# GUIDANCE IN THE FREQUENCY DOMAIN ENABLES HIGH-FIDELITY SAMPLING AT LOW CFG SCALES

# **Anonymous authors**Paper under double-blind review

#### **ABSTRACT**

Classifier-free guidance (CFG) has become an essential component of modern conditional diffusion models. Although highly effective in practice, the underlying mechanisms by which CFG enhances quality, detail, and prompt alignment are not fully understood. This paper presents a novel perspective on CFG by analyzing its effects in the frequency domain, showing that low and high frequencies have distinct impacts on generation quality. Specifically, low-frequency guidance governs global structure and condition alignment, while high-frequency guidance mainly enhances visual fidelity. However, applying a uniform scale across all frequencies—as is done in standard CFG—leads to oversaturation and reduced diversity at high scales and degraded visual quality at low scales. Based on these insights, we propose frequency-decoupled guidance (FDG), an effective approach that decomposes CFG into low- and high-frequency components and applies separate guidance strengths to each component. FDG improves image quality at low guidance scales and avoids the drawbacks of high CFG scales by design, i.e., retaining the benefits of CFG over unguided generation without its common failures. Through extensive experiments across multiple datasets and models, we demonstrate that FDG consistently enhances sample fidelity while preserving diversity, leading to improved FID and recall compared to CFG, establishing our method as a plug-andplay alternative to standard classifier-free guidance.

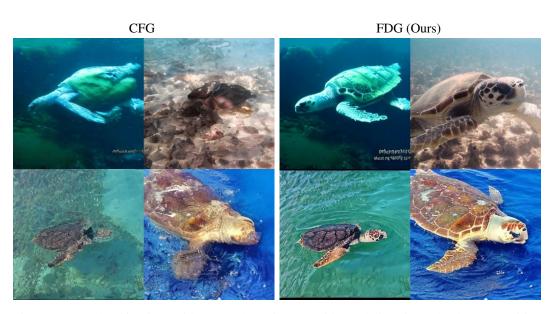


Figure 1: Low classifier-free guidance produces images with good diversity and color composition but often results in low-quality and blurry generations. We propose frequency-decoupled guidance (FDG), a novel modification to the CFG update rule in the frequency domain that significantly improves image quality at low guidance scales while, by design, avoiding common issues of high CFG scales such as reduced diversity. The examples here are generated using a class-conditional DiT-XL/2 model (Peebles & Xie, 2022) with a CFG scale of 1.2.

# 1 Introduction

Diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2021b) are a class of generative models that learn the data distribution by reversing a forward noising process that progressively corrupts data with increasing levels of Gaussian noise. While the theory suggests that simulating this reverse process should generate high-quality samples, in practice, unguided sampling often produces low-quality images that poorly match the input condition. To mitigate this, classifier-free guidance (CFG) (Ho & Salimans, 2022) has become a standard technique in modern diffusion models for improving quality and prompt alignment—though often at the cost of reduced diversity (Ho & Salimans, 2022; Sadat et al., 2024a) and excessive oversaturation (Sadat et al., 2025).

Current diffusion models typically rely on high guidance scales to achieve better image quality and prompt alignment. However, high guidance scales degrade sample diversity and introduce color saturation artifacts (Ho & Salimans, 2022). Conversely, low CFG scales tend to produce more diverse samples with natural color compositions but often suffer from poor global structure and lower visual fidelity. To address these trade-offs, several empirical strategies have been proposed to balance diversity and quality of CFG (Kynkäänniemi et al., 2024; Sadat et al., 2024a; 2025; WANG et al., 2024). Despite this progress, a systematic understanding of how CFG improves image quality and prompt alignment remains limited. Existing works neither fully explain the internal mechanisms of CFG nor explore how to improve generation quality at low guidance scales.

In this paper, our objective is to advance the understanding of how CFG works and improve image quality at low CFG scales, thereby avoiding the detrimental effects associated with high guidance scales. We begin by analyzing the CFG update rule in the frequency domain and show that CFG enhances image quality and prompt alignment through distinct frequency components. Specifically, we find that the low-frequency components of the CFG signal primarily govern global structure and condition alignment, while its high-frequency components mainly contribute to visual quality and details, with minimal impact on the overall composition of the image. We also observe that excessive guidance in the low-frequency domain leads to reduced diversity and oversaturation, whereas high-frequency components usually benefit from higher guidance scales. Since standard CFG applies a uniform scale across all frequencies, this explains why low CFG scales degrade quality, and high CFG scales boost detail at the cost of diversity and oversaturation.

Building on this insight, we propose a CFG scheme in the frequency domain, termed frequency-decoupled guidance (FDG), which disentangles the guidance scales applied to the low- and high-frequency components of the CFG update. We argue that low-frequency components should be guided more conservatively to avoid low-diversity and oversaturated generations, while high-frequency components can benefit from stronger guidance to enhance image quality. By assigning separate guidance strengths to these components, FDG improves image quality while retaining the diversity typically associated with low CFG scales. FDG is among the first approaches to systematically enhance the quality of low CFG scales, thereby avoiding the drawbacks of high CFG scales by design.

FDG introduces practically no additional sampling cost and can be applied to any pretrained diffusion model without extra training or fine-tuning. Through extensive experiments, we show that FDG consistently improves image quality and maintains diversity across a range of datasets, models, and metrics. In other words, FDG delivers the quality benefits of CFG over unguided generation while avoiding its usual drawbacks, such as reduced diversity. Moreover, FDG offers a deeper understanding of how CFG enhances image quality and prompt alignment by isolating the roles of low- and high-frequency components in the CFG update throughout the inference process. As such, we consider FDG as a superior plug-and-play alternative to the standard classifier-free guidance.

# 2 RELATED WORK

Score-based diffusion models (Song & Ermon, 2019; Song et al., 2021b; Sohl-Dickstein et al., 2015; Ho et al., 2020) learn data distributions by reversing a forward process that progressively corrupts data with Gaussian noise. They have rapidly surpassed prior generative techniques in fidelity and diversity (Nichol & Dhariwal, 2021; Dhariwal & Nichol, 2021), achieving state-of-the-art results across unconditional image synthesis (Dhariwal & Nichol, 2021; Karras et al., 2022; Yu et al., 2025a; Karras et al., 2024), text-to-image generation (Podell et al., 2023; Esser et al., 2024; Qin et al., 2025), video synthesis (Blattmann et al., 2023b;a; Bar-Tal et al., 2024; Wan et al., 2025), image-to-image

translation (Saharia et al., 2022a; Liu et al., 2023a; Xia et al., 2023), and audio synthesis (Chen et al., 2021; Huang et al., 2023; Liu et al., 2023b; Tian et al., 2025).

Recent studies have introduced numerous enhancements to the original DDPM framework (Ho et al., 2020), including improved network architectures (Hoogeboom et al., 2023; Karras et al., 2023; Peebles & Xie, 2022; Dhariwal & Nichol, 2021), novel sampling strategies (Song et al., 2021a; Karras et al., 2022; Liu et al., 2022b; Lu et al., 2022a; Salimans & Ho, 2022), and better training techniques (Nichol & Dhariwal, 2021; Karras et al., 2022; Song et al., 2021b; Salimans & Ho, 2022; Rombach et al., 2022). Despite these advancements, guidance methods—such as classifier-free guidance (Ho & Salimans, 2022)—remain essential for enhancing sample quality and improving alignment between conditioning information and generated outputs (Nichol et al., 2022).

Although several recent works have focused on improving diffusion model training for high-quality unguided generation (Karras et al., 2023; Yu et al., 2025b; Leng et al., 2025; Hoogeboom et al., 2024), modern diffusion models still rely on high guidance scales to enhance image quality and prompt alignment. However, this comes at the cost of reduced diversity (Ho & Salimans, 2022; Sadat et al., 2024a) and undesirable artifacts such as oversaturation (Sadat et al., 2025). Our work addresses this issue by disentangling the benefits of high CFG from its drawbacks. Leveraging frequency decomposition, we introduce a novel approach that improves image quality at lower CFG scales to avoid the trade-offs associated with high CFG scales. In other words, FDG enhances quality relative to unguided generation as the guidance scale increases, while avoiding the diversity loss and oversaturation effects of CFG.

Frequency decomposition techniques, such as Laplacian pyramids (Burt & Adelson, 1983) and wavelet transforms (Brewster, 1993), have recently been leveraged in generative models to enhance both quality and efficiency (Denton et al., 2015; Gal et al., 2021; Atzmon et al., 2024; Sadat et al., 2024b; Agarwal et al., 2025; Xiao et al., 2024). However, their potential for improving the sampling behavior of diffusion models remains underexplored. We demonstrate that applying guidance in the frequency domain unifies the advantages of both high and low CFG regimes (i.e., the fidelity of high scales and the diversity and color balance of low scales) into a single, inference-only method.

#### 3 BACKGROUND

**Diffusion models** Let  $x \sim p_{\text{data}}(x)$  denote a data sample, and let  $t \in [0,1]$  represent a continuous time variable. The forward diffusion process adds noise to the data as  $z_t = x + \sigma(t)\epsilon$ , where  $\sigma(t)$  is a time-dependent noise schedule. This schedule controls the degree of corruption, with  $\sigma(0) = 0$  (no noise) and  $\sigma(1) = \sigma_{\text{max}}$  (maximum noise). As shown by Karras et al. (2022), this forward process corresponds to the following ODE:

$$dz = -\dot{\sigma}(t)\sigma(t)\nabla_{z_t}\log p_t(z_t)dt, \tag{1}$$

where  $p_t(z_t)$  is the distribution over noisy samples at time t, with  $p_0 = p_{\text{data}}$  and  $p_1 = \mathcal{N}(\mathbf{0}, \sigma_{\text{max}}^2 \mathbf{I})$ . Given access to the time-dependent score function  $\nabla_{z_t} \log p_t(z_t)$ , one can sample from the original data distribution by integrating the ODE in reverse from t = 1 to t = 0. Since this score function is unknown, it is approximated by a neural denoiser  $D_{\theta}(z_t, t)$ , which is trained to recover the clean data x from its noisy counterpart  $z_t$ . Conditional generation is enabled by augmenting the denoiser with an auxiliary input y, such as class labels or text prompts, yielding  $D_{\theta}(z_t, t, y)$ .

