

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 FIXING MODEL-FITTING: COMPRESSING GUIDANCE FOR BETTER DIFFUSION SAMPLING

Anonymous authors

Paper under double-blind review

## ABSTRACT

Model-fitting occurs when samples are overly adjusted to satisfy with the guidance model rather than the true conditions, often leading to poor outcomes. The root cause of this problem is the consecutiveness of guidance timesteps throughout the diffusion sampling process. In this work, We quantify this effect and show that breaking the consecutiveness of standard guidance alleviates the problem. Based on this insight, our method, Compress Guidance, distributes a small number of guidance steps across the full sampling process, yielding substantial improvements in image quality and diversity while cutting guidance cost by over 80%. Experiments on both label-conditional image generation and text-to-image generation, across multiple datasets and models, confirm that Compress Guidance consistently surpasses baselines in image quality with significantly lower computational overhead.

## 1 INTRODUCTION

Guidance is mainly divided into classifier-free guidance in Ho & Salimans (2022) and classifier guidance in Dhariwal & Nichol (2021). Although both of these methods significantly improve the performance of the diffusion samples Dhariwal & Nichol (2021); Ho & Salimans (2022); Bansal et al. (2023); Liu et al. (2023); Epstein et al. (2023), they both suffer from high computation cost. For classifier guidance, the act of gradients calculation backwards through a classifier is costly. On the other hand, forwarding through a diffusion model twice at every timestep also costs significant computation in classifier-free guidance.

This work challenges the necessity of the current complex process based on several key observations. First, we find that the guidance loss is predominantly active during the early stages of the sampling process, when the image lacks a well-defined structure. As the model progresses and shifts its focus to refining image details, the guidance loss tends to approach zero. Additionally, when evaluating intermediate samples with an additional classifier not used for guidance, we observe that the loss from this external classifier does not decrease in the same way as it does for the guidance-specific classifier. This suggests that the generated samples are tailored to fit the features of the guiding classifier rather than producing generalized features applicable to different classifiers. We define this issue as *model-fitting*, where the generated image pixels are optimized to satisfy the guiding classifier’s criteria rather than generalizing to the intended conditions. The problem is validated by three pieces of evidence in section 4.1.

These observations prompt us to question whether guidance is necessary at every timestep and how reducing the frequency of guidance could enhance generative quality. In Section 4.2, we further explore the properties of guidance in ensuring sample quality. Based on this analysis, we propose a simple yet effective method called Compress Guidance (CompG), which mitigates the issue by reducing the number of time steps that invoke gradient calculation. This approach not only improves sample quality but also significantly accelerates the overall process as shown in Fig.1. In most parts of the works, we utilize classifier guidance as the main object for observations due to the explicit loss given by the classifier. However, the methods can be applied to classifier-free guidance as well.

Concurrent works have explored relevant ideas. Wang et al. (2024) shows that early guidance can cause conflicts and degrade outputs, while IntervalGuidance Kynkänniemi et al. (2025) reports the best results at mid-range noise timesteps. However, as demonstrated in Section 4.2, these insights are not always correct. Moreover, the cost reduction in IntervalGuidance arises only as a byproduct




Figure 1: *Stable Diffusion with classifier-free guidance. The left figure is the vanilla classifier-free guidance with application on all 50 timesteps. Our proposed Compress Guidance method is the right figure, where we only apply guidance on 8 out of 50 steps. Our methods are superior to classifier-free guidance regarding image quality, quantitative performance, and efficiency. The efficiency is calculated by sampling 30000 images with 1 GPU. More comparison is in Appendix H (Figure 13, 14, 15, 16)*

of avoiding low/high noise conflicts; it lacks an explicit mechanism for controlling guidance cost. In contrast, our method directly regulates guidance steps through a defined equation, improving efficiency without sacrificing performance. To the best of our knowledge, we are the first to explicitly reduce guidance cost without relying on additional distillation training, and we do so by principled analysis of model-fitting challenges. Furthermore, our plug-and-play approach is compatible with any model. Throughout most of our study, we observe on classifier guidance because the classifier provides explicit loss signals (unlike classifier-free guidance). Nevertheless, our observations extend naturally to classifier-free guidance, as discussed in Section 4.4.

Overall, the contributions of our works are three-fold: **(1)** Explore and quantify the model-fitting problem in guidance and the redundant computation resulting from current guidance methods. **(2)** Propose a simple but effective method to contain the model-fitting problem and improve computational time. **(3)** Extensive analysis and experimental results for different datasets and generative tasks on both classifier and classifier-free guidance perspectives.

## 2 RELATED WORK

Diffusion Generative Models (DGMs) Ho et al. (2020); Song et al. (2020b); Vahdat et al. (2021); Song & Ermon (2020); Lipman et al. (2022) have recently become one of the most popular generative models in many tasks such as image editing Kawar et al. (2023); Huang et al. (2024), text-to-image sampling Rombach et al. (2022); Podell et al. (2023); Ramesh et al. (2022) or image/videos generation Ho et al. (2022); Blattmann et al. (2023). Guidance is often utilized to improve the performance of DGMs Dhariwal & Nichol (2021); Ho & Salimans (2022); Bansal et al. (2023); Liu et al. (2023); Epstein et al. (2023); Wang et al. (2024); Karras et al. (2023; 2022). Besides improving the performance, the guidance also offers a trade-off between image quality and diversity Dhariwal & Nichol (2021); Ho & Salimans (2022); Ma et al. (2023), which helps users tune their sampling process up to their expectations. Gradient views of guidance are also well explored in the literature. Zheng et al. (2022) explores the gradient vanishing problem in classifier guidance, while Dinh et al. (2023a) examines conflicts in guidance sampling, and Dinh et al. (2023b) investigates guidance uncertainty during sampling. Our work introduces a new perspective on guidance—model-fitting—drawing an analogy to the overfitting problem in neural network training.

Although guidance is beneficial in many forms, it faces severe serious drawbacks in running time. For classifier guidance, the running time is around 80% higher compared to the original diffusion model

108 sampling time due to the evaluation of gradients at every sampling step. In contrast, classifier-free  
 109 guidance requires the process to forward to the expensive diffusion model twice at every timestep.  
 110 Previous works on improving the running time of DGMs involve the reduction of sampling steps  
 111 Song et al. (2020a); Zhang & Chen (2022); Song et al. (2023) and latent-based diffusion models  
 112 Rombach et al. (2022); Peebles & Xie (2023). Recently, the research community has focused on  
 113 distilling from a large number of timesteps to a smaller number of timesteps Salimans & Ho (2022);  
 114 Sauer et al. (2023); Li et al. (2024); Yin et al. (2024b;a); Song & Dhariwal (2023); Wang et al. (2025);  
 115 Meng et al. (2023); Heek et al. (2024); Yan et al. (2024); Luhman & Luhman (2021); Ren et al.  
 116 (2025); Zhou et al. (2024) or reducing the architectures of diffusion models Li et al. (2024); Tang  
 117 et al. (2023); Zhang et al. (2024). However, most of these works mainly solve the problem of the  
 118 expensive diffusion samplings, not the cost resulted by guidance.  
 119

120 Prior work Wang et al. (2024); Kynkänniemi et al. (2025) shows that early guidance can harm  
 121 generation in conditional diffusion models due to conflicts with conditional inputs. However, this  
 122 does not apply to unconditional models, where early guidance is essential for quality. These studies  
 123 also overlook the high computational cost of guidance. In contrast, our work addresses both quality  
 124 and efficiency, generalizing model-fitting issues to all guided diffusion and proposing a plug-and-play  
 125 solution to reduce guidance overhead without sacrificing performance.  
 126

### 3 BACKGROUND

127 **Diffusion Models** Ho et al. (2020) have the form of:  $p_\theta := p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$  where  
 128  $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) := \mathcal{N}(\mathbf{x}_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t))$  supporting the reverse process from  $\mathbf{x}_T$  to  $\mathbf{x}_0$ . This  
 129 process is denoising process where starting from the  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{x}_T; 0, \mathbf{I})$  to gradually move  
 130 to  $\mathbf{x}_0 \sim q(\mathbf{x}_0)$ . This process is trained to be matched with the forward diffusion process  
 131  $q(\mathbf{x}_{1:T}|\mathbf{x}_0) := \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1})$  given  $q(\mathbf{x}_t|\mathbf{x}_{t-1})$  as  $q(\mathbf{x}_t|\mathbf{x}_{t-1}) := \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I})$  or  
 132 we can write the conditional distribution of  $\mathbf{x}_t$  given  $\mathbf{x}_0$  as below:  
 133

$$q(\mathbf{x}_t|\mathbf{x}_0) := \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t) \mathbf{I}) \quad (1)$$

