

Beta Sampling is All You Need: Efficient Image Generation Strategy for Diffusion Models using Stepwise Spectral Analysis

Haeil Lee* Hansang Lee* Seoyeon Gye Junmo Kim

School of Electrical Engineering, KAIST, South Korea

{haeil.lee, hansanglee, sawyun, junmo.kim}@kaist.ac.kr

Abstract

Generative diffusion models have emerged as a powerful tool for high-quality image synthesis, yet their iterative nature demands significant computational resources. This paper proposes an efficient time step sampling method based on an image spectral analysis of the diffusion process, aimed at optimizing the denoising process. Instead of the traditional uniform distribution-based time step sampling, we introduce a Beta distribution-like sampling technique that prioritizes critical steps in the early and late stages of the process. Our hypothesis is that certain steps exhibit significant changes in image content, while others contribute minimally. We validated our approach using Fourier transforms to measure frequency response changes at each step, revealing substantial low-frequency changes early on and high-frequency adjustments later. Experiments with ADM and Stable Diffusion demonstrated that our Beta Sampling method consistently outperforms uniform sampling, achieving better FID and IS scores, and offers competitive efficiency relative to state-of-the-art methods like AutoDiffusion. This work provides a practical framework for enhancing diffusion model efficiency by focusing computational resources on the most impactful steps, with potential for further optimization and broader application.

1. Introduction

Generative diffusion models have emerged as a powerful tool for high-quality image generation, producing results that rival or surpass traditional generative adversarial networks (GANs) [7]. These models iteratively refine images through a diffusion process, where noise is progressively added and then removed, ultimately generating realistic images from random noise [11]. However, this iterative nature comes with a significant computational cost, necessitating numerous time steps to achieve high-quality outputs. Effi-




Figure 1. **An Overview.** (a) We analyzed the Fourier transform of images generated at each time step during the denoising process and found that the changes in low- and high-frequency components are concentrated in the early and the later stages, respectively. (b) Based on this, we propose a Beta distribution-like sampling method that focuses on key stages with significant frequency changes. (c) Experiments show our method generates higher quality images at lower steps compared to uniform sampling.

*These authors contributed equally to this work.

cient methods are therefore crucial to reduce computational burden while maintaining the quality of generated images.

Previous efforts to improve the efficiency of diffusion models have focused on reducing the number of sampling steps in denoising process to improve the efficiency of diffusion models. Some approaches model the sampling process as ordinary differential equations (ODEs), which enables fewer steps using first-order and higher-order solvers [16, 18, 30]. Another effective strategy is to use knowledge distillation techniques repeatedly to condense multiple steps into a single step [3, 19, 26]. This enables the generation of high-quality images in 10 steps or fewer. Recently, new approaches have been proposed based on the idea that the sampling process of diffusion models plays a different role at each step [2, 23, 35]. Some works aim to reduce the number of sampling steps by selecting only the optimal step for particular diffusion models [8, 14, 31], while others aim to increase efficiency by using multiple smaller but step-specific models for each step [13, 14, 34, 35].

In this paper, we propose a novel Beta Sampling method for improving the efficiency of generative diffusion models. By conducting an image spectral analysis of the diffusion process using Fourier transform, we identified that significant changes in image content occur predominantly at the early and late stages of the denoising process. Based on this insight, we introduce a Beta distribution-like time step sampling method that emphasizes critical steps in the early and late stages of the process. Our key hypothesis is that certain steps in the diffusion process exhibit significant changes in image content, while others contribute minimally. By focusing on these impactful time steps, we aim to enhance the efficiency of image generation without compromising quality, in contrast to traditional uniform sampling. We validate our approach through experiments with ADM and Stable Diffusion models, demonstrating that our Beta Sampling method consistently achieves better FID and IS scores compared to uniform sampling and exhibits competitive efficiency against the state-of-the-art AutoDiffusion method. Our findings highlight the potential for substantial computational savings and quality improvements in image generation tasks.

Our contributions are as follows.