Classifier-free guidance Classifier-free guidance (CFG) is an inference technique designed to enhance the quality of generated samples by interpolating between conditional and unconditional model predictions (Ho & Salimans, 2022). Let  $y_{\text{null}} = \emptyset$  denote a null condition representing the unconditional case. CFG modifies the denoiser output at each sampling step as follows:

$$\hat{D}_{CFG}(\boldsymbol{z}_t, t, \boldsymbol{y}) = D_{\boldsymbol{\theta}}(\boldsymbol{z}_t, t, \boldsymbol{y}_{\text{null}}) + w(D_{\boldsymbol{\theta}}(\boldsymbol{z}_t, t, \boldsymbol{y}) - D_{\boldsymbol{\theta}}(\boldsymbol{z}_t, t, \boldsymbol{y}_{\text{null}})),$$
(2)

Here, w=1 corresponds to the unguided case (i.e., sampling with  $D_{\theta}(\boldsymbol{z}_t, t, \boldsymbol{y})$ ). The unconditional model  $D_{\theta}(\boldsymbol{z}_t, t, \boldsymbol{y}_{\text{null}})$  is typically trained by randomly replacing the condition  $\boldsymbol{y}$  with  $\boldsymbol{y}_{\text{null}}$  during training. Similar to the truncation trick in GANs (Brock et al., 2019), CFG improves image quality, but often at the cost of sample diversity (Murphy, 2023) and oversaturation (Sadat et al., 2025).

**Frequency decompositions** Multi-level frequency decompositions, such as Laplacian pyramids and wavelet transforms, are commonly used to separate an image into different frequency bands.

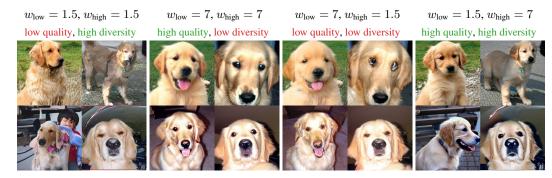


Figure 2: Illustration of how different frequency components of CFG affect generation. Sampling with low CFG scales results in diverse generations but lower overall quality. Increasing the CFG scale improves quality but reduces diversity. We show that the low-frequency component of the CFG signal primarily drives the reduction in diversity, while the high-frequency component contributes to quality enhancement without affecting diversity.

These techniques decompose an image into coarse, low-frequency structures and fine, high-frequency details. The low-frequency components capture the global characteristics of the image, including object placement, overall geometry, and color distribution, while the high-frequency components represent localized information such as edges, textures, and fine structural details.

# 4 CLASSIFIER-FREE GUIDANCE IN THE FREQUENCY DOMAIN

We now describe how guidance in the frequency domain can enhance the characteristics of CFG. Our first goal is to understand how different frequency components in the CFG prediction  $\hat{D}_{\text{CFG}}(\boldsymbol{z}_t, t, \boldsymbol{y})$  influence the final generation. Let  $\psi[\cdot]$  denote a linear and invertible frequency transformation, such as a Laplacian pyramid or wavelet transform, which decomposes each input  $\boldsymbol{x}$  into low- and high-frequency components, denoted by  $\psi_{\text{low}}[\boldsymbol{x}]$  and  $\psi_{\text{high}}[\boldsymbol{x}]$ , respectively. For notational convenience, we define  $D_c(\boldsymbol{z}_t) \doteq D_{\boldsymbol{\theta}}(\boldsymbol{z}_t, t, \boldsymbol{y}), \ D_u(\boldsymbol{z}_t) \doteq D_{\boldsymbol{\theta}}(\boldsymbol{z}_t, t, \boldsymbol{y}_{\text{null}}),$  and  $\hat{D}_{\text{CFG}}(\boldsymbol{z}_t) \doteq \hat{D}_{\text{CFG}}(\boldsymbol{z}_t, t, \boldsymbol{y})$  to represent the conditional, unconditional, and CFG outputs, respectively. As a consequence of the assumed properties of  $\psi$ , the CFG update rule can be expressed as

$$\hat{D}_{CFG}(\boldsymbol{z}_t) = \psi^{-1}[\psi[\hat{D}_{CFG}(\boldsymbol{z}_t)]]$$
(3)

$$= \psi^{-1}[\psi[D_u(z_t)] + w(\psi[D_c(z_t)] - \psi[D_u(z_t)])]. \tag{4}$$

This implies that, in the frequency domain, the CFG update affects both the low- and high-frequency components of  $\psi[\hat{D}_{CFG}(z_t)] = \{\hat{D}_{CFG}^{low}(z_t), \hat{D}_{CFG}^{high}(z_t)\}$  as follows:

$$\hat{D}_{CFG}^{low}(\boldsymbol{z}_t) = \psi_{low}[D_u(\boldsymbol{z}_t)] + w(\psi_{low}[D_c(\boldsymbol{z}_t)] - \psi_{low}[D_u(\boldsymbol{z}_t)]), \tag{5}$$

$$\hat{D}_{CFG}^{high}(\boldsymbol{z}_t) = \psi_{high}[D_u(\boldsymbol{z}_t)] + w(\psi_{high}[D_c(\boldsymbol{z}_t)] - \psi_{high}[D_u(\boldsymbol{z}_t)]). \tag{6}$$

Thus, in standard CFG, both low- and high-frequency components are guided using the same scale w throughout the sampling process. However, we argue that this approach is suboptimal, as the low- and high-frequency components exhibit different behaviors and influence distinct aspects of the generated image. We find that strong guidance on the low-frequency component  $\hat{D}_{\mathrm{CFG}}^{\mathrm{low}}(z_t)$  leads to oversaturation and reduced diversity, whereas high guidance on the high-frequency component  $\hat{D}_{\mathrm{CFG}}^{\mathrm{high}}(z_t)$  primarily enhances image quality. On the other hand, low scales for  $w_{\mathrm{low}}$  keeps diversity and realistic color composition while low values for  $w_{\mathrm{high}}$  degrade the visual details of the image and result in reduced sample quality. Motivated by this, we propose a generalized CFG scheme, called frequency-decoupled guidance (FDG), that employs separate guidance scales— $w_{\mathrm{low}}$  for low-frequency and  $w_{\mathrm{high}}$  for high-frequency components.

Figure 2 illustrates the effect of low- and high-frequency components on generated images. For CFG  $(w_{\text{low}} = w_{\text{high}})$ , low guidance scales lead to poor global structure and visual degradation, while high guidance reduces diversity. We observe that strong guidance on the low-frequency signal (i.e., a large  $w_{\text{low}}$ ) primarily causes diversity issues, whereas increasing  $w_{\text{high}}$  enhances quality without adverse effects on diversity. These findings highlight the limitations of using a single scale w in standard CFG: low w produces blurry or incoherent outputs, while high w reduces diversity and causes oversaturation. Our results therefore advocate for asymmetric guidance, where we set  $w_{\text{low}} < w_{\text{high}}$ .

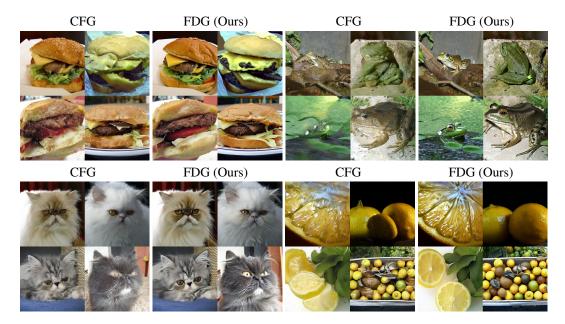


Figure 3: Class-conditional generation results using EDM2 with  $w=w_{\rm low}=1.25$ . FDG enhances image quality while maintaining the diversity of low CFG scales.

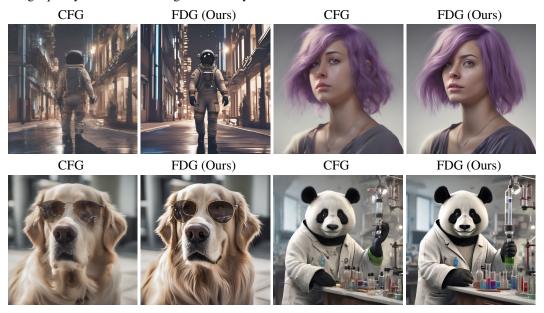


Figure 4: Text-to-image generation results using Stable Diffusion XL with  $w = w_{low} = 2$ . FDG enhances the details of the CFG image while maintaining the overall structure and color palette.

Implementation details We use Laplacian pyramids (Burt & Adelson, 1983) as the frequency transform  $\psi$  in our experiments. The complete algorithm for applying FDG is given in Algorithm 1, and the pseudocode of our method is provided in Algorithm 2. Notably, FDG requires only minor modifications to the standard CFG sampling procedure, introduces no significant computational overhead, and is readily compatible with all pretrained diffusion models.