134  $\beta_t$  is the fixed variance scheduled before the process starts, Ho et al. (2020) denotes  $\alpha_t := 1 - \beta_t$   
 135 and  $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$  used in Eq.1. We have the  $\mathbf{x}_{t-1}$  conditioned on  $\mathbf{x}_0$  and  $\mathbf{x}_t$  as:  
 136

$$q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_{t-1}; \tilde{\mu}_t(\mathbf{x}_t, \mathbf{x}_0), \tilde{\beta}_t \mathbf{I}) \quad (2)$$

137 where  $\tilde{\mu}_t(\mathbf{x}_t, \mathbf{x}_0) := \frac{\sqrt{\bar{\alpha}_{t-1}} \beta_t}{1 - \bar{\alpha}_t} \mathbf{x}_0 + \frac{\sqrt{\alpha_t} (1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \mathbf{x}_t$  and  $\tilde{\beta}_t := \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t$ . To train the diffusion  
 138 model, the lower bound loss is utilized as below:  
 139

$$\mathbb{E}_{q}[-\log p_\theta(\mathbf{x}_0)] \leq \mathbb{E}_{q}[-\log p(\mathbf{x}_T) - \sum_{t \geq 1} \log \frac{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)}{q(\mathbf{x}_t|\mathbf{x}_{t-1})}] \quad (3)$$

140 Rewrite Eq. 3 as  $\mathbb{E}_q[D_{KL}(q(\mathbf{x}_T|\mathbf{x}_0)||p(\mathbf{x}_T)) + \sum_{t>1} D_{KL}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)||p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)) -$   
 141  $\log p_\theta(\mathbf{x}_0|\mathbf{x}_1)]$  The training process actually optimize the  $\sum_{t>1} D_{KL}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)||p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t))$   
 142 where the diffusion model try to match the distribution of  $\mathbf{x}_{t-1}$  by using only  $\mathbf{x}_t$ . There are several  
 143 implementations for optimising the 3. However, the  $\theta$  as parameters of the noise predictor  $\epsilon_\theta(\mathbf{x}_t, t)$  is  
 144 the most popular choice. After the  $\theta$  are trained using Eq. 3, we have the sampling equation:  
 145

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} (\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(\mathbf{x}_t, t)) + \sigma_t \mathbf{z} \quad (4)$$

146 **Guidance** in the Diffusion model offers conditional information and image quality enhancement.  
 147 Given a classifier  $p_\phi(y|\mathbf{x}_t)$  that match with the labels distribution conditioned on images  $\mathbf{x}_t$ , we have  
 148 the sampling equation with guidance as:  
 149

$$\mathbf{x}_{t-1} \sim \mathcal{N}(\mu_t + s \sigma_t^2 \nabla_{\mathbf{x}_t} \log p_\phi(y|\mathbf{x}_t), \sigma_t) \quad (5)$$

150 with  $s$  is the guidance scale. Besides the classifier guidance as Eq.5, Ho & Salimans (2022) proposes  
 151 another version named classifier-free guidance. This guidance method does not base the information  
 152 on a classifier. Instead, the guidance depends on the conditional information from a conditional  
 153 diffusion model. The sampling equation has the form:  
 154

$$\mathbf{x}_{t-1} \sim \mathcal{N}(\tilde{\mu}_t(\mathbf{x}_t, \frac{\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \tilde{\epsilon}_t}{\sqrt{\bar{\alpha}_t}}), \sigma_t) \quad (6)$$

155 given  $\tilde{\epsilon} = (1 + w) \epsilon_\theta(\mathbf{x}_t, c) - w \epsilon_\theta(\mathbf{x}_t)$  with  $w$  is the guidance scale.  
 156

162 **4 MODEL-FITTING IN GUIDANCE**  
 163

164 We begin by modelling the sampling equation as two distinct optimization objectives, illustrating  
 165 that the sampling process functions as a form of “training”, where parameters  $\mathbf{x}_t$  are optimized  
 166 over  $T$  timesteps. We then analyze the “training” of  $\mathbf{x}_t$  in light of these objectives, highlighting the  
 167 model-fitting problem that arises in the current guidance-driven sampling process. From Eq.4:

$$168 \mathbf{x}_{t-1} = \frac{(1 - \alpha_t) \sqrt{\bar{\alpha}_{t-1}} \mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(\mathbf{x}_t, t)}{1 - \bar{\alpha}_t} + \frac{(1 - \bar{\alpha}_{t-1}) \sqrt{\alpha_t}}{1 - \bar{\alpha}_t} \mathbf{x}_t + \sigma_t z \quad (7)$$

171 **Distribution matching objective:** Assuming that  $\epsilon_\theta(\mathbf{x}_t, t)$  is learned perfectly to match random  
 172 noise  $\epsilon$  at timestep  $t$ , we have  $\frac{\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(\mathbf{x}_t, t)}{\sqrt{\bar{\alpha}_t}} = \mathbf{x}_0$  is the exact prediction of  $\mathbf{x}_0$  at timestep  $t$   
 173 according to Eq.1. Denoting  $\tilde{\mathbf{x}}_0$  is the prediction of  $\mathbf{x}_0$  at timestep  $t$ , we can re-write the equation as  
 174 bellow:

$$175 \mathbf{x}_{t-1} = \frac{(1 - \alpha_t) \sqrt{\bar{\alpha}_{t-1}} \tilde{\mathbf{x}}_0 + (1 - \bar{\alpha}_{t-1}) \sqrt{\alpha_t}}{1 - \bar{\alpha}_t} \mathbf{x}_t + \sigma_t z \quad (8)$$

177 This equation 8 can be derived from  $q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0)$  in Eq. 2 with parameterized trick for Gaussian Dis-  
 178 tribution. Thus, the first aim of the sampling process is to match the distribution  $q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0)$ . Nev-  
 179 ertheless, the Eq.8 is based on the assumption that  $\tilde{\mathbf{x}}_0 \sim \mathbf{x}_0$ , which often does not hold when  $t \rightarrow T$ .  
 180 Given  $\tilde{\mathbf{x}}_0 = \frac{\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(\mathbf{x}_t, t)}{\sqrt{\bar{\alpha}_t}}$ , this formulation is rooted from  $\tilde{\mathbf{x}}_0 \sim \mathcal{N}(\frac{1}{\sqrt{\bar{\alpha}}}\mathbf{x}_t; \frac{\bar{\alpha}-1}{\bar{\alpha}}\mathbf{I})$  with assump-  
 181 tion that  $\epsilon_\theta(\mathbf{x}_t, t) \sim \epsilon$ . However,  $\epsilon_\theta(\mathbf{x}_t, t)$  is trained to minimize  $D_{KL}[q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) || p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t)]$   
 182 as in Ho et al. (2020) which actually causes a significantly distorted information if  $\epsilon_\theta(\mathbf{x}_t, t)$  is utilized  
 183 to sample  $\tilde{\mathbf{x}}_0$  from  $\mathbf{x}_t$  if  $t \rightarrow T$ . A smaller  $t$  would result in a better prediction of  $\mathbf{x}_0$  and with  $t = 0$ ,  
 184 we have  $\bar{\alpha} = 1$  resulting in  $\tilde{\mathbf{x}}_0 = \mathbf{x}_t$ .

185 **Theorem 1.** Assuming  $\epsilon_\theta$  is trained to converge with noise prediction error magnitude at a timestep  
 186  $t$  is approximate  $\Delta$ , the sampling process  $\mathbf{x}_{t-1} \sim q(\mathbf{x}_{t-1} | \mathbf{x}_t, \tilde{\mathbf{x}}_0)$  from  $T$  to 0 is equivalent to the  
 187 minimization of  $\|\mathbf{x}_0 - \tilde{\mathbf{x}}_0\|$  .wrt.  $\mathbf{x}_t$ .