- We provide a spectral analysis of the diffusion process through Fourier transform, identifying that the significant changes in frequency component are concentrated in the early and later stages of the denoising process.
- Based on our analysis, we introduce a Beta distribution-like sampling method that prioritizes steps with substantial changes in low and high-frequency components.
- We demonstrate through experiments with ADM and Stable Diffusion that our proposed sampling method

consistently achieves improved FID and IS scores compared to uniform sampling, and offers competitive efficiency against state-of-the-art methods like AutoDiffusion in terms of computational complexity.

2. Related works

2.1. Denoising process of diffusion models

There is an active research effort to analyze and improve the denoising process of diffusion models from various aspects. So far, most analyses agree that the diffusion model’s denoising process proceeds in a coarse-to-fine manner, i.e., it generates the overall structure of the image by focusing on low-frequency components at the beginning, and then completes the details through changes in high-frequency components at the end, resulting in progressively better images.

Several efforts have been focused on analyzing the denoising process of the diffusion models. First, Choi *et al.* [4] analyzed the process of adding noise to an image using LPIPS distance, with the idea that the diffusion model has a different pretask for each step, and proposed that the denoising process of the diffusion model can be divided into three stages: 1) creating coarse features 2) generating perceptually rich contents 3) and removing remaining noise.

To analyze how the latent structure of the diffusion model varies with the diffusion timestep, Park *et al.* [21] identified the frequency domain of the local latent basis by power spectral density (PSD) analysis, and confirmed the shift from low-frequency to high-frequency as the denoising process progresses. By analyzing the frequency domain of the image obtained at each step, we also confirmed that the low-frequency part is restored at the beginning and the high-frequency detail is restored at the end [13, 35]. Furthermore, Li *et al.* [14] argued that the difficulty and importance of each of these different steps is different, and that efficient sampling can be achieved by finding an optimal time step.

More recently, researchers have tried to visually interpret the denoising process through a text to image diffusion model [20]. This paper investigated the spatial recovery level at each timestep in the denoising process and showed that the focal region of the model changes from semantic information to fine-grained regions.

2.2. Frequency Analysis on Diffusion Models

Frequency analysis has long been a fundamental tool for image processing and analysis in the field of computer vision. It enables the separation of high and low-frequency components of an image. The Fourier transform has been widely used for frequency analysis in identifying major patterns in images, removing noise, and compressing images.

Research has shown that deep neural networks initially adapt to low-frequency signals, while learning high-frequency details more gradually [22, 32, 33]. This phe-

nomenon, known as *spectral bias*, has also been observed in deep generative models such as GANs [9, 12, 28]. A similar spatial frequency bias has been found in diffusion models. Choi *et al.* [4] suggests that the denoising process of diffusion models shows three distinguishable phases. The first phase captures coarse attributes, while the subsequent phases progressively incorporate finer details. Other studies have noted that throughout the progression of the denoising process, low-frequency components are preserved, while high-frequency components undergo rapid changes [29] or the ratio of high-frequency increases as the denoising progresses [21].

This observation indicates that the denoising task of diffusion models varies at each time step, leading to distinct characteristics in spectral analysis at each step. Recent studies have utilized this fact to enhance the quality of generated samples [29] or computational efficiency [13, 35]. For example, some studies give more weight to a specific frequency at each step [29, 35], while Lee *et al.* [13] utilizes different smaller models at each step, each specialized for a particular range of frequencies.

3. Spectral Analysis of Denoising Processes in Diffusion Models

The denoising process in diffusion models has been demonstrated to encompass several distinct stages, each characterized by unique model behaviors [4, 5]. This variability in behavior across time steps suggests that the importance of each step in the denoising process may not be uniform. To investigate the relative importance of each step in the denoising process, we conducted a comprehensive spectral analysis of the denoising procedure in pre-trained diffusion models. This analytical approach allows us to inspect which parts of the denoising process contribute most significantly to meaningful changes in the generated image. By decomposing the process into its spectral components, we aim to provide insights into the differential contributions of each denoising step to the final output quality.

Diffusion Models. To investigate the frequency characteristics of diffusion models, we conducted a comprehensive analysis on two prominent models: ADM-G [7] and Stable Diffusion [23]. The ADM-G model was trained on ImageNet 64×64 [6], while Stable Diffusion was trained on the LAION-5B [27] dataset.