# 5 EXPERIMENTS

**Setup** We primarily conduct experiments on text-to-image generation using Stable Diffusion models (Rombach et al., 2022; Podell et al., 2023; Esser et al., 2024), and class-conditional ImageNet (Russakovsky et al., 2015) generation using EDM2 (Karras et al., 2023) and DiT-XL/2 (Peebles & Xie, 2022). In all cases, we adopt the official pretrained checkpoints and codebases to maintain

Table 1: Quantitative comparison between CFG and FDG. FDG consistently improves FID and recall, showing high generation quality while maintaining good diversity.

| Model                                       | Guidance          | FID↓                  | Precision ↑         | Recall ↑            |
|---|-------------------|-----------------------|---------------------|---------------------|
| EDM2-S (Karras et al., 2023)                | CFG<br>FDG (Ours) | 9.77<br><b>5.44</b>   | <b>0.85</b> 0.83    | 0.52<br><b>0.68</b> |
| EDM2-XXL (Karras et al., 2023)              | CFG<br>FDG (Ours) | 8.65<br><b>4.99</b>   | <b>0.83</b> 0.82    | 0.56<br><b>0.68</b> |
| DiT-XL/2 (Peebles & Xie, 2022)              | CFG<br>FDG (Ours) | 9.31<br><b>5.33</b>   | <b>0.89</b><br>0.84 | 0.54<br><b>0.65</b> |
| Stable Diffusion 2.1 (Rombach et al., 2022) | CFG<br>FDG (Ours) | 24.99<br><b>23.33</b> | 0.68<br><b>0.69</b> | 0.44<br><b>0.49</b> |
| Stable Diffusion XL (Podell et al., 2023)   | CFG<br>FDG (Ours) | 25.23<br><b>24.60</b> | <b>0.64</b> 0.62    | 0.49<br><b>0.52</b> |
| Stable Diffusion 3 (Esser et al., 2024)     | CFG<br>FDG (Ours) | 30.32<br><b>27.68</b> | 0.76<br>0.76        | 0.37<br><b>0.42</b> |

Table 2: Quantitative comparison of CFG and FDG across various evaluation metrics for text-to-image models. FDG consistently outperforms CFG on all metrics and models.

| Benchmark                         | Model                | Guidance          | ImageReward ↑          | HPSv2↑                | PickScore ↑         | CLIP Score ↑          |
|-----------------------------------|----------------------|-------------------|------------------------|-----------------------|---------------------|-----------------------|
| DrawBench (Saharia et al., 2022b) | Stable Diffusion 2.1 | CFG<br>FDG (Ours) | -0.112<br><b>0.101</b> | 0.273<br><b>0.279</b> | 0.45<br><b>0.55</b> | 30.82<br><b>31.81</b> |
|                                   | Stable Diffusion XL  | CFG<br>FDG (Ours) | 0.323<br><b>0.595</b>  | 0.277<br><b>0.286</b> | 0.38<br><b>0.62</b> | 31.99<br><b>32.78</b> |
|                                   | Stable Diffusion 3   | CFG<br>FDG (Ours) | 0.185<br><b>0.803</b>  | 0.273<br><b>0.288</b> | 0.33<br><b>0.67</b> | 31.26<br><b>33.08</b> |
| Parti Prompts (Yu et al., 2022)   | Stable Diffusion 2.1 | CFG<br>FDG (Ours) | 0.109<br><b>0.332</b>  | 0.270<br><b>0.277</b> | 0.42<br><b>0.58</b> | 31.55<br><b>32.09</b> |
|                                   | Stable Diffusion XL  | CFG<br>FDG (Ours) | 0.442<br><b>0.663</b>  | 0.276<br><b>0.283</b> | 0.48<br><b>0.52</b> | 32.22<br><b>32.81</b> |
|                                   | Stable Diffusion 3   | CFG<br>FDG (Ours) | 0.527<br><b>0.988</b>  | 0.272<br><b>0.286</b> | 0.31<br><b>0.69</b> | 31.84<br><b>32.83</b> |
| HPS Prompts (Wu et al., 2023)     | Stable Diffusion 2.1 | CFG<br>FDG (Ours) | 0.035<br><b>0.282</b>  | 0.269<br><b>0.276</b> | 0.35<br><b>0.65</b> | 32.73<br><b>33.56</b> |
|                                   | Stable Diffusion XL  | CFG<br>FDG (Ours) | 0.603<br><b>0.824</b>  | 0.278<br><b>0.286</b> | 0.35<br><b>0.65</b> | 33.99<br><b>34.77</b> |
|                                   | Stable Diffusion 3   | CFG<br>FDG (Ours) | 0.470<br><b>1.00</b>   | 0.272<br><b>0.288</b> | 0.29<br><b>0.71</b> | 32.22<br><b>33.73</b> |

consistency with the original implementations. We provide extensive additional experiments and thorough ablations in Appendices B and C. Also, more details on the experimental setup and the hyperparameters used for each experiment are given in Appendix D.

**Evaluation metrics** For class-conditional models, we adopt Fréchet Inception Distance (FID) (Heusel et al., 2017) as the main metric to assess both image quality and diversity, given its strong correlation with human perception. To account for FID's sensitivity to implementation details, we evaluate all models under a consistent setup. Additionally, we report precision as a supplementary quality metric and recall to capture diversity (Kynkäänniemi et al., 2019). For text-to-image tasks, we further use ImageReward (Xu et al., 2023), HPSv2 (Wu et al., 2023), PickScore (Kirstain et al., 2023), and CLIP Score (Hessel et al., 2021) to evaluate image quality and prompt alignment.

#### 5.1 Main results

**Qualitative results** Figures 3 and 4 qualitatively compare the generations of FDG and CFG for the EDM2 and Stable Diffusion XL models at low CFG scales. FDG enhances generation quality while preserving the overall structure and color palette of the CFG output. Thus, FDG improves the quality at low CFG scales while avoiding the issues caused by high CFG values.

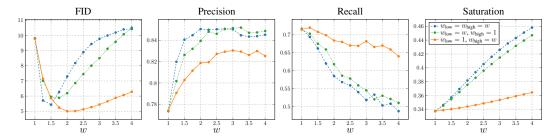


Figure 5: Illustrating the effect of different frequency components on CFG behavior. Although CFG improves quality, it rapidly restricts diversity and increases saturation, leading to higher FID and lower recall. Note that the low-frequency component is mainly responsible for these adverse effects, whereas the high-frequency component enhances quality while preserving diversity and color composition, resulting in better FID and recall.

**Quantitative results** Table 1 provides the metrics for FDG and CFG across several models at guidance values typically used in practice. We observe that FDG consistently improves FID and recall while largely maintaining the precision of the CFG outputs. We therefore conclude that FDG enhances quality while preserving diversity, leading to significantly better FID. Additionally, we evaluate generation quality at lower CFG values in Table 2 using metrics specifically designed to reflect human preferences for text-to-image models. As FDG boosts quality without introducing the detrimental effects of high CFG scales, it significantly outperforms CFG across all quality metrics for the Stable Diffusion models.

## 5.2 EFFECT OF DIFFERENT FREQUENCY COMPONENTS ON THE GENERATED DISTRIBUTION

To directly measure the effect of different frequency components on CFG outputs, we compared CFG with two sampling variants that use only  $\hat{D}_{\mathrm{CFG}}^{\mathrm{low}}(z_t)$  or  $\hat{D}_{\mathrm{CFG}}^{\mathrm{high}}(z_t)$ . This was achieved by setting  $w_{\mathrm{high}}=1$  or  $w_{\mathrm{low}}=1$  in FDG. As shown in Figure 5, increasing  $w_{\mathrm{low}}$  is the main cause of the reduced diversity in CFG, as evidenced by the higher FID and lower recall values. In contrast, increasing  $w_{\mathrm{high}}$  improves precision while maintaining recall, leading to significantly better FID across most guidance scales. These results suggest that excessive low-frequency guidance is the main contributor to the adverse effects of high CFG, whereas increasing high-frequency guidance generally enhances generation quality. Additionally, we observed that the low-frequency component is the primary cause of oversaturation, explaining why high guidance scales result in color artifacts. Therefore, our method adopts low values for  $w_{\mathrm{low}}$  and high values for  $w_{\mathrm{high}}$ .

# 5.3 Frequency analysis of prompt alignment in CFG

We next demonstrate how different frequency components of the CFG signal influence the alignment between generated images and the input condition. Figure 6 shows that although both low- and high-frequency components improve alignment as the guidance scale increases, the low-frequency component is the primary driver of this effect. Therefore, FDG can achieve comparable or superior prompt alignment to CFG by appropriately combining low- and high-frequency components, as shown in Table 2.

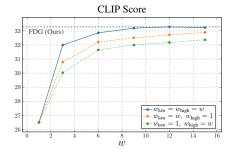


Figure 6: Effect of frequency components of CFG on prompt alignment.

#### 5.4 ADDITIONAL EXPERIMENTS

Relation to variable guidance scale Several works have explored the use of time-dependent guidance scales to balance diversity and quality in classifier-free guidance (Sadat et al., 2024a; Kynkäänniemi et al., 2024; WANG et al., 2024). For example, guidance interval (GI) (Kynkäänniemi et al., 2024) applies the guidance scale only during a limited range of sampling steps identified via grid search. We argue that the improvement in diversity offered by these methods is closely related to the frequency decomposition of the guidance signal. To support this, we provide the norms of

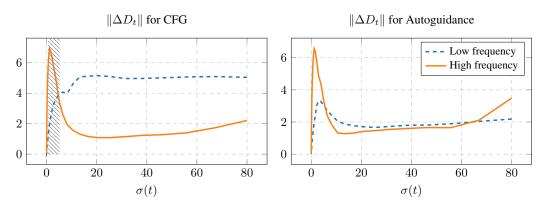


Figure 7: Illustration of how the norm of the guidance signal  $\Delta D_t$  behaves in the frequency domain for CFG and Autoguidance. For CFG, low-frequency components are dominant during most early steps (high  $\sigma(t)$ ), which can be harmful. Guidance interval (Kynkäänniemi et al., 2024) improves this by limiting CFG to the shaded region (found by grid search), i.e., the steps where high-frequency components are dominant. In contrast, Autoguidance maintains strong high-frequency norms throughout the sampling process, making the low-frequency signal useful at all steps (Karras et al., 2024).