188 Full proof is shown in Appendix F. If we consider  $\mathbf{x}_t$  of the Eq.8 as the set of optimization parameters,  
 189 the sampling process will have the objective  $\min_{\mathbf{x}_t} \|\mathbf{x}_0 - \tilde{\mathbf{x}}_0\|$ . The Eq.8 turns into:

$$191 \mathbf{x}_{t-1} = \mathbf{x}_t - \underbrace{\left( \frac{\sqrt{\alpha_t} - 1}{\sqrt{\bar{\alpha}_t}} \mathbf{x}_t + \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t} \sqrt{\alpha_t}} \epsilon_\theta(\mathbf{x}_t, t) - \sigma_t z \right)}_{\gamma_1 \nabla_{\mathbf{x}_t} \|\mathbf{x}_0 - \tilde{\mathbf{x}}_0\|} \quad (9)$$

194 Eq.9 turns the sampling process into a stochastic gradient descent process where the  $\mathbf{x}_t$  is the  
 195 parameter of the model at the timestep  $t$ , the updated direction into  $\mathbf{x}_t$  aims to satisfy  $\min_{\mathbf{x}_t} \|\mathbf{x}_0 - \tilde{\mathbf{x}}_0\|$ .  
 196 We denote gradient  $\nabla_{\mathbf{x}_t} \|\mathbf{x}_0 - \tilde{\mathbf{x}}_0\|$  as denoising gradient  $\nabla_{\mathbf{x}_t} d$ .

197 **Classification objective:** From Eq.5, we have the term  $s\sigma_t^2 \nabla_{\mathbf{x}_t} \log p_\phi(y | \mathbf{x}_t)$  is added to the sam-  
 198 pling equation for guidance. This term can be written in full form as  $s\sigma_t^2 \nabla_{\mathbf{x}_t} (q(y) \log q(y) -$   
 199  $q(y) \log p_\phi(y | \mathbf{x}_t))$  which is equivalent to  $-s\sigma_t^2 \nabla D_{KL}[q(y) || p_\phi(\hat{y} | \mathbf{x}_t)]$ . Combine Eq.9 with guid-  
 200 ance information in Eq.5, we have:

$$202 \mathbf{x}_{t-1} = \mathbf{x}_t - \underbrace{\left( \frac{\sqrt{\alpha_t} - 1}{\sqrt{\bar{\alpha}_t}} \mathbf{x}_t + \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t} \sqrt{\alpha_t}} \epsilon_\theta(\mathbf{x}_t, t) - \sigma_t z \right)}_{\gamma_1 \nabla_{\mathbf{x}_t} d} - \underbrace{(-s\sigma_t^2 \nabla_{\mathbf{x}_t} \log p_\phi(y | \mathbf{x}_t))}_{\gamma_2 \nabla_{\mathbf{x}_t} D_{KL}[q(y) || p_\phi(\hat{y} | \mathbf{x}_t)]} \quad (10)$$

205 As a result, the process of updating  $\mathbf{x}_t$  to  $\mathbf{x}_{t-1}$  is a “training” step to optimize to objective functions  
 206  $\|\mathbf{x}_0 - \tilde{\mathbf{x}}_0\|$  and  $D_{KL}[q(y) || p_\phi(\hat{y} | \mathbf{x}_t)]$  with two gradients respecting to  $\mathbf{x}_t$  as Eq.10. Since this  
 207 is similar to the training process, it is expected to face some problems in training deep neural  
 208 networks. In this work, the problem of model fitting is detected by observing the losses given by the  
 209 classification objective during the sampling process. For convenience, from later on,  $\nabla$  is denoted for  
 210  $\nabla_{\mathbf{x}_t}$ . Although classifier-free guidance does not have explicit loss function like classifier guidance,  
 211 the observation on classifier guidance still can be applied to classifier-free guidance.

213 **4.1 MODEL-FITTING**  
 214

215 Based on the optimization problem from the sampling process in the previous section, we first define  
 216 *on-sampling loss* and *off-sampling loss* for observation.




Figure 2: ImageNet256x256. The top row is the vanilla guidance (obtained by ADM-G Dhariwal & Nichol (2021)), where all the timesteps got the guidance information. The second and third rows are our proposed method, which only applies 35 time steps. The second row distributes the timesteps uniformly, while the third row distributes the timesteps toward the early stage of the sampling process. The Compress Guidance performs significantly better than the original guidance method. One blue stick means one guidance step.

**Definition 1.** *On-sampling loss/accuracy* refers to the loss or accuracy evaluated on the generated samples  $\mathbf{x}_t$  at timestep  $t$  during the diffusion sampling process, which consists of  $T$  timesteps. This loss is defined as  $-\log p_\phi(\hat{y}|\mathbf{x}_t)$  by the classifier parameters  $\phi$  that provide guidance throughout the sampling process.

**Definition 2.** *Off-sampling loss/accuracy* refers to the loss or accuracy evaluated on the generated samples  $\mathbf{x}_t$  at timestep  $t$  during the diffusion sampling process, which consists of  $T$  timesteps. This loss is defined as  $-\log p_{\phi'}(\hat{y}|\mathbf{x}_t)$  by the classifier parameters  $\phi'$  that **do not** provide guidance throughout the sampling process.

We set up the off-sampling classifier  $\phi'$  with the same architecture and performance as the on-sampling classifier  $\phi$  used for guidance. The only difference between the two models is the parameters. Off-sampling classifier is initialized as the parameters of the on-sampling classifier. We fine-tune the off-sampling model with 10000 timesteps with the same loss for training the on-sampling classifier. The testing accuracy between the off-sampling classifier and the on-sampling classifier is shown in Table 10 in Appendix D.

We visualize the *on-sampling* loss obtained from the noise-aware ADM classifier from Dhariwal & Nichol (2021) on ImageNet256x256, as shown in Figure 4. Our results indicate that classification information is predominantly active during the early stages and converges within the first 120 timesteps. In contrast, the *off-sampling* loss follows a different trend, converging only after the denoising process is nearly ended. This observation suggests that generated samples behave inconsistently when evaluated with classifiers of similar performance but different parameters, highlighting an over-adjustment to the model’s parameters rather than to the true characteristics or conditions of the generated images.

**Definition 3.** *Model-fitting* occurs when sampled images  $\mathbf{x}_t$  at timestep  $t$  is updated to maximize  $p_\phi(y|\mathbf{x}_t)$  or to satisfy the parameters of the  $\phi$  only instead of the real distribution  $q(y|\mathbf{x}_t)$ .

In practice, a pretrained  $p_\phi(y|\mathbf{x}_t)$  is only able to capture part of the  $q(y|\mathbf{x}_t)$ . Fitting solely with  $p_\phi(y|\mathbf{x}_t)$  limits the sample’s generalisation ability, leading to incorrect features or overemphasising certain details due to misclassification or overfocusing of the guidance classifier. Three pieces of evidence support that the vanilla guidance suffers from **model-fitting** problem.

**Evidence 1:** From Fig.4, we see that while the on-sampling loss converges around the 120<sup>th</sup> timestep, the off-sampling loss remains high until the diffusion model converges later. This indicates that samples  $\mathbf{x}_t$  at timestep  $t$  satisfy only the on-sampling classifier but not the off-sampling classifier, despite their identical performance and architecture. Although the off-sampling loss decreases by the end, a significant gap between the off-sampling and on-sampling losses persists. This supports our hypothesis that the guidance sampling process produces features that fit only the guidance classifier, not the conditional information.




Figure 3: **Left:**  $G$ ,  $UG$ ,  $ES$ ,  $DG$  represent vanilla guidance, uniform skipping, early stopping, and delayed guidance.  $UG$  suffers from non-convergence;  $ES$  from forgetting;  $DG$  is dominated by early signal. **Right:** Solid = on-sampling loss, dashed = off-sampling. Delaying guidance amplifies negative effects.

**Evidence 2:** Table 1 illustrates the model-fitting problem through accuracy metrics. With vanilla guidance, the accuracy is about 90.80% for the on-sampling classifier. However, the same samples evaluated by the off-sampling classifier or Resnet152 achieve only around 62.5% and 34.2% accuracy, respectively. This indicates that many features generated by the model are specific to the guidance classifier and do not generalize to other models.

**Evidence 3:** Figure 2 (first row) shows samples from vanilla guidance, where every sampling step receives guidance information. Applying guidance at all timesteps forces the model to fit the on-sampling classifier's perception. Often, this makes the model colour-sensitive, focusing only on generating the "orange" feature for Goldfish and ignoring other details such as the shape, position and texture. From the three pieces of evidence we can observe, we can conclude that the vanilla guidance scheme has suffered from the model-fitting problem.