For the ADM-G model, we observed the denoising process for a total of 10,000 image generations. Specifically, we generated 10 images for each class in the ImageNet dataset, ensuring a broad representation across all categories. In the case of Stable Diffusion, we randomly selected 1,000 captions from the validation set of COCO [15] dataset to serve as prompts for image generation.

Spectral Analysis. To conduct spectral analysis of the denoising process, we applied a 2D Fourier transform to the




Figure 2. Spectral analysis of denoising process in ADM-G [7] and Stable Diffusion (SD) [23]. The trend in the changes of high-frequency and low-frequency components during the denoising steps demonstrates that the core of the diffusion model’s image denoising process lies in the early and late stages.

images generated at each step of the diffusion model. We then calculated the relative log magnitude of each frequency component. For an image x , we denote this result as $RLM(x)$. When it comes to Stable Diffusion, a latent diffusion model, we utilized a pre-trained variational autoencoder to transform the latent at each step into the image domain for frequency component analysis. The visualized results of $RLM(x_t)$ across all time step t are depicted in Figs. 2a and 2b. As anticipated, the frequency components of the image x_{1000} at $t = 1000$, which represents pure Gaussian noise, exhibited a flat distribution. However, as the denoising process progressed, we observed changes in the frequency components of the images. To quantify these changes, we defined the frequency change at a specific step t as the difference between the frequency components of consecutive steps: $RLM(x_t) - RLM(x_{t+1})$. These differ-

ences by time step t are shown in Figs. 2c and 2d. This approach allows us to track the evolution of frequency components throughout the denoising process. By examining these spectral changes, we can gain insights into how different frequency bands are affected at various stages of image generation.

Major Changes Occur in the Early and Late Stages. The visualization of our spectral analysis is presented in Fig. 2. Consistent with previous works, we observe that the early stages of the denoising process (closer to $t = 1000$) primarily form coarse, low-frequency components, as evident in the lower right area of Figs. 2c and 2d. The red regions indicating an increase in the heatmap demonstrate that the low-frequency components are formed predominantly in the early stages. Conversely, the later stages (approaching $t = 0$) are dominated by changes in fine, high-frequency components, as shown in the upper left area of these graphs. To further clarify this phenomenon, we isolated the increasing low-frequency components in the early stages and the decreasing high-frequency components in the later stages, and show their respective averages in Figs. 2e and 2f. This separation demonstrates that low-frequency elements undergo significant changes early in the process, while high-frequency elements experience substantial modifications in the later stages. Notably, the intermediate stages show relatively minor changes in image information from a frequency perspective.

4. Beta Sampling for Step Reduction

Inspired by spectral analysis, we propose a novel sampling technique for step reduction in diffusion models. Our previous investigations revealed a significant phenomenon in the image-denoising process of diffusion models: the most critical and substantial changes are predominantly concentrated in the initial and final stages of the model’s operation. This observation suggests that when selecting a reduced number of steps compared to the total time steps used during training, there is a compelling need to develop a method that can allocate more steps to the early and late stages of the process. Consequently, we propose a novel sampling strategy that allocates a higher density of steps to the initial and final stages of the denoising process based on Beta distribution. This method aims to concentrate computational resources on the periods where the most crucial transformations occur.

To sample Beta distribution-like time steps in the denoising process, our method leverages the Probability Integral Transform (PIT), a well-established method for generating samples from a target distribution using uniform random variables [1]. We begin with a set of uniformly sampled time steps t_i , where $i = 1, 2, \dots, N$, and $t_i \in [0, T - 1]$. Then we normalize the time steps: $t'_i = t_i/T$. These time steps are initially equidistant, following a Uniform distribu-

tion $U(0, 1)$. We choose a Beta distribution $B(\alpha, \beta)$ as our target distribution for the time steps. The Beta distribution is selected for its flexibility in modeling various shapes over a finite interval $[0, 1]$, which aligns well with the normalized time step range. The CDF of the beta distribution, denoted as $F(x; \alpha, \beta)$, is computed. We apply the PIT to transform our uniform samples into samples from the Beta distribution. For each uniformly sampled time step t'_i , we compute:

$$t_i^B = F^{-1}(t'_i; \alpha, \beta) \quad (1)$$

where F^{-1} is the inverse CDF (or Percent Point Function, PPF) of the Beta distribution. If necessary, we rescale the transformed time steps t_i^B to ensure they span the full range of the diffusion process. Our method does not involve random sampling from the Beta distribution. Instead, it employs distribution equalization through the PIT to achieve fixed-point sampling. This approach ensures a deterministic and consistent allocation of time steps that adheres to the desired Beta distribution, offering a more controlled and reproducible sampling process compared to random sampling methods.