Table 3: Comparing guidance interval (Kynkäänniemi et al., 2024) with FDG based on Stable Diffusion XL (Podell et al., 2023). FDG achieves better quality metrics due to having high-frequency guidance in earlier sampling steps.

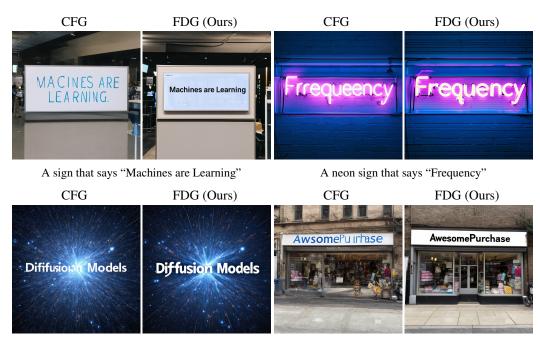
| Method            | ImageReward ↑ | HPSv2↑ | PickScore ↑ | CLIP Score ↑ |
|-------------------|---------------|--------|-------------|--------------|
| Guidance interval | 0.437         | 0.282  | 0.34        | 32.77        |
| FDG (Ours)        | 0.595         | 0.286  | 0.66        | 32.78        |

Table 4: Quantitative comparison between CFG and FDG using SDXL-Lightning (Lin et al., 2024) as an example of a distilled model that uses fewer sampling steps. FDG outperforms sampling both with and without CFG, achieving better quality and good prompt alignment.

| Method     | ImageReward ↑ | HPSv2↑ | CLIP Score ↑ |
|------------|---------------|--------|--------------|
| w/o CFG    | 0.535         | 0.282  | 31.77        |
| CFG        | 0.573         | 0.286  | 32.44        |
| FDG (Ours) | 0.672         | 0.292  | 32.44        |

the low- and high-frequency components of the guidance signal across sampling steps in Figure 7. Note that the norm of the high-frequency component increases over sampling, while the norm of the low-frequency component decreases. We observe that GI starts applying guidance when the norms of the two components become roughly equal. This suggests that applying guidance in the mid-stages, as GI does, implicitly increases the effective guidance on the high-frequency component relative to the low-frequency component. This analysis also provides a principled way to select an interval for applying GI, avoiding costly grid searches. Additionally, GI may still lead to quality degradation due to the absence of guidance at the beginning of sampling. As shown in Table 3, FDG outperforms GI in terms of image quality for Stable Diffusion XL by maintaining high-frequency guidance during the early sampling steps.

**Frequency analysis of Autoguidance** Autoguidance (Karras et al., 2024) proposed a modified version of CFG that replaces the unconditional prediction with a degraded version of the main diffusion model. Figure 7 presents a frequency analysis of the update provided by Autoguidance, showing that both low- and high-frequency components remain strong throughout the sampling process, unlike CFG where low-frequency response dominates. This likely explains why Autoguidance is effective at all steps, and why it provides a better guidance direction compared to CFG.



A cosmic sign that says "Diffusion Models"

A storefront with "AwesomePurchase" written on it

Figure 8: Compared with CFG, FDG can improve the quality of text rendering in generated images while also using a low guidance scale to preserve the realism of generated images.

**Compatibility with distilled models** CFG is often detrimental to distilled models that use a small number of sampling steps. In contrast, we show that FDG can be effectively applied to distilled models, such as SDXL-Lightning (Lin et al., 2024), without quality degradation. Table 4 provides a quantitative comparison between FDG and CFG. Compared to both baselines, FDG achieves higher quality metrics with good prompt alignment.

Improving text rendering in diffusion models We next demonstrate that FDG can enhance the quality of generated text in Stable Diffusion 3 (Esser et al., 2024). Generating high-quality text requires substantial details, which poses a challenge for standard CFG, since high guidance scales often result in unrealistic samples. In contrast, Figure 8 shows that FDG achieves realistic generation and correct spelling of text by separately controlling  $w_{\text{low}}$  and  $w_{\text{high}}$ .

# 6 CONCLUSION

In this work, we have taken a principled look at classifier-free guidance in the frequency domain and have shown that its beneficial effects on structural fidelity and fine details stem from strong guidance applied to the high-frequency components of the CFG signal, while its detrimental impact on diversity and oversaturation arises from excessive guidance on the low-frequency components. Building on this insight, we proposed frequency-decoupled guidance (FDG) to disentangle guidance strength across frequency bands, applying conservative scaling to low frequencies while exploiting stronger scaling at high frequencies. This approach preserves the diversity and color composition of low guidance scales while enhancing details akin to high guidance scales. As a result, FDG improves the quality of low CFG values while avoiding the adverse effects of high CFG scales by design. Importantly, FDG introduces practically no additional training or sampling cost and can be seamlessly integrated as a *plug-and-play* enhancement to any pretrained diffusion model using CFG. As with CFG itself, challenges remain in accelerating sampling and improving generation quality in extreme out-of-distribution domains, which we identify as promising directions for future research.

#### BROADER IMPACT STATEMENT

Our approach has the potential to enhance the realism and quality of outputs generated by diffusion models without the need for costly retraining. As such, it offers practical benefits for visual content creation. However, with the continued advancement of generative modeling, the generation and dissemination of fabricated or inaccurate data become increasingly accessible. While advancements in AI-generated content have the potential to enhance productivity and creativity, it remains essential to critically assess the accompanying risks and ethical considerations. For a comprehensive discussion on ethics and creativity within computer vision, we direct readers to Rostamzadeh et al. (2021).

#### REPRODUCIBILITY STATEMENT

Our work leverages the official implementations of the pretrained models referenced in the main text. The inference procedure for applying FDG is outlined in Algorithm 1, while the corresponding pseudocode is provided in Algorithm 2. Further implementation details, including the hyperparameters used in our main experiments, are described in Appendix D.

# REFERENCES

- Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chattopadhyay, Yongxin Chen, Yin Cui, Yifan Ding, et al. Cosmos world foundation model platform for physical ai. *arXiv preprint arXiv:2501.03575*, 2025.
- Donghoon Ahn, Hyoungwon Cho, Jaewon Min, Wooseok Jang, Jungwoo Kim, Seonhwa Kim, Hyun Hee Park, Kyong Hwan Jin, and Seungryong Kim. Self-rectifying diffusion sampling with perturbed-attention guidance. *CoRR*, abs/2403.17377, 2024. doi: 10.48550/ARXIV.2403.17377. URL https://doi.org/10.48550/arXiv.2403.17377.
- Yuval Atzmon, Maciej Bala, Yogesh Balaji, Tiffany Cai, Yin Cui, Jiaojiao Fan, Yunhao Ge, Siddharth Gururani, Jacob Huffman, Ronald Isaac, et al. Edify image: High-quality image generation with pixel space laplacian diffusion models. *arXiv preprint arXiv:2411.07126*, 2024.
- Omer Bar-Tal, Hila Chefer, Omer Tov, Charles Herrmann, Roni Paiss, Shiran Zada, Ariel Ephrat, Junhwa Hur, Guanghui Liu, Amit Raj, et al. Lumiere: A space-time diffusion model for video generation. In *SIGGRAPH Asia 2024 Conference Papers*, pp. 1–11, 2024.
- Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, Varun Jampani, and Robin Rombach. Stable video diffusion: Scaling latent video diffusion models to large datasets. *CoRR*, abs/2311.15127, 2023a. doi: 10.48550/ARXIV.2311.15127. URL https://doi.org/10.48550/arXiv.2311.15127.
- Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22563–22575, 2023b.
- M. E. Brewster. An introduction to wavelets (charles k. chui). *SIAM Rev.*, 35(2):312–313, 1993. doi: 10.1137/1035061. URL https://doi.org/10.1137/1035061.
- Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net, 2019. URL https://openreview.net/forum?id=B1xsqj09Fm.
- Peter J. Burt and Edward H. Adelson. The laplacian pyramid as a compact image code. volume 31, pp. 532–540. 1983. doi: 10.1109/TCOM.1983.1095851. URL https://doi.org/10.1109/TCOM.1983.1095851.
- Nanxin Chen, Yu Zhang, Heiga Zen, Ron J. Weiss, Mohammad Norouzi, and William Chan. Wavegrad: Estimating gradients for waveform generation. In 9th International Conference on Learning

Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net, 2021. URL https://openreview.net/forum?id=NsMLjcFaO8O.

- Emily L Denton, Soumith Chintala, Rob Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. *Advances in neural information processing systems*, 28, 2015.
- Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat gans on image synthesis. In Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 8780–8794, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/49ad23d1ec9fa4bd8d77d02681df5cfa-Abstract.html.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. *arXiv preprint arXiv:2403.03206*, 2024.
- Rinon Gal, Dana Cohen Hochberg, Amit Bermano, and Daniel Cohen-Or. SWAGAN: a style-based wavelet-driven generative model. *ACM Trans. Graph.*, 40(4):134:1–134:11, 2021. doi: 10.1145/3450626.3459836. URL https://doi.org/10.1145/3450626.3459836.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. *CoRR*, abs/2104.08718, 2021. URL https://arxiv.org/abs/2104.08718.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 6626–6637, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/8a1d694707eb0fefe65871369074926d-Abstract.html.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *CoRR*, abs/2207.12598, 2022. doi: 10.48550/arXiv.2207.12598. URL https://doi.org/10.48550/arXiv.2207.12598.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html.
- Emiel Hoogeboom, Jonathan Heek, and Tim Salimans. simple diffusion: End-to-end diffusion for high resolution images. *CoRR*, abs/2301.11093, 2023. doi: 10.48550/arXiv.2301.11093. URL https://doi.org/10.48550/arXiv.2301.11093.
- Emiel Hoogeboom, Thomas Mensink, Jonathan Heek, Kay Lamerigts, Ruiqi Gao, and Tim Salimans. Simpler diffusion (sid2): 1.5 FID on imagenet512 with pixel-space diffusion. *CoRR*, abs/2410.19324, 2024. doi: 10.48550/ARXIV.2410.19324. URL https://doi.org/10.48550/arXiv.2410.19324.
- Qingqing Huang, Daniel S. Park, Tao Wang, Timo I. Denk, Andy Ly, Nanxin Chen, Zhengdong Zhang, Zhishuai Zhang, Jiahui Yu, Christian Havnø Frank, Jesse H. Engel, Quoc V. Le, William Chan, and Wei Han. Noise2music: Text-conditioned music generation with diffusion models. *CoRR*, abs/2302.03917, 2023. doi: 10.48550/arXiv.2302.03917. URL https://doi.org/10.48550/arXiv.2302.03917.
- Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. 2022. URL https://openreview.net/forum?id=k7FuTOWMOc7.
- Tero Karras, Miika Aittala, Jaakko Lehtinen, Janne Hellsten, Timo Aila, and Samuli Laine. Analyzing and improving the training dynamics of diffusion models, 2023.