**Analogy to overfitting:** In neural network training, we have a dataset  $\mathbf{x}$  and a classifier  $f_\theta(\mathbf{x})$  to approximate the posterior distribution  $p(y|\mathbf{x})$ . Let  $\mathbf{x}_{\text{train}}$  be the training data and  $\mathbf{x}_{\text{test}}$  the testing data. Overfitting occurs when  $f_\theta$  is tailored to fit  $\mathbf{x}_{\text{train}}$  but fails to generalize to the entire dataset  $\mathbf{x}$ . This is observed by the gap between training loss/accuracy and testing loss/accuracy on  $\mathbf{x}_{\text{train}}$  and  $\mathbf{x}_{\text{test}}$ .

Table 3: Overfitting vs. Model-Fitting

Aspect	Overfitting	Model-fitting
Train Data	$\mathbf{x}_{\text{train}}$	$f_{\phi_g}$
Test Data	$\mathbf{x}_{\text{test}}$	$f_{\phi_o}$
Parameters	$f_\phi$	$\mathbf{x}$

only the specific  $f_\phi$  instead of generalizing conditional information. Here,  $f_\phi$  is the on-sampling "data", and off-sampling "data"  $f_{\phi'}$  is used to observe the model-fitting, analogous to using training and testing data for overfitting observation.

**Note:** Similar model-fitting analysis for classifier-free guidance is in Appendix C.

## 4.2 ANALYSIS

Gradient over-calculation is the main reason for model-fitting. Thus, **gradient balance**, which is not to call too many times of gradient calculation, is required. A straightforward solution is to eliminate the gradient calculations for the later timesteps, which have been found to be less active, as shown in Figure 4. This approach is referred to as Early Stopping (ES), where guidance is halted from the 200<sup>th</sup> timestep onwards, continuing until the 0<sup>th</sup> timestep.

**Early Stopping:** Figure 3 demonstrates that ES suffers from the *forgetting* problem, where on-sampling classification loss increases during the remaining sampling process, negatively impacting the generative outputs. This suggests that the guidance requires the property of **continuity**, meaning the gap between consecutive guidance steps must not be too large to prevent the *forgetting* problem.

**Uniform skipping guidance:** We tried an alternative approach named Uniform Skipping Guidance (UG). In UG, 50 guidance steps are evenly distributed across 250 sampling steps, with guidance




Figure 4: (left)  $OADM-C$ , (right)  $Resnet152$  off-sampling loss. The On-sampling loss converges very early while leaving the off-sampling loss converges at the end of the process after the conclusion of the denoising process.

Table 1: A significant gap exists between the on-sampling and the off-sampling classifier in terms of accuracy, indicating model-fitting.

Evaluation Model	Accuracy
On-sampling classifier	90.8%
Off-sampling classifier	62.5%
Off-sampling Resnet152	34.2%

In the diffusion model's sampling process, the classifier  $f_\phi$  is pretrained or fixed. The aim is to adjust the samples  $\mathbf{x}$  to match the trained posterior  $p_\phi(y|\mathbf{x})$ . This process also uses Stochastic Gradient Descent with different roles:  $f_\phi$  acts as the fixed data, and  $\mathbf{x}$  are the trainable parameters. The model-fitting problem arises when  $\mathbf{x}$  is adjusted to fit

only the specific  $f_\phi$  instead of generalizing conditional information. Here,  $f_\phi$  is the on-sampling "data", and off-sampling "data"  $f_{\phi'}$  is used to observe the model-fitting, analogous to using training and testing data for overfitting observation.

324 applied every five steps. This ensures continuity throughout the sampling process, mitigating the  
 325 *forgetting* problem. However, as shown in Figure 4, UG encounters the issue of *non-convergence*,  
 326 where the classification magnitude is too weak and becomes overshadowed by the denoising signals  
 327 from the diffusion models, leading to poor conditional information. Thus, guidance must require  
 328 another property, which is **magnitude sufficiency**.

329 **Delayed Guidance (DG):** Prior work Kynkänniemi et al. (2025); Wang et al. (2024) suggests  
 330 delaying guidance to avoid conflicts with the diffusion model. However, as shown in Figure 3 (right),  
 331 longer delays worsen performance. *Why does this contradict Wang et al. (2024)?* That study assumes  
 332 a conditional diffusion model, where applying guidance too early causes conflicts between guidance  
 333 and conditional information of diffusion model. But in an unconditional model, there's no such  
 334 conflict—so delaying guidance only harms performance.

335 In summary, vanilla guidance faces the issue of *model-fitting*, while ES and UG fail due to the  
 336 *forgetting* and *non-convergence* problems, respectively. Therefore, the primary goal of our proposed  
 337 method is to meet three key conditions which are **gradient balance**, **guidance continuity** and  
 338 **magnitude sufficiency**.

### 340 4.3 COMPRESS GUIDANCE

342 To avoid calculating gradient too frequently, we propose to utilize the gradient from the previous  
 343 guidance step at several next sampling steps, given that the gradient magnitude difference between two  
 344 consecutive sampling steps is not too significant. By doing this, we can satisfy **magnitude sufficiency**  
 345 without re-calculating the gradient at every sampling step. Note that the gradient directions have not  
 346 been updated since the last guidance step, resulting in the **gradient balance**. Since all the sampling  
 347 step receives a guidance signal, the **continuity** is guaranteed.

348 The hypothesis for utilizing the gradient from the previous timestep is three-fold. First of all, the  
 349 avoidance of re-calculation of gradients frequently through the classifier prevents the generated  
 350 samples from capturing the classification pattern of the classifier and helps to avoid model-fitting.  
 351 Second, in the early stage, the avoidance of frequent gradient updates helps to avoid the noisy updated  
 352 direction given by noisy samples. Finally, when the image is clear in the later timesteps, it is safe to  
 353 skip the gradient calculation since the value of the gradient is less active during this stage as in Fig.4.

$$354 \quad \mathbf{x}_{t-1} = \begin{cases} \mathbf{x}_t - \gamma_1 \nabla d - \gamma_2 \nabla D_{KL}[q(\hat{y}|\mathbf{x}_t) || q(y)], & \text{if } t \in G \\ \mathbf{x}_t - \gamma_1 \nabla d - \gamma_2 \Gamma_t, & \text{otherwise} \end{cases} \quad (11)$$

357 The set  $G$  is the set of time-steps for which the gradient will be calculated.  $\Gamma$  is a variable  
 358 used to store the calculated gradient from the previous sampling step,  $\Gamma_t$  is updated by  
 359  $\Gamma_{t-1} = \begin{cases} \nabla D_{KL}[q(\hat{y}|\mathbf{x}_t) || q(y)], & \text{if } t \in G \\ \Gamma_t, & \text{otherwise.} \end{cases}$ . In practice, we find out that instead of duplicating  
 360 gradients as in Eq. 11, we can slightly improve the performance by compressing the duplicated  
 361 gradients into one guidance step instead of providing guidance to all sampling as in Eq.12. We name  
 362 this method as *Compress Guidance*. We modify the sampling equation as below:

$$364 \quad \mathbf{x}_{t-1} = \begin{cases} \mathbf{x}_t - \gamma_1 \nabla d - \gamma_2 \sum_{t=G_i}^{G_{i+1}} \Gamma_t, & \text{if } t = a_i \\ \mathbf{x}_t - \gamma_1 \nabla d, & \text{otherwise} \end{cases} \quad (12)$$

367 One of the algorithm's assumptions is that the magnitude is mostly the same for two consecutive  
 368 sampling steps. From Appendix G, we observe that the classification gradient magnitude difference  
 369 between two consecutive sampling steps is often larger in the early stage of the sampling process.  
 370 Thus, we propose a method that distributes more guidance toward the early sampling stage and  
 371 sparingly at the end of the process. This will help to avoid the significant accumulation of magnitude  
 372 differences in the early stage and help to deliver better performance as well as reduce the number of  
 373 guidance steps. The scheme is defined as Eq. 13.

$$374 \quad G_i = T - \lfloor \frac{T}{|G|^k} i^k \rfloor \quad \forall 0 \leq i \leq l, k \in [0; +\infty] \quad (13)$$

377 From the eq. 13, we have two main properties. First, when  $k \rightarrow +\infty$ , guidance timesteps are  
 distributed toward the early stage of the sampling process. Second, when  $k < 1$  and  $k \rightarrow 0$ , guidance




Figure 5: *ImageNet256x256*. Left: *Vanilla guidance applied at all timesteps*. Right: *Compact Guidance (ours) applied at 50 out of 250 timesteps*. *Compact Guidance* reduces over-emphasized features, correcting weird and incorrect details. Further results are in AppendixH

timesteps are distributed toward the late stage of the process. The proposed solution to select the timesteps for guidance as Eq.13 allows us to choose the number of timesteps we will do guidance and how to distribute these timesteps along the sampling process by adjusting the  $k$  values. The full proof of these properties is written in the Appendix F. The full algorithm is in Algorithm 1 (Appendix).