By carefully selecting the α and β parameters of the Beta distribution, we can precisely control the density of time steps at different stages of the denoising process. This adaptive sampling strategy enables a more efficient allocation of computational resources, focusing on the most critical phases of the diffusion process while maintaining a smooth transition throughout the entire range.

5. Experiments

5.1. Experimental Setup

Diffusion Models. To validate the effectiveness of our method, we conducted experiments using representative pre-trained diffusion models without retraining or fine-tuning. As in Sec. 3, we employed the ADM-G [7] trained on ImageNet 64×64 and the Stable Diffusion [23] for our experiments. For ADM-G, we utilized the DDIM [30] sampler, while for Stable Diffusion, we employed PLMS [16].

Metrics. We used Fréchet Inception Distance (FID) [10] and Inception Score (IS) [25] as our evaluation metrics, consistent with most previous works in this field. FID measures the quality difference between generated images and target images, with lower values indicating greater similarity. Higher IS values suggest that the generated images exhibit high quality and diversity.

Baselines. We conducted a comparative analysis with uniform sampling and AutoDiffusion [14]. uniform sampling is the most commonly used technique when decreasing the number of time steps. AutoDiffusion, introduced to address the limitations of uniform sampling, employs an evolutionary search algorithm to identify optimal time steps that minimize the FID. We implemented AutoDiffusion as described

in the original paper, conducting a 10 epoch evolutionary search and selecting the optimal time step configuration.

Implementation Details. In our experiments with ADM-G, we generated 50,000 samples using random class guidance. For Stable Diffusion, we produced 10,000 samples using captions from the COCO dataset. We then calculated FID and IS scores for these generated samples to assess the performance of our method.

5.2. Results

Beta Sampling Outperforms Uniform Sampling and Competes with AutoDiffusion. Table 1 and Table 2 present a comparison of FID and IS scores for generated images using different sampling methods (Beta Sampling, AutoDiffusion, and uniform sampling) in ADM-G and Stable Diffusion (SD) models, respectively. In both ADM and SD experiments, the three sampling methods show similar trends. At 50 and 100 steps, AutoDiffusion’s genetic algorithm search consumes significant time, and since the step count is high enough to approach a solution-like image, the FID and IS scores for both uniform and Beta Sampling become similar. [24] However, at 10, 15, and 20 steps, Beta Sampling shows clear improvements over uniform sampling. In ADM-G, Beta Sampling even outperforms AutoDiffusion, while in Stable Diffusion, AutoDiffusion performs slightly better. Nevertheless, considering AutoDiffusion’s long search time, Beta Sampling proves more efficient. In very low step settings of 4 and 6, AutoDiffusion yields better FID and IS scores than Beta Sampling. However, while uniform sampling shows a significant performance drop, Beta Sampling exhibits only marginal performance degradation. In summary, at higher step counts (50, 100), the difference between sampling methods is negligible, but at practical step counts (10, 15, 20), Beta Sampling excels in both quality and efficiency, whereas in near one-shot environments (4, 6 steps), AutoDiffusion performs better.

Figure 3 shows generated examples from identical initial noise. At 4 steps, Beta Sampling produces blurrier results compared to AutoDiffusion. However, at 6 and 10 steps, both Beta Sampling and AutoDiffusion yield clearer images than uniform sampling. Given the computationally intensive search process of AutoDiffusion, Beta Sampling emerges as a more efficient alternative. At 15 and 20 steps, all strategies produce clear images, with both uniform and Beta Sampling offering efficient options. Overall, except for 4 steps, Beta Sampling demonstrates competitive image quality without the additional time or computational burdens.