- Tero Karras, Miika Aittala, Tuomas Kynkäänniemi, Jaakko Lehtinen, Timo Aila, and Samuli Laine. Guiding a diffusion model with a bad version of itself. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=bg6fVPVs3s.
- Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Pick-a-pic: An open dataset of user preferences for text-to-image generation. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=G5RwHpBUv0.
- Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada, pp. 3929–3938, 2019. URL https://proceedings.neurips.cc/paper/2019/hash/0234c510bc6d908b28c70ff313743079-Abstract.html.
- Tuomas Kynkäänniemi, Miika Aittala, Tero Karras, Samuli Laine, Timo Aila, and Jaakko Lehtinen. Applying guidance in a limited interval improves sample and distribution quality in diffusion models. *CoRR*, abs/2404.07724, 2024. doi: 10.48550/ARXIV.2404.07724. URL https://doi.org/10.48550/arXiv.2404.07724.
- Xingjian Leng, Jaskirat Singh, Yunzhong Hou, Zhenchang Xing, Saining Xie, and Liang Zheng. Repa-e: Unlocking vae for end-to-end tuning with latent diffusion transformers. *arXiv* preprint *arXiv*:2504.10483, 2025.
- Shanchuan Lin, Anran Wang, and Xiao Yang. Sdxl-lightning: Progressive adversarial diffusion distillation. *CoRR*, abs/2402.13929, 2024. doi: 10.48550/ARXIV.2402.13929. URL https://doi.org/10.48550/arXiv.2402.13929.
- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In David J. Fleet, Tomás Pajdla, Bernt Schiele, and Tinne Tuytelaars (eds.), Computer Vision ECCV 2014 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V, volume 8693 of Lecture Notes in Computer Science, pp. 740–755. Springer, 2014. doi: 10.1007/978-3-319-10602-1\_48. URL https://doi.org/10.1007/978-3-319-10602-1\_48.
- Guan-Horng Liu, Arash Vahdat, De-An Huang, Evangelos A. Theodorou, Weili Nie, and Anima Anandkumar. I<sup>2</sup>sb: Image-to-image schrödinger bridge. *CoRR*, abs/2302.05872, 2023a. doi: 10.48550/arXiv.2302.05872. URL https://doi.org/10.48550/arXiv.2302.05872.
- Haohe Liu, Zehua Chen, Yi Yuan, Xinhao Mei, Xubo Liu, Danilo Mandic, Wenwu Wang, and Mark D Plumbley. Audioldm: Text-to-audio generation with latent diffusion models. *arXiv* preprint arXiv:2301.12503, 2023b.
- Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on manifolds. 2022a. URL https://openreview.net/forum?id=PlKWVd2yBkY.
- Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on manifolds. In *The Tenth International Conference on Learning Representations, ICLR* 2022, *Virtual Event, April* 25-29, 2022. OpenReview.net, 2022b. URL https://openreview.net/forum?id=PIKWVd2yBkY.
- 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 *NeurIPS*, 2022a. URL http://papers.nips.cc/paper\_files/paper/2022/hash/260a14acce2a89dad36adc8eefe7c59e-Abstract-Conference.html.
- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *CoRR*, abs/2211.01095, 2022b. doi: 10.48550/arXiv.2211.01095. URL https://doi.org/10.48550/arXiv.2211.01095.

- 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.
- Kevin P. Murphy. *Probabilistic Machine Learning: Advanced Topics*. MIT Press, 2023. URL http://probml.github.io/book2.
- Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8162–8171. PMLR, 2021. URL http://proceedings.mlr.press/v139/nichol21a.html.
- Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. GLIDE: towards photorealistic image generation and editing with text-guided diffusion models. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (eds.), *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pp. 16784–16804. PMLR, 2022. URL https://proceedings.mlr.press/v162/nichol22a.html.
- William Peebles and Saining Xie. Scalable diffusion models with transformers. *CoRR*, abs/2212.09748, 2022. doi: 10.48550/arXiv.2212.09748. URL https://doi.org/10.48550/arXiv.2212.09748.
- Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: improving latent diffusion models for high-resolution image synthesis. *CoRR*, abs/2307.01952, 2023. doi: 10.48550/ARXIV.2307.01952. URL https://doi.org/10.48550/arXiv.2307.01952.
- Qi Qin, Le Zhuo, Yi Xin, Ruoyi Du, Zhen Li, Bin Fu, Yiting Lu, Jiakang Yuan, Xinyue Li, Dongyang Liu, et al. Lumina-image 2.0: A unified and efficient image generative framework. *arXiv* preprint *arXiv*:2503.21758, 2025.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pp. 10674–10685. IEEE, 2022. doi: 10.1109/CVPR52688.2022.01042. URL https://doi.org/10.1109/CVPR52688.2022.01042.
- Negar Rostamzadeh, Emily Denton, and Linda Petrini. Ethics and creativity in computer vision. *CoRR*, abs/2112.03111, 2021. URL https://arxiv.org/abs/2112.03111.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y. URL https://doi.org/10.1007/s11263-015-0816-y.
- Seyedmorteza Sadat, Jakob Buhmann, Derek Bradley, Otmar Hilliges, and Romann M. Weber. CADS: Unleashing the diversity of diffusion models through condition-annealed sampling. In *The Twelfth International Conference on Learning Representations*, 2024a. URL https://openreview.net/forum?id=zMoNrajk2X.
- Seyedmorteza Sadat, Jakob Buhmann, Derek Bradley, Otmar Hilliges, and Romann M. Weber. LiteVAE: Lightweight and efficient variational autoencoders for latent diffusion models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024b. URL <a href="https://openreview.net/forum?id=mTAbl8kUzq">https://openreview.net/forum?id=mTAbl8kUzq</a>.
- Seyedmorteza Sadat, Otmar Hilliges, and Romann M. Weber. Eliminating oversaturation and artifacts of high guidance scales in diffusion models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=e2ONKX6qzJ.

Chitwan Saharia, William Chan, Huiwen Chang, Chris A. Lee, Jonathan Ho, Tim Salimans, David J. Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In Munkhtsetseg Nandigjav, Niloy J. Mitra, and Aaron Hertzmann (eds.), SIGGRAPH '22: Special Interest Group on Computer Graphics and Interactive Techniques Conference, Vancouver, BC, Canada, August 7-11, 2022, pp. 15:1–15:10. ACM, 2022a. doi: 10.1145/3528233.3530757. URL https://doi.org/10.1145/3528233.3530757.

- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. 2022b. URL http://papers.nips.cc/paper\_files/paper/2022/hash/ec795aeadae0b7d230fa35cbaf04c041-Abstract-Conference.html.
- Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022.* OpenReview.net, 2022. URL https://openreview.net/forum?id=TIdIXIpzhoI.
- Chenyang Si, Ziqi Huang, Yuming Jiang, and Ziwei Liu. Freeu: Free lunch in diffusion u-net. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pp. 4733–4743. IEEE, 2024. doi: 10.1109/CVPR52733.2024.00453. URL https://doi.org/10.1109/CVPR52733.2024.00453.
- Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. 37:2256–2265, 2015. URL http://proceedings.mlr.press/v37/sohl-dickstein15.html.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net, 2021a. URL https://openreview.net/forum?id=St1giarCHLP.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 11895–11907, 2019. URL https://proceedings.neurips.cc/paper/2019/hash/3001ef257407d5a371a96dcd947c7d93-Abstract.html.
- Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net, 2021b. URL https://openreview.net/forum?id=PxTIG12RRHS.
- Richard Szeliski. Computer Vision Algorithms and Applications, Second Edition. Texts in Computer Science. Springer, 2022. ISBN 978-3-030-34371-2. doi: 10.1007/978-3-030-34372-9. URL https://doi.org/10.1007/978-3-030-34372-9.
- Zeyue Tian, Yizhu Jin, Zhaoyang Liu, Ruibin Yuan, Xu Tan, Qifeng Chen, Wei Xue, and Yike Guo. Audiox: Diffusion transformer for anything-to-audio generation, 2025.
- Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025.
- Xi WANG, Nicolas Dufour, Nefeli Andreou, Marie-Paule CANI, Victoria Fernandez Abrevaya, David Picard, and Vicky Kalogeiton. Analysis of classifier-free guidance weight schedulers. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=SUMtDJqicd.
- Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li. Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *arXiv preprint arXiv:2306.09341*, 2023.