#### 4.4 COMPRESS GUIDANCE ON CLASSIFIER-FREE GUIDANCE

We start from the noise sampling equation of the classifier-free guidance as:  $\tilde{\epsilon} = (1+w)\epsilon_\theta(\mathbf{x}_t, c, t) - w\epsilon_\theta(\mathbf{x}_t, t) = \epsilon_\theta(\mathbf{x}_t, c, t) + w(\epsilon_\theta(\mathbf{x}_t, c, t) - \epsilon_\theta(\mathbf{x}_t, t)) = \epsilon_\theta(\mathbf{x}_t, c, t) + wC$ .  $C$  could stand for classification information as mentioned in Dinh et al. (2023c). Replace the  $\tilde{\epsilon}$  to Eq.9, we have:

$$\mathbf{x}_{t-1} = \mathbf{x}_t - \underbrace{\left( \frac{\sqrt{\alpha_t} - 1}{\sqrt{\alpha_t}} \mathbf{x}_t + \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t} \sqrt{\alpha_t}} \epsilon_\theta(\mathbf{x}_t, c, t) - \sigma_t \mathbf{z} \right)}_{\gamma_1 \nabla d \quad (\text{match with Eq. 9})} - \underbrace{\frac{\alpha_t - 1}{\sqrt{1 - \bar{\alpha}_t}} wC}_{\text{classification information}} \quad (14)$$

From this perspective, we can further apply the technique from Compress Guidance to the classification term in classifier-free guidance with the compression of classification information  $\frac{\alpha_t - 1}{\sqrt{1 - \bar{\alpha}_t}} C$ .

## 5 EXPERIMENTAL RESULTS

**Setup** Experiments are conducted on pretrained Diffusion models on *ImageNet 64x64*, *ImageNet 128x128*, *ImageNet 256x256*, *ImageNet 512x512* Deng et al. (2009) and *MSCOCO* Lin et al. (2014). The base Diffusion models utilized for label condition sampling task are ADM Dhariwal & Nichol (2021) and CADM Dhariwal & Nichol (2021) for classifier guidance, EDM2 Karras et al. (2023) DiTPeebles & Xie (2023) for classifier-free guidance (CFG) Ho & Salimans (2022), GLIDENichol et al. (2021) for CLIP text-to-image guidance and Stable Diffusion Rombach et al. (2022) for text-to-image classifier-free guidance. Other baselines we also do comparison is BigGAN Brock et al. (2018), VAQ-VAE-2 Zhao et al. (2020), LOGAN Wu et al. (2019), DCTransformers Nash et al. (2021). FID/sFID, Precision and Recall are utilized to evaluate image quality and diversity measurements. We denote Compress Guidance as “-CompG” and “-G” as vanilla guidance, “-CFG” is the CFG, and “-CompCFG” is our proposed Compress Guidance applying on CFG. Full results with details of the experimental set-up are discussed in Appendix D and E.

### 5.1 CLASSIFIER & CLASSIFIER-FREE GUIDANCE

Guidance in unconditional diffusion models enhances both image quality and diversity by providing conditional information during sampling, as shown in Table 4. CompressGuidance (CG) significantly improves FID, sFID, and Recall metrics, supported by qualitative evidence in Figures 5 and 11, and reduces guidance steps by 5x, leading to a 42% and 23% decrease in runtime on ImageNet 64x64 and 256x256, respectively (trade-off IS/FID can be observed in Fig. 8 and 9). In contrast, guidance in conditional diffusion models mainly boosts diversity, with smaller overall impact due to the model’s inherent conditional structure. As shown in Table 13, CompG still improves Recall and reduces guidance steps by 5x, with notable runtime savings of 39.79%, 29.63%, and 22% on ImageNet

432 64x64, 128x128, and 256x256 resolutions, respectively. From section 4.4, we also apply the CompG  
 433 technique on classifier-free guidance (CompCFG) and demonstrate the results in Table 5.  
 434

435 **Table 4: Unconditional guidance: CompG reduces  
 436 guidance by 5 $\times$  and improves performance.**

Model	$ G  (\downarrow)$	GPU hours ( $\downarrow$ )	FID ( $\downarrow$ )	sFID ( $\downarrow$ )	Prec ( $\uparrow$ )	Rec ( $\uparrow$ )
<b>ImageNet 64x64</b>						
ADM (No guidance)	0	26.33	9.95	6.58	0.60	0.65
ADM-G	250	54.86	6.40	9.67	0.73	0.54
ADM-CompG	50	31.80	5.91	8.26	0.71	0.56
<b>ImageNet 256x256</b>						
ADM (No guidance)	0	245.37	26.21	6.35	0.61	0.63
ADM-G	250	334.25	11.96	10.28	0.75	0.45
ADM-CompG	50	258.33	11.65	8.52	0.75	0.48

## 437 5.2 TEXT-TO-IMAGE GUIDANCE

438 **Table 6: Stable Diffusion on MSCOCO 256x256.**  
 439 **CompG improves quality (Fig. 1) and all metrics.**

Model	$ G  (\downarrow)$	GPU hrs ( $\downarrow$ )	FID ( $\downarrow$ )	IS ( $\uparrow$ )	CLIP ( $\uparrow$ )	GenEval ( $\uparrow$ )
SD-CFG	50	54	16.04	32.34	30	0.42
SD-CompCFG	8	35	14.04	35.90	31	0.43

## 440 5.3 ABLATION STUDY

441 **Distribution guidance timesteps toward the early stage of the**  
 442 **process:** According to the eq. 13, by adjusting  $k$ , we can distribute  
 443 the timesteps toward the early stage or the late stage of the sampling  
 444 process. Table 7 shows the comparison between  $k$  values. With  
 445  $k = 1.0$ , guidance steps are distributed uniformly. Larger  $k$  results  
 446 in comparable performance but more fruitful running time and the  
 447 number of guidance steps.  
 448

449 **Trade-off between computation and image quality** Compact rate  
 450 is the total number of sampling steps over the number of guidance  
 451 steps  $\frac{T}{|G|}$ . The larger the compact rate, the lower the model’s guid-  
 452 ance, hence the lower running time. Figure 12 shows the effect of  
 453 fewer timesteps on IS, FID and Recall as in Figure 12a, 12b and 12c.  
 454

455 **Table 7: ImageNet64x64.** Varying  $k$  shows improved  
 456 efficiency and quality with fewer guidance steps and  
 457 lower compute.