The quality of examples generated with Stable Diffusion can be observed in Fig. 4. At 4 and 6 steps, uniform sampling exhibits structural defects and poor color representation. Beta Sampling mitigates these issues but remains inferior to AutoDiffusion, which employs a more extensive search process for large models. At 10 steps, Beta Sampling

Steps	Sampling Strategies	FID (↓)	IS (↑)
4	Uniform	138.66	7.09
	AutoDiffusion	17.86	34.88
	Ours (Beta)	31.64	25.27
6	Uniform	23.71	31.53
	AutoDiffusion	11.17	43.47
	Ours (Beta)	13.12	41.30
10	Uniform	8.86	46.50
	AutoDiffusion	6.24	57.85
	Ours (Beta)	6.13	58.15
15	Uniform	5.38	54.82
	AutoDiffusion	4.92	64.03
	Ours (Beta)	4.43	66.28
20	Uniform	4.35	58.41
	AutoDiffusion	3.93	68.05
	Ours (Beta)	3.93	71.42
50	Uniform	3.20	69.32
	Ours (Beta)	3.17	73.85
100	Uniform	3.19	71.50
	Ours (Beta)	3.19	73.85

Table 1. FID (↓) and IS (↑) scores for ADM-G [7] on ImageNet 64×64 with various number of time steps and sampling strategies. For Beta Sampling, we set $\alpha = \beta = 0.5$ for steps from 4 to 20, and $\alpha = \beta = 0.9$ for steps 50 and 100.

produces samples comparable in quality to AutoDiffusion. Interestingly, at 20 steps, AutoDiffusion occasionally introduces undesired artifacts into the examples. While there is considerable fluctuation depending on the prompt, Beta Sampling demonstrates competitive efficiency in all cases except at 4 steps.

Ablation Study (1): Symmetric Beta Sampling Performs Better. Tab. 3 and Fig. 5 present the FID and IS evaluations and sample images generated using these three Beta Sampling with different parameters. For Beta(2,5) with $\alpha < \beta$, which concentrates sampling in the later stages, the generated images lacked the necessary low-frequency generation, resulting in noisy and incomprehensible images, as evidenced by very poor FID and IS scores. This demonstrates that the early steps are crucial for forming low-frequency components, which need to be shaped before the high-frequency components are generated. Beta(5,2) with $\alpha > \beta$ focused on early stages, producing images with good low-frequency content but lacking high-frequency details, resulting in blurred edges and inferior FID and IS scores compared to uniform sampling. In contrast, our proposed Beta(0.5,0.5) with $\alpha = \beta$ balanced early and late stage sampling, preserving both content and detail, and achieving the best FID and IS scores. These results confirm our hypothesis that the early and late steps of the denoising process




Figure 3. Examples generated by ADM-G [7] on ImageNet 64×64 with various sampling strategies.



Figure 4. Examples generated by Stable Diffusion [23] with various sampling strategies. The text prompts used for generation are "A man who is wearing a suit and tie" and "Two large elephants are standing beside each other".

are crucial for enhancing low-frequency and high-frequency components, respectively.

Ablation Study (2): Hyperparameter Analysis. Figure 6 illustrates the changes in FID and IS performance as a function of the hyperparameter $\alpha = \beta$. When $\alpha = \beta = 1$, the distribution is equivalent to a uniform distribution. The observed changes in FID and IS performance as the hyperparameter decreases from the uniform distribution demonstrate the effectiveness of Beta Sampling. Specifically, we set $\alpha = \beta = 0.5$ for ADM-G, and $\alpha = \beta = 0.6$ for Stable Diffusion except for the 50- and 100-step scenarios. For this relatively large NFE region, we identified the optimal parameters individually.

Ablation Study (3): Performance by Guidance Scale. Figure 7 shows FID scores for various guidance scales with and without Beta Sampling on Stable Diffusion. Across multiple steps, the Beta Distribution consistently reduces FID scores at all scales, demonstrating its stable performance.




Figure 5. Examples generated by ADM-G [7] on ImageNet 64×64 using Beta Sampling with various distribution shapes. The frequency column provides a relative frequency analysis of the 20-step generation, highlighting differences in the process.