- Bin Xia, Yulun Zhang, Shiyin Wang, Yitong Wang, Xinglong Wu, Yapeng Tian, Wenming Yang, Radu Timotfe, and Luc Van Gool. Diffi2i: Efficient diffusion model for image-to-image translation, 2023.
- Jun Xiao, Zihang Lyu, Hao Xie, Cong Zhang, Yakun Ju, Changjian Shui, and Kin-Man Lam. Frequency-aware guidance for blind image restoration via diffusion models. *CoRR*, abs/2411.12450, 2024. doi: 10.48550/ARXIV.2411.12450. URL https://doi.org/10.48550/arXiv.2411.12450.
- Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: learning and evaluating human preferences for text-to-image generation. In Proceedings of the 37th International Conference on Neural Information Processing Systems, pp. 15903–15935, 2023.
- Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, Ben Hutchinson, Wei Han, Zarana Parekh, Xin Li, Han Zhang, Jason Baldridge, and Yonghui Wu. Scaling autoregressive models for content-rich text-to-image generation. *Trans. Mach. Learn. Res.*, 2022, 2022. URL https://openreview.net/forum?id=AFDcYJKhND.
- Sihyun Yu, Sangkyung Kwak, Huiwon Jang, Jongheon Jeong, Jonathan Huang, Jinwoo Shin, and Saining Xie. Representation alignment for generation: Training diffusion transformers is easier than you think. 2025a.
- Sihyun Yu, Sangkyung Kwak, Huiwon Jang, Jongheon Jeong, Jonathan Huang, Jinwoo Shin, and Saining Xie. Representation alignment for generation: Training diffusion transformers is easier than you think. In *The Thirteenth International Conference on Learning Representations*, 2025b. URL https://openreview.net/forum?id=DJSZGGZYVi.
- Wenliang Zhao, Lujia Bai, Yongming Rao, Jie Zhou, and Jiwen Lu. Unipc: A unified predictor-corrector framework for fast sampling of diffusion models. *CoRR*, abs/2302.04867, 2023. doi: 10.48550/arXiv.2302.04867. URL https://doi.org/10.48550/arXiv.2302.04867.

# A BACKGROUND ON FREQUENCY DECOMPOSITIONS

**Laplacian Pyramids** The Laplacian pyramid (Burt & Adelson, 1983) is a multi-scale representation of an image based on successive band-pass filtering. Starting from an input image x, a Gaussian pyramid  $\{G_0, G_1, \ldots, G_N\}$  is constructed, where  $G_0 = x$ , and each level is a progressively downsampled (and low-pass filtered) version of the previous:

$$G_{i+1} = \mathsf{Downsample}(\mathsf{GaussianBlur}(G_i)).$$
 (7)

The Laplacian pyramid  $\{L_0, L_1, \dots, L_{N-1}\}$  is then formed by subtracting the upsampled version of each Gaussian level from its corresponding higher-resolution level:

$$L_i = G_i - \mathsf{Upsample}(G_{i+1}). \tag{8}$$

The top level of the pyramid is typically the final low-resolution image  $G_N$ . The decomposition can be inverted to reconstruct the original signal by sequentially upsampling and summing the Laplacian levels:

$$G_i = L_i + \mathsf{Upsample}(G_{i+1}). \tag{9}$$

This approach enables localized manipulation of image details at different spatial scales and is commonly used in image compression, enhancement, and blending (Szeliski, 2022).

Wavelet transforms Wavelet transforms (Brewster, 1993) are widely used in signal processing to extract spatial-frequency characteristics from input data. They rely on a pair of filters: a low-pass filter L and a high-pass filter H. In the case of 2D signals, four filters are derived as  $LL^{\top}$ ,  $LH^{\top}$ ,  $HL^{\top}$ , and  $HH^{\top}$ . When applied to an image x, the 2D wavelet transform decomposes it into a low-frequency sub-band  $x_L$  and three high-frequency sub-bands  $\{x_H, x_V, x_D\}$ , which capture horizontal, vertical, and diagonal details, respectively. For an input image of size  $H \times W$ , each sub-band has a spatial size of  $H/2 \times W/2$ . Multi-resolution analysis can be performed by recursively applying the wavelet transform to  $x_L$ . The transform is also invertible, allowing the reconstruction of the original image x from the set  $\{x_L, x_H, x_V, x_D\}$  using the inverse wavelet transform. Moreover, fast wavelet transform (FWT) (Mallat, 1989) makes it possible to compute wavelet sub-bands with linear complexity relative to the number of pixels in x.

#### B ADDITIONAL EXPERIMENTS

We present additional experiments in this section to further demonstrate the effectiveness of FDG across various scenarios and its compatibility with several alternative guidance methods.

#### B.1 EFFECT OF LOW- AND HIGH-FREQUENCY COMPONENTS ON GENERATIONS

To further demonstrate that low-frequency components govern global structure and high-frequency components contribute to details, we conducted an experiment in which we explicitly set either the high- or low-frequency portion of the CFG signal to zero. The results in Figure 9 show that when high frequencies are removed, the final generation retains the overall structure of the base image. Conversely, when low frequencies are removed, the generated image roughly shows the high-frequency details of the base image such as edges. These findings suggest that low-frequency components in the CFG signal control global structure, while high-frequency components determine more localized details.

#### B.2 COMPATIBILITY WITH ADAPTIVE PROJECTED GUIDANCE

Adaptive projected guidance (APG) (Sadat et al., 2025) is a method designed to reduce oversaturation and artifacts caused by high guidance scales. APG complements our methods, as its update rule can be integrated with the frequency decomposition in FDG to get better guidance directions. We demonstrate this combined approach in Figure 10 by incorporating the orthogonal projection from APG into FDG. The results indicate that this projection continues to effectively produce more realistic color compositions and fewer artifacts in generations.

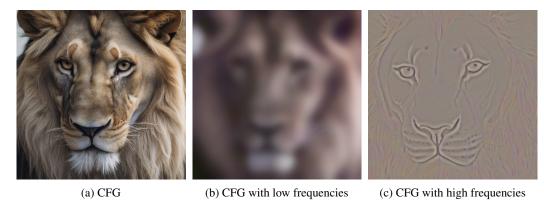


Figure 9: Effect of low- and high-frequency components of CFG on generated results. In this experiment, we explicitly set either low or high frequencies to zero to isolate the effect of the other component. Note that low frequencies determine the overall structure of the output, while high frequencies contribute to details.



Figure 10: Combining FDG with APG. The APG can be seamlessly integrated with FDG to produce generations with more realistic colors. For this experiment, we set  $w_{\text{low}} = 5$  and  $w_{\text{high}} = 15$ .

#### B.3 CHANGING THE WEIGHTS IN FDG

In this part, we further evaluate the effect of varying  $w_{\text{low}}$  and  $w_{\text{high}}$  on generation quality based on the EDM2-S model. For use  $w_{\text{high}} \in \{1.5, 2, 2.5\}$ , and we set  $w_{\text{low}} = r(w_{\text{high}} - 1) + 1$  with  $r \in \{0, 0.25, 0.5, 0.75, 1\}$ . Figure 11 shows that across all settings, increasing  $w_{\text{high}}$  and reducing the strength of the low-frequency component (r < 1) improves FID and recall while maintaining a comparable level of precision. Based on this, we recommend setting  $w_{\text{high}}$  relatively high and choosing  $r \leq 0.5$  across models. This experiment further demonstrates that by lowering the weight of the low-frequency component, FDG can simultaneously capture the benefits of both low and high CFG scales, i.e., achieving high precision and recall along with low FID.

#### B.4 COMBINING FDG WITH CADS

CADS (Sadat et al., 2024a) is an inference method designed to increase the diversity of diffusion models at high guidance scales by perturbing the conditional embedding with Gaussian noise. In this section, we show that CADS is compatible with our method, and that their combination outperforms either approach used in isolation. Table 5 supports this finding using the DiT-XL/2 model as a benchmark. We therefore conclude that the benefits of FDG are complementary to those of CADS.

#### B.5 Using different diffusion samplers

We also show that the effectiveness of FDG is not limited to a specific diffusion sampler. Table 6 compares the metrics of FDG and CFG across several popular diffusion samplers for DiT-XL/2, demonstrating that FDG consistently outperforms CFG across all setups.

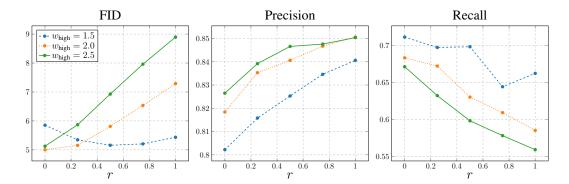


Figure 11: Quantitative evaluation of varying  $w_{\rm low}$  relative to  $w_{\rm high}$  (defined as  $w_{\rm low} = r(w_{\rm high}-1)+1$ ). Compared to the CFG baseline (r=1), increasing the high-frequency weight  $(w_{\rm high})$  while reducing the low-frequency weight (r<1) consistently improves FID and recall while maintaining precision. The best trade-off is achieved when  $w_{\rm high}$  is relatively high and r<0.5.

Table 5: Effectiveness of CADS on FDG using DiT-XL/2 at a guidance scale of 5. Combining FDG with CADS yields the best FID, outperforming each method used in isolation.

| Guidance          | FID↓  | Precision ↑ | Recall ↑ |
|-------------------|-------|-------------|----------|
| CFG               | 21.48 | 0.92        | 0.30     |
| CFG + CADS        | 14.53 | 0.88        | 0.48     |
| FDG (Ours)        | 13.90 | 0.90        | 0.48     |
| FDG + CADS (Ours) | 8.98  | 0.81        | 0.61     |

Table 6: Impact of applying FDG with popular diffusion samplers on the class-conditional ImageNet model (DiT-XL/2). FDG achieves improved FID and recall across all samplers.

|  | FDG (Ours)                           |                                      | CFG                                  |                                      |
|--|--------------------------------------|--------------------------------------|--------------------------------------|--------------------------------------|
| Sampler  | FID↓ Recall ↑                        |                                      | FID↓                                 | Recall ↑                             |
| DDIM (Song et al., 2021a)<br>DPM++ (Lu et al., 2022b)<br>SDE-DPM++ (Lu et al., 2022b)<br>PNDM (Liu et al., 2022a)<br>UniPC (Zhao et al., 2023) | 4.84<br>4.77<br>5.06<br>4.75<br>4.76 | 0.69<br>0.69<br>0.68<br>0.69<br>0.69 | 6.91<br>7.11<br>9.10<br>7.02<br>7.16 | 0.60<br>0.61<br>0.56<br>0.61<br>0.60 |

# B.6 CHANGING THE NUMBER OF SAMPLING STEPS

We also evaluated the performance of FDG and CFG across different numbers of sampling steps. Figure 12 shows that FDG maintains a consistent advantage over CFG across various sampling budgets, leading to improved FID and recall while preserving a similar level of precision. Therefore, we conclude that the observed improvements in FDG hold across different sampling budgets.