Model	$k$	$ G  (\downarrow)$	GPU hours ( $\downarrow$ )	FID ( $\downarrow$ )	sFID ( $\downarrow$ )	Prec ( $\uparrow$ )	Rec ( $\uparrow$ )
CADM (No guidance)	-	0	26.64	2.07	4.29	0.73	0.63
CADM-CompG	1.0	50	32.22	1.91	4.38	0.77	0.61
CADM-CompG	5.0	32	29.81	1.82	4.31	0.76	0.62
CADM-CompG	6.0	28	29.12	1.93	4.35	0.75	0.62

458 **Comparison with other guidance variants:** Table 8 compares our proposed CompCFG with Interval  
 459 Guidance methods from Kynkänniemi et al. (2025); Wang et al. (2024). CompCFG achieves results  
 460 comparable to IntG in Kynkänniemi et al. (2025), but with broader applicability. Unlike IntG, which  
 461 is limited to conditional diffusion models or classifier-free guidance, CompCFG can be integrated  
 462 into any diffusion model, delivering improved image quality and computational cost, as demonstrated  
 463 in Tables 4 and 5. Comparision with Dinh et al. (2023a;b); Zheng et al. (2022) is in Appendix E.1.  
 464


## 465 6 CONCLUSION

466 This paper quantifies model-fitting in diffusion model sampling, analogous to overfitting phenomenon,  
 467 by analyzing on- and off-sampling loss. To address this, we propose Compress Guidance, which  
 468 enhances generative performance while reducing guidance steps by at least fivefold and cutting  
 469 runtime by approximately 40%. Broader impacts and safeguards are discussed in Appendix A.  
 470

471 **Table 5: Conditional diffusion: CompCFG yield  
 472 lower FID and runtime with fewer guidance steps.**

Model	$ G  (\downarrow)$	GPU hours ( $\downarrow$ )	FID ( $\downarrow$ )	sFID ( $\downarrow$ )	Prec ( $\uparrow$ )	Rec ( $\uparrow$ )
<b>ImageNet 256x256</b>						
DIT (No guidance)	0	36.33	10.94	6.02	0.69	0.63
DIT-CFG	250	75.04	2.25	4.56	0.82	0.59
DIT-CompCFG	22	42.20	2.19	4.74	0.82	0.60
<b>ImageNet 512x512</b>						
EDM2 (No guidance)	0	4.22	2.23	5.21	0.75	0.62
EDM2-CFG	32	8.63	1.84	4.06	0.83	0.59
EDM2-CompCFG	6	5.06	1.63	3.91	0.80	0.61

473 We apply the CompG on this task with two  
 474 types of guidances, which are CLIP-based guid-  
 475 ance (GLIDE) Nichol et al. (2021) and classifier-  
 476 free guidance (Stable Diffusion) Rombach et al.  
 477 (2022). The results are shown in Table 12 and  
 478 6 and Figure 1.



479 **Figure 6: From 150 to 250 sampling steps. CompG narrows the  
 480 loss gap, mitigating overfitting. ES halts at 50 steps, leading to for-  
 481 getting problem and loss increase.**

482 **Table 8: ImageNet512x512.** Interval guidance with  
 483 CompCFG improves performance and diversity while  
 484 reducing steps.

Model	$ G  (\downarrow)$	Guidance Interval	FID ( $\downarrow$ )	sFID ( $\downarrow$ )	Prec ( $\uparrow$ )	Rec ( $\uparrow$ )
EDM2-IntG	6	[17, 22]	1.44	3.91	0.81	0.61
EDM2-CompCFG	6	[17, +∞)	1.44	3.88	0.81	0.62
EDM2-CompCFG	5	[17, +∞)	1.44	3.86	0.81	0.63
EDM2-CompCFG	4	[17, +∞)	1.45	3.87	0.80	0.63

486 REFERENCES  
487

488 Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas  
489 Geiping, and Tom Goldstein. Universal guidance for diffusion models. In *Proceedings of the*  
490 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 843–852, 2023.

491 Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler,  
492 and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion  
493 models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,  
494 pp. 22563–22575, 2023.

495 Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural  
496 image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.

497 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale  
498 hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*,  
499 pp. 248–255. Ieee, 2009.

500 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances*  
501 *in neural information processing systems*, 34:8780–8794, 2021.

502 Anh-Dung Dinh, Daochang Liu, and Chang Xu. Pixelsparam: A gradient view on diffusion sampling  
503 with guidance. In *International Conference on Machine Learning*, pp. 8120–8137. PMLR, 2023a.

504 Anh-Dung Dinh, Daochang Liu, and Chang Xu. Rethinking conditional diffusion sampling with  
505 progressive guidance. *Advances in Neural Information Processing Systems*, 36:42285–42297,  
506 2023b.

507 Anh-Dung Dinh, Daochang Liu, and Chang Xu. Pixelsparam: A gradient view on diffusion sampling  
508 with guidance. In *International Conference on Machine Learning*, pp. 8120–8137. PMLR, 2023c.

509 Dave Epstein, Allan Jabri, Ben Poole, Alexei Efros, and Aleksander Holynski. Diffusion self-  
510 guidance for controllable image generation. *Advances in Neural Information Processing Systems*,  
511 36:16222–16239, 2023.

512 Jonathan Heek, Emiel Hoogeboom, and Tim Salimans. Multistep consistency models. *arXiv preprint*  
513 *arXiv:2403.06807*, 2024.

514 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*,  
515 2022.

516 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in*  
517 *neural information processing systems*, 33:6840–6851, 2020.

518 Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J  
519 Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–8646,  
520 2022.

521 Yi Huang, Jiancheng Huang, Yifan Liu, Mingfu Yan, Jiaxi Lv, Jianzhuang Liu, Wei Xiong, He Zhang,  
522 Shifeng Chen, and Liangliang Cao. Diffusion model-based image editing: A survey, 2024.

523 Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-  
524 based generative models. In *Proc. NeurIPS*, 2022.

525 Tero Karras, Miika Aittala, Jaakko Lehtinen, Janne Hellsten, Timo Aila, and Samuli Laine. Analyzing  
526 and improving the training dynamics of diffusion models. *arXiv preprint arXiv:2312.02696*, 2023.

527 Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and  
528 Michal Irani. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the*  
529 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6007–6017, 2023.

530 Tuomas Kynkänniemi, Miika Aittala, Tero Karras, Samuli Laine, Timo Aila, and Jaakko Lehtinen.  
531 Applying guidance in a limited interval improves sample and distribution quality in diffusion  
532 models. *Advances in Neural Information Processing Systems*, 37:122458–122483, 2025.

540 Alexander C Li, Mihir Prabhudesai, Shivam Duggal, Ellis Brown, and Deepak Pathak. Your diffusion  
 541 model is secretly a zero-shot classifier. In *Proceedings of the IEEE/CVF International Conference*  
 542 *on Computer Vision*, pp. 2206–2217, 2023.

543 Yanyu Li, Huan Wang, Qing Jin, Ju Hu, Pavlo Chemerys, Yun Fu, Yanzhi Wang, Sergey Tulyakov,  
 544 and Jian Ren. Snapfusion: Text-to-image diffusion model on mobile devices within two seconds.  
 545 *Advances in Neural Information Processing Systems*, 36, 2024.

546 Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr  
 547 Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer vision–*  
 548 *ECCV 2014: 13th European conference, zurich, Switzerland, September 6–12, 2014, proceedings,*  
 549 *part v 13*, pp. 740–755. Springer, 2014.

550 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching  
 551 for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

552 Xihui Liu, Dong Huk Park, Samaneh Azadi, Gong Zhang, Arman Chopikyan, Yuxiao Hu, Humphrey  
 553 Shi, Anna Rohrbach, and Trevor Darrell. More control for free! image synthesis with semantic  
 554 diffusion guidance. In *Proceedings of the IEEE/CVF Winter Conference on Applications of*  
 555 *Computer Vision*, pp. 289–299, 2023.

556 Eric Luhman and Troy Luhman. Knowledge distillation in iterative generative models for improved  
 557 sampling speed. *arXiv preprint arXiv:2101.02388*, 2021.

558 Jiajun Ma, Tianyang Hu, Wenjia Wang, and Jiacheng Sun. Elucidating the design space of classifier-  
 559 guided diffusion generation. *arXiv preprint arXiv:2310.11311*, 2023.

560 Chenlin Meng, Robin Rombach, Ruiqi Gao, Diederik Kingma, Stefano Ermon, Jonathan Ho, and Tim  
 561 Salimans. On distillation of guided diffusion models. In *Proceedings of the IEEE/CVF Conference*  
 562 *on Computer Vision and Pattern Recognition*, pp. 14297–14306, 2023.

563 Charlie Nash, Jacob Menick, Sander Dieleman, and Peter W Battaglia. Generating images with  
 564 sparse representations. *arXiv preprint arXiv:2103.03841*, 2021.

565 Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew,  
 566 Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with  
 567 text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021.

568 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of*  
 569 *the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023.

570 Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe  
 571 Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image  
 572 synthesis. *arXiv preprint arXiv:2307.01952*, 2023.

573 Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-  
 574 conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.

575 Yuxi Ren, Xin Xia, Yanzuo Lu, Jiacheng Zhang, Jie Wu, Pan Xie, Xing Wang, and Xuefeng Xiao.  
 576 Hyper-sd: Trajectory segmented consistency model for efficient image synthesis. *Advances in*  
 577 *Neural Information Processing Systems*, 37:117340–117362, 2025.