Steps	Sampling Strategies	FID ($\downarrow$)	IS ($\uparrow$)
4	Uniform	38.22	16.06
	AutoDiffusion	20.18	23.10
	Ours (Beta)	34.71	16.95
6	Uniform	32.40	17.99
	AutoDiffusion	17.57	23.83
	Ours (Beta)	22.48	21.83
10	Uniform	19.16	22.04
	AutoDiffusion	13.20	26.52
	Ours (Beta)	16.45	25.51
20	Uniform	14.57	27.40
	AutoDiffusion	13.08	25.21
	Ours (Beta)	15.10	27.46
50	Uniform	15.92	27.76
	Ours (Beta)	15.89	27.85
100	Uniform	15.66	27.76
	Ours (Beta)	15.65	27.78

Table 2. FID ($\downarrow$) and IS ($\uparrow$) scores for Stable Diffusion [23] with various number of time steps and sampling strategies. For Beta Sampling, we set $\alpha = \beta = 0.6$ for steps from 4 to 20, $\alpha = \beta = 0.7$ for 50 steps and $\alpha = \beta = 0.9$ for 100 steps.

Steps	Sampling Strategies	FID ($\downarrow$)	IS ($\uparrow$)
4	Beta (0.5,0.5)	31.64	25.27
	Beta (5,2)	52.65	16.89
	Beta (2,5)	417.03	1.16
6	Beta (0.5,0.5)	13.12	41.30
	Beta (5,2)	31.74	25.18
	Beta (2,5)	419.33	1.18
10	Beta (0.5,0.5)	6.13	58.15
	Beta (5,2)	19.82	34.00
	Beta (2,5)	420.46	1.21
15	Beta (0.5,0.5)	4.43	66.28
	Beta (5,2)	15.17	39.53
	Beta (2,5)	418.37	1.28
20	Beta (0.5,0.5)	3.93	71.42
	Beta (5,2)	13.21	42.31
	Beta (2,5)	416.31	1.32

Table 3. FID ($\downarrow$) and IS ($\uparrow$) scores for ADM-G [7] on ImageNet 64×64 with Beta Sampling of various distribution shape and their corresponding sampled step distributions.

6. Discussion

Beta Sampling is Highly Efficient. One key feature of Beta Sampling is that it is both training-free and highly efficient. Various denoising speed-up techniques typically rely on either distillation from existing models or require additional training. For instance, in the case of InstaFlow [17],




Figure 6. FID ($\downarrow$) and IS ($\uparrow$) scores for ADM-G [7] on ImageNet 64×64 with Beta Sampling $Beta(\alpha, \beta)$ of various hyperparameter $\alpha = \beta$ and various time steps.




Figure 7. FID ($\downarrow$) scores for Stable Diffusion [23] with Beta Sampling across various guidance scale.

Method	Steps	Search	Cost	FID
AutoDiffusion	10	Genetic Algo.	31h+	6.24
Ours	10	Param. Search	19m	6.10

Table 4. Efficiency comparison with AutoDiffusion




Figure 8. FID scores measured with 2k samples and 50k samples across $\alpha = \beta$ of Beta distribution.

the process takes 199 days on an A100. In contrast, both AutoDiffusion and our method are training-free, search-based approaches, but differ significantly in efficiency. Table 4 demonstrates the superior efficiency of Beta Sampling compared to AutoDiffusion. For the ADM-G 10-step setting, AutoDiffusion requires a minimum of 31 hours for the search, while our method takes only 19 minutes with a sin-




Figure 9. A cumulative histogram of occurrence number of time steps sampled by AutoDiffusion [14] on ADM-G [7] with various number of time steps. To enhance the visibility of the histogram, the time steps were sampled 150 times for 4 and 6 steps, and 100 times for 10 and 15 steps, respectively.

gle A100. As shown in Fig. 8, by sampling just 2k images for parameter search, we were able to effectively determine the optimal α . This process includes denoising sample generation, and the gap in efficiency becomes even more pronounced for high-resolution models or when a higher number of steps is involved.

AutoDiffusion Behaves Like Beta Sampling. Our experiments revealed that Beta Sampling achieves results comparable to AutoDiffusion when using 10 or more time steps, but underperforms with extremely small time steps, such as 4 or 6 steps. To better understand these differences, we analyzed the distribution of time steps extracted from AutoDiffusion by examining the cumulative histogram of step frequencies, as shown in Fig. 9. For enhanced visibility, we repeated step sampling 150 times for 4 and 6 steps, and 50 times for 10 and 15 steps. The histogram revealed that for 4 and 6 steps, the distribution was nearly uniform, whereas for 10 and 15 steps, the distribution resembled a Beta distribution. This indicates that AutoDiffusion effectively operates similarly to Beta Sampling when using 10 or more steps. Consequently, for scenarios requiring 10 or more time steps, Beta Sampling emerges as a more efficient choice compared to AutoDiffusion, which involves a longer searching time.

