-based models. In *International Conference on Learning Representations*, 2021. 6, 7
- [49] Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with denoising diffusion gans. In *International Conference on Learning Representations*, 2022. 1, 2, 3, 4, 6, 7, 8
- [50] Mengping Yang, Zhe Wang, Ziqiu Chi, and Wenyi Feng. Wavegan: An frequency-aware gan for high-fidelity few-shot image generation. In *ECCV*, 2022. 2
- [51] Mengping Yang, Zhe Wang, Ziqiu Chi, and Yanbing Zhang. FreGAN: Exploiting frequency components for training GANs under limited data. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. 2
- [52] Ting Yao, Yingwei Pan, Yehao Li, Chong-Wah Ngo, and Tao Mei. Wave-vit: Unifying wavelet and transformers for visual representation learning. In *European Conference on Computer Vision*, pages 328–345. Springer, 2022. 2
- [53] Jaejun Yoo, Youngjung Uh, Sanghyuk Chun, Byeongkyu Kang, and Jung-Woo Ha. Photorealistic style transfer via wavelet transforms. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9036–9045, 2019. 2
- [54] Yingchen Yu, Fangneng Zhan, Shijian Lu, Jianxiong Pan, Feiying Ma, Xuansong Xie, and Chunyan Miao. Wavefill: A wavelet-based generation network for image inpainting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14114–14123, 2021. 2
- [55] Bowen Zhang, Shuyang Gu, Bo Zhang, Jianmin Bao, Dong Chen, Fang Wen, Yong Wang, and Baining Guo. Styleswin: Transformer-based gan for high-resolution image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11304–11314, 2022. 2
- [56] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable augmentation for data-efficient gan training. In *Advances in neural information processing systems*, 2020. 1, 6, 8