578 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-  
 579 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-*  
 580 *ence on computer vision and pattern recognition*, pp. 10684–10695, 2022.

581 Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv*  
 582 *preprint arXiv:2202.00512*, 2022.

583 Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial diffusion  
 584 distillation. *arXiv preprint arXiv:2311.17042*, 2023.

585 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv*  
 586 *preprint arXiv:2010.02502*, 2020a.

594 Yang Song and Prafulla Dhariwal. Improved techniques for training consistency models. *arXiv*  
 595 *preprint arXiv:2310.14189*, 2023.  
 596

597 Yang Song and Stefano Ermon. Improved techniques for training score-based generative models.  
 598 *Advances in neural information processing systems*, 33:12438–12448, 2020.

599 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben  
 600 Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint*  
 601 *arXiv:2011.13456*, 2020b.

602 Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *ArXiv*, 2023.

603 Siao Tang, Xin Wang, Hong Chen, Chaoyu Guan, Yansong Tang, et al. Lightweight diffusion models  
 604 with distillation-based block neural architecture search. *arXiv preprint arXiv:2311.04950*, 2023.

605 Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space.  
 606 *Advances in neural information processing systems*, 34:11287–11302, 2021.

607 Fu-Yun Wang, Zhaoyang Huang, Alexander Bergman, Dazhong Shen, Peng Gao, Michael Lingelbach,  
 608 Keqiang Sun, Weikang Bian, Guanglu Song, Yu Liu, et al. Phased consistency models. *Advances*  
 609 in *Neural Information Processing Systems*, 37:83951–84009, 2025.

610 Xi Wang, Nicolas Dufour, Nefeli Andreou, Marie-Paule Cani, Victoria Fernández Abrevaya, David  
 611 Picard, and Vicky Kalogeiton. Analysis of classifier-free guidance weight schedulers. *arXiv*  
 612 *preprint arXiv:2404.13040*, 2024.

613 Yan Wu, Jeff Donahue, David Balduzzi, Karen Simonyan, and Timothy Lillicrap. Logan: Latent  
 614 optimisation for generative adversarial networks. *arXiv preprint arXiv:1912.00953*, 2019.

615 H Yan, X Liu, J Pan, JH Liew, Q Liu, and J Feng. Perflow: Piecewise rectified flow as universal  
 616 plug-and-play accelerator. *arXiv preprint arXiv:2405.07510*, 2024.

617 Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Frédéric Durand, and  
 618 William T Freeman. Improved distribution matching distillation for fast image synthesis. In  
 619 *NeurIPS*, 2024a.

620 Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Frédéric Durand, William T Freeman,  
 621 and Taesung Park. One-step diffusion with distribution matching distillation. In *CVPR*, 2024b.

622 Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator.  
 623 *arXiv preprint arXiv:2204.13902*, 2022.

624 Yang Zhang, Er Jin, Yanfei Dong, Ashkan Khakzar, Philip Torr, Johannes Stegmaier, and Kenji  
 625 Kawaguchi. Effortless efficiency: Low-cost pruning of diffusion models. *arXiv preprint*  
 626 *arXiv:2412.02852*, 2024.

627 Yang Zhao, Chunyuan Li, Ping Yu, Jianfeng Gao, and Changyou Chen. Feature quantization improves  
 628 gan training. *arXiv preprint arXiv:2004.02088*, 2020.

629 Guangcong Zheng, Shengming Li, Hui Wang, Taiping Yao, Yang Chen, Shouhong Ding, and Xi Li.  
 630 Entropy-driven sampling and training scheme for conditional diffusion generation. In *European*  
 631 *Conference on Computer Vision*, pp. 754–769. Springer, 2022.

632 Mingyuan Zhou, Huangjie Zheng, Zhendong Wang, Mingzhang Yin, and Hai Huang. Score identity  
 633 distillation: Exponentially fast distillation of pretrained diffusion models for one-step generation.  
 634 In *Forty-first International Conference on Machine Learning*, 2024.

635

636

637

638

639

640

641

642

643

644

645

646

647

---

**Algorithm 1** Compress Guidance
 

---

```

648
649 Input: class labels  $y$ , classification scale  $s$ 
650  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
651  $\Gamma \leftarrow 0$ 
652  $G \leftarrow$  Using Eq.13
653 for  $t = T, \dots, 1$  do
654  $z \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
655 if  $t \in G$  then
656  $g \leftarrow s \nabla_{\mathbf{x}_t} \log p_{\phi}(y | \mathbf{x}_t)$ 
657  $G'_t \leftarrow$  the next guidance step
658  $\Gamma \leftarrow g \times |t - G'_t|$ 
659  $\mathbf{x}_{t-1} \leftarrow \frac{1}{\sqrt{\alpha_t}} (\mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\alpha_t}} \epsilon_{\theta}(\mathbf{x}_t, t)) + \sigma_t^2 \Gamma + \sigma_t z$ 
660 else
661  $\mathbf{x}_{t-1} \leftarrow \frac{1}{\sqrt{\alpha_t}} (\mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\alpha_t}} \epsilon_{\theta}(\mathbf{x}_t, t)) + \sigma_t z$ 
662 end if
663 end for
664
665
```

---

**A BROADER IMPACT AND SAFEGUARD**

666 The work does not have concerns about safeguarding since it does not utilize the training data. The  
 667 paper only utilizes the pre-trained models from DiT Peebles & Xie (2023), ADMDhariwal & Nichol  
 668 (2021), GLIDE Nichol et al. (2021) and Stable Diffusion Rombach et al. (2022). The work fastens  
 669 the sampling process of the diffusion model and contributes to the population of the diffusion model  
 670 in reality. However, the negative impact might be on the research on a generative model where bad  
 671 people use that to fake videos or images.  
 672

**B FULL ALGORITHMS**

673 Algorithm 1 shows full algorithm. The full source code will be released beyond acceptance.  
 674

**C MODEL-FITTING ANALYSIS FOR CLASSIFIER-FREE GUIDANCE**

675 Different from classifier guidance, classifier-free guidance does not have an explicit classifier loss.  
 676 However, that does not mean that classifier-free guidance does not suffer from model-fitting. We  
 677 expand the analysis to CFG using the following steps:  
 678

- 679 1. Define the classifier in CFG
- 680 2. Define the observable loss
- 681 3. Observe on/off sampling

682 **Define the classifier in CFG** Alexander *et al.* Li et al. (2023) proved that a conditional diffusion  
 683 model is itself a classifier with the classification objective:

$$684 \arg \min_c \mathbb{E}_{t, \epsilon} [\|\epsilon_{\theta}(x_t, c) - \epsilon(x_t)\|], \quad (15)$$

685 where  $x_t$  is a noisy version of  $x_0$  with random noise  $\epsilon(x_t)$ .  
 686

687 Given this objective, we similarly define the objective function to “optimize” the parameter  $x_t$ , as  
 688

$$689 x_t^* = \arg \min_{x_t} \Delta_c(x_t) \quad (16)$$

$$690 \text{s.t. } \Delta_c(x_t) \leq \Delta_{c'}(x_t), \quad \forall c' \neq c, c' \in \mathcal{C}, \quad (17)$$

691 with  
 692

$$693 \Delta_{c'}(x_t) = \mathbb{E}_{\epsilon} [\|\epsilon_{\theta}(x_t, c') - \epsilon(x_t)\|_2^2], \quad c' \in \mathcal{C}, \quad (18)$$

694 as the logit for each class  $c'$ .  
 695

702 When denoising from  $x_t$ , we normally do not have  $\epsilon(x_t)$ . We replace this with the predicted  $\epsilon_\theta(x_t)$ ,  
 703 giving

$$\Delta_{c'}(x_t) = \mathbb{E}_\epsilon \left[ \|\epsilon_\theta(x_t, c') - \epsilon_\theta(x_t)\|_2^2 \right]. \quad (19)$$

706 Thus, the diffusion model itself acts as a classifier.

707 **Define observable loss:** We convert the objective into the observable cross-entropy loss as follows:

$$709 p(c' | x_t) = \frac{\exp(-\Delta_{c'}(x_t))}{\sum_{k \in \mathcal{C}} \exp(-\Delta_k(x_t))} \implies p(c | x_t) = \text{softmax}(-\Delta(x_t))_c. \quad (20)$$


$$712 L_{\text{CE}}(x_t, c) = -\log p(c | x_t) = \Delta_c(x_t) + \log \sum_{k \in \mathcal{C}} \exp(-\Delta_k(x_t)). \quad (21)$$

715 **On/Off sampling observation** We observe On/Off sampling loss using  $L_{\text{CE}}$ . We use two public  
 716 models from EDM2<sup>1</sup> which share the same architecture but differ in training hyperparameters:  
 717 EDM-S-0.025 and EDM-S-0.085. Their performance is reported below:

Model	FID	Accuracy
EDM-S-0.025	2.29	61%
EDM-S-0.085	2.40	63%

722 Table 9: General performance of diffusion models on generative task (FID) and classification task  
 723 (Accuracy).

725 The two models share similar performance. EDM-S-0.025 is used for on-sampling loss observation  
 726 (joined in the guidance process), and EDM-S-0.085 is used for off-sampling observation (not used  
 727 for guidance).