ADM-G vs. Stable Diffusion. We validated the effectiveness of Beta Sampling using two diffusion models, ADM-G and Stable Diffusion. Our results demonstrated that Beta Sampling is more advantageous than uniform sampling for both models and is particularly more efficient than AutoDiffusion for ADM-G with steps greater than 10. How-

ever, in the case of Stable Diffusion, Beta Sampling showed marginally inferior results to AutoDiffusion even at 10 steps. This discrepancy can be attributed to the differences in the spectral analysis of ADM-G and Stable Diffusion. As depicted in Fig. 2, ADM’s frequency changes exhibit peaks in both high and low frequencies concentrated at the extremities of the steps with a steep decline in variance. In contrast, Stable Diffusion’s frequency changes show a high-frequency peak around the 100-step mark, with low-frequency changes tapering off more gradually. Therefore, while Beta Sampling is well-suited to ADM’s frequency change distribution, it is not entirely compatible with Stable Diffusion’s distribution. Future improvements in generation performance for Stable Diffusion could be achieved by identifying hyperparameters or distributions that better align with its frequency change characteristics.

Limitations and Future Works. One limitation is that our approach relies on pre-determined spectral analyses, which may not fully capture the dynamic nature of the diffusion process across different datasets and model architectures. Future work could address these limitations by adaptively sampling the most critical steps to further enhance the efficiency and generalizability of our method.

7. Conclusion

In this paper, we introduced an efficient time step sampling method for generative diffusion models, leveraging frequency domain analysis to optimize the image generation process. Our approach replaces traditional uniform sampling with a Beta distribution-like method, emphasizing critical steps in the early and late stages of diffusion. We validated our hypothesis—that significant changes in image content occur at specific steps—through Fourier transform analysis, revealing substantial low-frequency changes early on and high-frequency adjustments later. Experiments with ADM-G and Stable Diffusion showed our method consistently outperforms uniform sampling, achieving better FID and IS scores, and offers competitive efficiency compared to state-of-the-art techniques like AutoDiffusion. This work provides a practical framework for enhancing diffusion model efficiency by focusing computational resources on the most impactful steps, with potential for further optimization and adaptive techniques in future research.

Acknowledgements

This work was partly supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (2022R1I1A1A01071970) and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.RS-2021-II212068, Artificial Intelligence Innovation Hub).