#### B.7 COMPATIBILITY WITH OTHER GUIDANCE METHODS

While our approach primarily focuses on beneficial weighting of frequency components tailored to the unique dynamics of CFG, the underlying principle of decoupling these components during guidance is more broadly applicable. FDG can be naturally extended to other guidance methods, though the optimal weighting schemes may differ depending on the specific characteristics of each algorithm. To investigate this generalization, we conducted additional experiments applying frequency-specific scaling to two alternative guidance strategies: perturbed attention guidance (PAG) (Ahn et al., 2024) with Stable Diffusion XL and autoguidance (Karras et al., 2024) with EDM2. The corresponding results in Tables 7 and 8 highlight the impact of FDG beyond CFG.

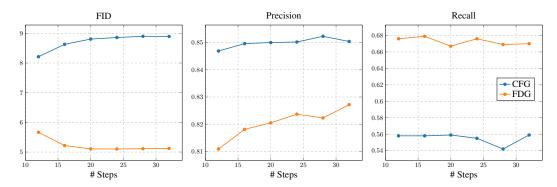


Figure 12: Comparison between FDG and CFG across different numbers of sampling steps. FDG consistently outperforms CFG, maintaining a clear FID improvement at all sampling budgets.

Table 7: Comparison of PAG and PAG+FDG with different weighting configurations.

| Method   | ImageReward ↑ | HPSv2↑ | HPS Win Rate ↑ | CLIP Score ↑ |
|--|---------------|--------|----------------|--------------|
| PAG (w = 4)  | -0.13         | 0.270  | 0.34           | 29.28        |
| +FDG ( $w_{\text{low}} = 4, w_{\text{high}} = 7$ ) | -0.05         | 0.272  | 0.63           | 29.33        |
| +FDG ( $w_{\text{low}} = 7, w_{\text{high}} = 4$ ) | -0.03         | 0.273  | 0.66           | 29.33        |

Table 8: Comparison of Autoguidance and FDG on EDM2 with different weighting schemes.

| Method   | FID↓ | Precision ↑ | Recall ↑ |
|--|------|-------------|----------|
| autoguidance ( $w = 3$ )                             | 4.82 | 0.77        | 0.73     |
| +FDG ( $w_{\text{low}} = 3, w_{\text{high}} = 1.5$ ) | 4.63 | 0.79        | 0.73     |

Table 9: Comparison of constant vs. dynamic weighting for FDG based on EDM2-S.

| Method                              | FID↓                | Precision ↑         | Recall ↑            |
|-------------------------------------|---------------------|---------------------|---------------------|
| CFG ( $w = 2.5$ )<br>FDG (constant) | 7.45<br><b>5.12</b> | 0.84<br><b>0.83</b> | 0.57<br><b>0.67</b> |
| FDG (dynamic)                       | 5.23                | 0.83                | 0.66                |

Interestingly, while the standard FDG configuration improves the performance of PAG, we observed that PAG exhibits sampling dynamics that differ from those of CFG. In particular, low-frequency components are underrepresented relative to the norm of the update, whereas in CFG the low-frequency components tend to dominate. Consequently, applying stronger guidance to low frequencies relative to high frequencies proved beneficial for PAG. A similar trend was observed with autoguidance, where performance improved with  $w_{\rm low}>w_{\rm high}$ . This is consistent with the analysis in Figure 7, which illustrates that autoguidance dynamics fundamentally differ from CFG by exhibiting stronger high-frequency components in the update steps. Overall, these results demonstrate that employing distinct guidance scales for low- and high-frequency components (i.e.,  $w_{\rm low}\neq w_{\rm high}$ ) is beneficial across both PAG and autoguidance.

#### B.8 Adaptive weights for FDG

Our pipeline supports time-dependent values for  $w_{\text{low}}$  and  $w_{\text{high}}$ . While the main experiments used constant scaling for simplicity, adaptive weighting of these parameters is also possible. To illustrate this, we conducted an experiment with EDM2-S using time-dependent weighting for FDG. For the dynamic variant, we set  $w_{\text{low}} = 1 + 0.5(1 - t)$  and  $w_{\text{high}} = 1 + 1.5(1 - t)$ , where  $t \in [0, 1]$ . Table 9 shows that both FDG variants outperform CFG, and the constant-weight configuration achieves a lower FID. As such, we opted for the constant weighting strategy in our main experiments.

Table 10: Effect of combining FDG with FreeU based on Stable Diffusion 2.1.

| Configuration | Image Reward ↑ | HPSv2↑ | CLIP Score ↑ |
|---------------|----------------|--------|--------------|
| CFG           | -0.112         | 0.273  | 30.82        |
| CFG + FreeU   | -0.043         | 0.274  | 30.84        |
| FDG           | 0.101          | 0.280  | 31.81        |
| FDG + FreeU   | 0.158          | 0.280  | 31.64        |

Table 11: Ablation of the frequency decomposition operator  $\psi$  based on the EDM2 model.

(a) The choice of the frequency decomposition function

| (ł | o) mu | lti-lev | el vs | sing | le-l | level | l transf | ormat | ion |
|----|-------|---------|-------|------|------|-------|----------|-------|-----|
|----|-------|---------|-------|------|------|-------|----------|-------|-----|

| $\psi$            | $FID\downarrow$ | Precision ↑ | Recall ↑ |
|-------------------|-----------------|-------------|----------|
| CFG               | 8.89            | 0.85        | 0.55     |
| Laplacian pyramid | 5.12            | 0.83        | 0.67     |
| Wavelet transform | 5.26            | 0.81        | 0.71     |

| Config       | FID↓ | Precision ↑ | Recall ↑ |
|--------------|------|-------------|----------|
| CFG          | 8.89 | 0.85        | 0.55     |
| single-level | 5.12 | 0.83        | 0.67     |
| multi-level  | 5.25 | 0.83        | 0.67     |

#### B.9 COMPLEMENTARITY WITH FREEU

FreeU (Si et al., 2024) is a method designed to improve the denoising quality of UNet-based diffusion models and is complementary to classifier-free guidance (CFG). Since FreeU operates on top of CFG, we can directly replace CFG with FDG in FreeU. Table 10 confirms that the effects of the two methods are indeed complementary. Notably, FDG alone outperforms both CFG and CFG + FreeU, while FDG + FreeU achieves the highest overall quality. These findings demonstrate that FDG and FreeU address complementary aspects of the generation process, distinguishing our contribution from Si et al. (2024). It is also worth noting that FreeU is applicable only to UNet-based models, whereas FDG can be applied to any architecture.

# C ABLATION STUDIES

Effect of the frequency decomposition operator We next test the performance of FDG for two choices of  $\psi$  in Table 11a. We note that FDG is not sensitive to this design choice as long as the operation provides an informative low- and high-frequency component. We chose the Laplacian pyramid, as it slightly outperformed DWT in our experiments.

**Using multi-level frequency decomposition** We also experimented with a multi-level Laplacian pyramid for frequency decomposition, applying different guidance scales to each frequency level. Table 11b shows both multi-level and single-level approaches outperform CFG. For simplicity, our main experiments used a single-level pyramid, though multi-level decomposition can help control separate high-frequency bands and remains viable for applications needing such control.

**Sampling time** On Stable Diffusion 3 (Esser et al., 2024), both CFG and FDG achieve a throughput of 1.22 iterations per second with identical memory usage when sampling with a batch size of four. This result clearly demonstrates that the frequency operations are lightweight and introduce virtually no additional overhead to the sampling process.

# D IMPLEMENTATION DETAILS

The sampling algorithm of FDG is provided in Algorithm 1, and the corresponding PyTorch implementation is given in Algorithm 2. Compared to CFG, FDG only adds a few extra lines of code and does not incur any noticeable computational overhead. As stated in the code, we convert the model's output to the denoised estimate  $D_{\theta}(z_t, t, y)$  (also known as the  $x_0$  prediction), apply the guidance step, and then convert it back to the original output format at each sampling step. For latent diffusion models, the frequency decomposition is performed in the latent space. The guidance scales are selected in the same way practitioners typically choose CFG values for a model, i.e., generating a few samples and visually inspecting them. This is standard practice in diffusion model use, requires

Table 12: Guidane parameters used to compare the performance of FDG with CFG.