References

- [1] John E. Angus. The probability integral transform and related results. *SIAM Review*, 36(4):652–654, 1994. 4
- [2] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Qinsheng Zhang, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, Tero Karras, and Ming-Yu Liu. ediff-i: Text-to-image diffusion models with an ensemble of expert denoisers, 2023. 2
- [3] David Berthelot, Arnaud Autef, Jierui Lin, Dian Ang Yap, Shuangfei Zhai, Siyuan Hu, Daniel Zheng, Walter Talbott, and Eric Gu. Tract: Denoising diffusion models with transitive closure time-distillation, 2023. 2
- [4] Jooyoung Choi, Jungbeom Lee, Chaehun Shin, Sungwon Kim, Hyunwoo Kim, and Sungroh Yoon. Perception prioritized training of diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11472–11481, June 2022. 2, 3
- [5] Kamil Deja, Anna Kuzina, Tomasz Trzcinski, and Jakub Tomczak. On analyzing generative and denoising capabilities of diffusion-based deep generative models. *Advances in Neural Information Processing Systems*, 35:26218–26229, 2022. 3
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 3
- [7] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 1, 3, 4, 5, 6, 7, 8
- [8] Gongfan Fang, Xinyin Ma, and Xinchao Wang. Structural pruning for diffusion models. In A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 16716–16728. Curran Associates, Inc., 2023. 2
- [9] Joel Frank, Thorsten Eisenhofer, Lea Schönherr, Asja Fischer, Dorothea Kolossa, and Thorsten Holz. Leveraging frequency analysis for deep fake image recognition. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 3247–3258. PMLR, 13–18 Jul 2020. 3
- [10] 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*, 30, 2017. 4
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020. 1
- [12] Mahyar Khayatkhoei and Ahmed Elgammal. Spatial frequency bias in convolutional generative adversarial networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(7):7152–7159, Jun. 2022. 3
- [13] Yunsung Lee, Jin-Young Kim, Hyojun Go, Myeongho Jeong, Shinhyeok Oh, and Seungtaek Choi. Multi-architecture multi-expert diffusion models, 2023. 2, 3
- [14] Lijiang Li, Huixia Li, Xiawu Zheng, Jie Wu, Xuefeng Xiao, Rui Wang, Min Zheng, Xin Pan, Fei Chao, and Rongrong Ji. Autodiffusion: Training-free optimization of time steps and architectures for automated diffusion model acceleration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7105–7114, October 2023. 2, 4, 8
- [15] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 3
- [16] Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on manifolds. In *International Conference on Learning Representations*, 2022. 2, 4
- [17] Xingchao Liu, Xiwen Zhang, Jianzhu Ma, Jian Peng, et al. InstafLOW: One step is enough for high-quality diffusion-based text-to-image generation. In *The Twelfth International Conference on Learning Representations*, 2023. 7
- [18] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan LI, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 5775–5787. Curran Associates, Inc., 2022. 2
- [19] Chenlin Meng, Robin Rombach, Ruiqi Gao, Diederik Kingma, Stefano Ermon, Jonathan Ho, and Tim Salimans. On distillation of guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14297–14306, June 2023. 2
- [20] Ji-Hoon Park, Yeong-Joon Ju, and Seong-Wan Lee. Explaining generative diffusion models via visual analysis for interpretable decision-making process. *Expert Systems with Applications*, 248:123231, 2024. 2
- [21] Yong-Hyun Park, Mingi Kwon, Jaewoong Choi, Junghyo Jo, and Youngjung Uh. Understanding the latent space of diffusion models through the lens of riemannian geometry. In A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 24129–24142. Curran Associates, Inc., 2023. 2, 3
- [22] Nasim Rahaman, Devansh Arpit, Aristide Baratin, Felix Draxler, Min Lin, Fred A Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of deep neural networks. *arXiv preprint arXiv:1806.08734*, 4, 2018. 2
- [23] 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 (CVPR)*, pages 10684–10695, June 2022. 2, 3, 4, 6, 7

- [24] Amirmojtaba Sabour, Sanja Fidler, and Karsten Kreis. Align your steps: Optimizing sampling schedules in diffusion models. *arXiv preprint arXiv:2404.14507*, 2024. 5
- [25] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29, 2016. 4
- [26] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*, 2022. 2
- [27] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 3
- [28] Katja Schwarz, Yiyi Liao, and Andreas Geiger. On the frequency bias of generative models. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 18126–18136. Curran Associates, Inc., 2021. 3
- [29] Chenyang Si, Ziqi Huang, Yuming Jiang, and Ziwei Liu. Freeu: Free lunch in diffusion u-net. *arXiv preprint arXiv:2309.11497*, 2023. 3
- [30] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. 2, 4
- [31] Daniel Watson, William Chan, Jonathan Ho, and Mohammad Norouzi. Learning fast samplers for diffusion models by differentiating through sample quality. In *International Conference on Learning Representations*, 2022. 2
- [32] 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. 2
- [33] Zhi-Qin John Xu, Yaoyu Zhang, and Yanyang Xiao. Training behavior of deep neural network in frequency domain. In Tom Gedeon, Kok Wai Wong, and Minhoo Lee, editors, *Neural Information Processing*, pages 264–274, Cham, 2019. Springer International Publishing. 2
- [34] Shuai Yang, Yukang Chen, Luozhou WANG, Shu Liu, and Ying-Cong Chen. Denoising diffusion step-aware models. In *The Twelfth International Conference on Learning Representations*, 2024. 2
- [35] Xingyi Yang, Daquan Zhou, Jiashi Feng, and Xinchao Wang. Diffusion probabilistic model made slim. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22552–22562, June 2023. 2, 3