(a) Guidance parameters used for Table 1.

| (b) Guidance parameters used for Table | (b) G | Guidance | parameters | used | for | Table | 2 |
|--|-------|----------|------------|------|-----|-------|---|
|--|-------|----------|------------|------|-----|-------|---|

| Model                | w  | $w_{\text{low}}$ | 2121 : 1 |
|----------------------|----|------------------|----------|
|                      |    | w <sub>low</sub> | wnign    |
| EDM2-S               | 3  | 1                | 3        |
| EDM2-XL              | 2  | 1                | 2        |
| DiT-XL/2             | 2  | 1                | 2        |
| Stable Diffusion 2.1 | 7  | 3                | 7        |
| Stable Diffusion XL  | 10 | 5                | 10       |
| Stable Diffusion 3   | 7  | 3                | 7        |

| Model                | w   | $w_{\mathrm{low}}$ | $w_{high}$ |
|----------------------|-----|--------------------|------------|
| Stable Diffusion 2.1 | 3   | 3                  | 12         |
| Stable Diffusion XL  | 3   | 3                  | 12         |
| Stable Diffusion 3   | 1.5 | 1.5                | 12         |

#### Algorithm 1 Guided sampling with FDG

**Require:** Frequency decomposition operators  $\psi[\cdot]$  and  $\psi^{-1}[\cdot]$  (e.g., Laplacian pyramid)

**Require:** Guidance weights  $w_{low}$  (low-frequency),  $w_{high}$  (high-frequency)

**Require:** Conditioning input y1: Initialize:  $z_1 \sim \mathcal{N}(\mathbf{0}, \sigma_{\max}^2 I)$ 2: **for**  $t = \{1, 1 - \delta t, \dots, 0\}$  **do** 

3: Compute the frequency decomposition of the conditional and unconditional predictions:

$$\psi[D_c(\boldsymbol{z}_t)] = \{\psi_{\text{low}}[D_c(\boldsymbol{z}_t)], \psi_{\text{high}}[D_c(\boldsymbol{z}_t)]\}$$

$$\psi[D_u(\boldsymbol{z}_t)] = \{\psi_{\text{low}}[D_u(\boldsymbol{z}_t)], \psi_{\text{high}}[D_u(\boldsymbol{z}_t)]\}$$

4: Compute the low- and high-frequency components of FDG

$$\hat{D}_{\text{FDG}}^{\text{low}}(\boldsymbol{z}_t) = \psi_{\text{low}}[D_u(\boldsymbol{z}_t)] + w_{\text{low}}(\psi_{\text{low}}[D_c(\boldsymbol{z}_t)] - \psi_{\text{low}}[D_u(\boldsymbol{z}_t)])$$

$$\hat{D}_{\text{FDG}}^{\text{high}}(\boldsymbol{z}_t) = \psi_{\text{high}}[D_u(\boldsymbol{z}_t)] + w_{\text{high}}(\psi_{\text{high}}[D_c(\boldsymbol{z}_t)] - \psi_{\text{high}}[D_u(\boldsymbol{z}_t)])$$

5: Convert the guided prediction to the data space using the inverse transform:

$$\hat{D}_{\text{FDG}}(\boldsymbol{z}_t) = \psi^{-1}[\{\hat{D}_{\text{FDG}}^{\text{low}}(\boldsymbol{z}_t), \hat{D}_{\text{FDG}}^{\text{high}}(\boldsymbol{z}_t)\}]$$

6: Perform one sampling step (e.g., one step of DDIM):

$$z_{t-1} = \text{diffusion\_reverse}(\hat{D}_{FDG}, z_t, t)$$

7: end for

8: **return**  $z_0$ 

no retraining, and incurs negligible computational cost. Moreover, as shown in Figure 5, FDG outperforms CFG across a wide range of parameter values, demonstrating robustness to scale choice.

For evaluation, we mainly rely on the ADM evaluation suite (Dhariwal & Nichol, 2021) to calculate FID, precision, and recall. For class-conditional ImageNet models, FID is computed using 10,000 generated images along with the complete training dataset. In the case of text-to-image models, FID is measured using the validation split of MS COCO 2017 (Lin et al., 2014). To evaluate text-to-image quality metrics such as ImageReward (Xu et al., 2023), we follow the official implementations and use the authors' provided test datasets. For PickScore (Kirstain et al., 2023), we calculate the win probability and report a win if one image outperforms the other by a margin greater than 0.1; otherwise, the result is reported as a tie. Details of the hyperparameters used in our main experiments are listed in Table 12.

# E MORE VISUAL EXAMPLES

This section provides additional samples comparing the performance of FDG with CFG. Figure 13 presents further results for class-conditional generation using DiT-XL/2 and EDM2. For both models, we observe that FDG preserves the structure of the base image while significantly enhancing the details. Additional results for text-to-image models are shown in Figures 14 and 15.

# **Algorithm 2** PyTorch implementation of FDG.

1134

```
1136
1137
        import torch
1138
        import kornia
1139
        from kornia.geometry.transform import build_laplacian_pyramid
1140
1141
1142
        def project(
1143
            v0: torch.Tensor, # [B, C, H, W]
1144
            v1: torch.Tensor, # [B, C, H, W]
1145
        ):
1146
            dtype = v0.dtype
1147
            v0, v1 = v0.double(), v1.double()
1148
            v1 = torch.nn.functional.normalize(v1, dim=[-1, -2, -3])
1149
            v0_parallel = (v0 * v1).sum(dim=[-1, -2, -3], keepdim=True) * v1
1150
            v0_orthogonal = v0 - v0_parallel
1151
            return v0_parallel.to(dtype), v0_orthogonal.to(dtype)
1152
1153
1154
        def build_image_from_pyramid(pyramid):
1155
            img = pyramid[-1]
1156
            for i in range(len(pyramid) - 2, -1, -1):
1157
                img = kornia.geometry.pyrup(img) + pyramid[i]
1158
            return img
1159
1160
1161
        # We assume all model predictions are converted to "x_0" prediction.
1162
        def laplacian_guidance(
1163
            pred_cond: torch.Tensor,
                                       # [B, C, H, W]
1164
            pred_uncond: torch.Tensor, # [B, C, H, W]
1165
            guidance_scale=[1.0, 1.0], # Guidance scales from high- to low-frequency
1166
                                      # Optional weights for projection
            parallel_weights=None,
1167
        ):
1168
            levels = len(guidance_scale)
1169
            if parallel_weights = None:
1170
                parallel_weights = [1.0] * levels
1171
1172
            pred_cond_pyramid = build_laplacian_pyramid(pred_cond, levels)
1173
            pred_uncond_pyramid = build_laplacian_pyramid(pred_uncond, levels)
1174
1175
            pred_guided_pyramid = []
1176
            parameters = zip(
1177
                pred_cond_pyramid, pred_uncond_pyramid, guidance_scale, parallel_weights
1178
1179
            for idx, (p_cond, p_uncond, scale, par_weight) in enumerate(parameters):
1180
                diff = p_cond - p_uncond
1181
                diff_parallel, diff_orthogonal = project(diff, p_cond)
1182
                diff = par_weight * diff_parallel + diff_orthogonal
1183
                p_guided = p_cond + (scale - 1) * diff
1184
                pred_guided_pyramid.append(p_guided)
1185
            pred_guided = build_image_from_pyramid(pred_guided_pyramid)
1186
            return pred_guided.to(pred_cond.dtype)
1187
```

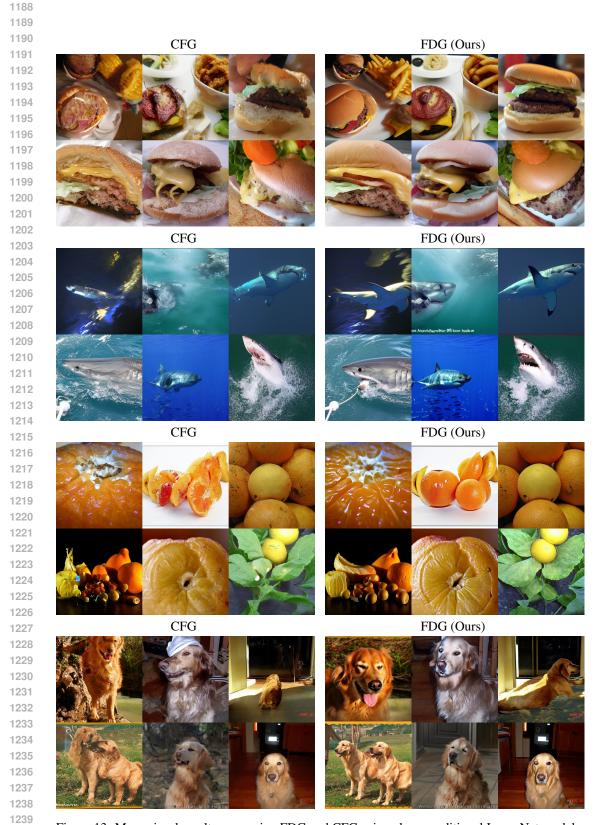


Figure 13: More visual results comparing FDG and CFG using class-conditional ImageNet models.

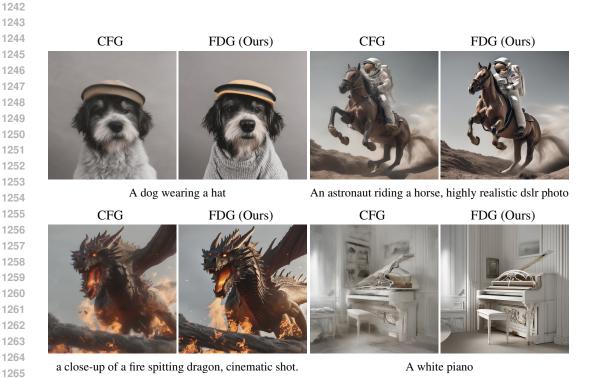


Figure 14: More visual examples comparing FDG with CFG using Stable Diffusion XL.

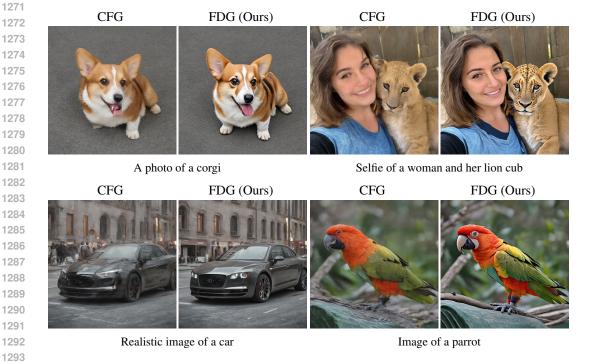


Figure 15: More visual examples comparing FDG with CFG using Stable Diffusion 3.