746 Figure 7: On-sampling and off-sampling loss across timesteps for EDM-S-0.025 and EDM-S-0.085  
 747 models. The On-sampling loss has a significant gap to the off-sampling loss. However, using our  
 748 proposed CompressGuidance (CompG) helps to close the gap between On/Off sampling loss.  
 749

750 The results show that model-fitting also occurs for CFG. Given the same diffusion models with similar  
 751 performance, the diffusion model used for guidance achieves much lower loss during sampling. The  
 752 use of CompG narrows the gap between off-sampling and on-sampling loss, indicating reduced  
 753 model-fitting. Furthermore, as shown in Section 5.3, our method significantly improves both runtime  
 754 and sampling quality.

755 <sup>1</sup><https://shorturl.at/uJxeV>




Figure 8: *FID curve given different guidance weight. This shows that the original vanilla guidance trades the quality with the diversity very significantly, while CompG helps to achieve the stability of the output.*

## D EXPERIMENTAL SETUP

**Off-sampling classifier:** Off-sampling classifier is initialized as the parameters of the on-sampling classifier. We fine-tune the model with 10000 timesteps with the same loss for training the on-sampling classifier. The testing accuracy between the off-sampling classifier and the on-sampling classifier is shown in Table 10

Evaluation Model	Accuracy
<i>On-sampling classifier</i>	64.5%
<i>Off-sampling classifier</i>	63.5%

Table 10: Evaluation of On-sampling classifier and Off-sampling classifier on ground-truth images.

Figure 11 shows all the hyperparameters used for all experiments in the paper. Normally, since we skip a lot of timesteps that do guidance, the process will fall into the case of forgetting. To avoid this situation, we would increase the guidance scale significantly. The value of the guidance scale is often based on the compact rate  $\frac{T}{|G|}$ . A larger compact rate also indicates a larger guidance scale. In Table 15 and Figure 6, to achieve a fair comparison, we tune the guidance scale of CompG to achieve a similar Recall value with vanilla guidance. The reason is that the higher the level of diversity, the harder features can be recognized, resulting in higher loss and lower accuracy. If we don't configure similar diversity between the two schemes, the one with higher diversity will always achieve lower accuracy and higher loss value. We want to avoid the case that the model only samples one good image for all.

For all the tables, the models which are in bold are the proposed.

**GPU hours:** All the GPU hours are calculated based on the time for sampling 50000 samples in ImageNet or 30000 samples in MSCoco.

All experiments are run on a cluster with 4 V100 GPUs.

## E FULL COMPARISON

Table 13 shows the full comparison with different famous baselines.

The stability of the CompG is visualized in Figure 8.

### E.1 ADDITIONAL ABLATION

In addition to Kynkänniemi et al. (2025), one of the most recent studies on guidance, we compare our proposed method with Dinh et al. (2023a;b); Zheng et al. (2022). Dinh et al. (2023a) addresses

810

811

812

MODEL	DATASET	$k$	$s$	$ G $	TIME-STEPS
<b>TABLE 4</b>					
ADM	IMAGENET 64x64	-	0.0	0	250
ADM-G	IMAGENET 64x64	-	4.0	250	250
ADM-COMPG	IMAGENET 64x64	1.0	4.0	50	250
ADM	IMAGENET 256x256	-	0.0	0	250
ADM-G	IMAGENET 256x256	-	4.0	250	250
ADM-COMPG	IMAGENET 256x256	1.0	4.0	50	250
<b>TABLE 5 &amp; 13</b>					
CADM	IMAGENET 64x64	-	0.0	0	250
CADM-G	IMAGENET 64x64	-	0.5	250	250
CADM-COMPG	IMAGENET 64x64	1.0	2.0	50	250
CADM-CFG	IMAGENET 64x64	-	0.1	250	250
CADM-COMPCFG	IMAGENET 64x64	5.0	0.1	25	250
CADM	IMAGENET 128x128	0.9	0.0	0	250
CADM-G	IMAGENET 128x128	-	0.5	250	250
CADM-CFG	IMAGENET 128x128	-	0.5	250	250
CADM	IMAGENET 256x256	-	0.0	0	250
CADM-G	IMAGENET 256x256	-	0.5	250	250
CADM-COMPG	IMAGENET 256x256	1.5	0.5	50	250
DIT-CFG	IMAGENET 256x256	-	1.5	250	250
DIT-COMPCFG	IMAGENET 256x256	1.2	1.5	22	250
EDM2-CFG	IMAGENET 256x256	-	1.2	32	32
EDM2-COMPCFG	IMAGENET 512x512	2.5	0.3	6	32
<b>TABLE 6 &amp; 12</b>					
GLIDE-G	MSCOCO 64x64	-	0.0	250	250
GLIDE-COMPG	MSCOCO 64x64	2.0	8.0	25	250
GLIDE-G	MSCOCO 256x256	-	0.0	250	250
GLIDE-COMPG	MSCOCO 256x256	2.0	5.5	35	250
SDIFF-CFG	MSCOCO 256x256	-	2.0 (FID, IS), 7.5 (CLIP, GENEVAL)	50	50
SDIFF-COMPCFG	MSCOCO 256x256	1.0	2.0 (FID, IS), 7.5 (CLIP, GENEVAL)	8	50
<b>TABLE 7</b>					
CADM	IMAGENET 64x64	-	0.0	0	250
CADM-G	IMAGENET 64x64	-	4.0	250	250
CADM-COMPG	IMAGENET 64x64	5.0, 6.0	4.0	50	250
<b>TABLE 8</b>					
EDM2-INTG	IMAGENET 256x256	-	2.0	6	32
EDM2-COMPCFG $ G  = 6$	IMAGENET 512x512	2.5	1.6	6	32
EDM2-COMPCFG $ G  = 5$	IMAGENET 512x512	2.5	1.7	5	32
EDM2-COMPCFG $ G  = 4$	IMAGENET 512x512	2.5	1.7	4	32

847

848


Table 11: All hyper-parameters required for reproducing the results.

849

850

851

852



853

854

855

856

857

858

859

860

861

862

Figure 9: IS-FID curve of ImageNet 256x256. The performance shows that CompG comes with high and stable IS with improves in FID more significantly than vanilla G given IS increases..

Model	$ G  (\downarrow)$	GPU hrs ( $\downarrow$ )	ZFID ( $\downarrow$ )
<b>MSCOCO 64x64</b>			
GLIDE-G	250	34.04	24.78
<b>GLIDE-CompG</b>	<b>25</b>	<b>20.93</b>	<b>24.5</b>
<b>MSCOCO 256x256</b>			
GLIDE-G	250	66.84	34.78
<b>GLIDE-CompG</b>	<b>35</b>	<b>37.55</b>	<b>33.12</b>

Table 12: Applying CompG on text-to-image GLIDE classifier-based guidance on MSCoco datasets.

Model	$ G  (\downarrow)$	GPU hours ( $\downarrow$ )	FID ( $\downarrow$ )	sFID ( $\downarrow$ )	Prec ( $\uparrow$ )	Rec ( $\uparrow$ )
<b>ImageNet 64x64</b>						
BigGAN	-	-	4.06	3.96	0.79	0.48
IDDPM	0	28.32	2.90	3.78	0.73	0.62
CADM (No guidance)	0	26.64	2.07	4.29	0.73	0.63
<b>CADM-<math>\bar{G}</math></b>	250	53.52	2.47	4.88	<b>0.80</b>	0.57
<b>CADM-CompG</b>	<b>50</b>	<b>32.22</b>	<b>1.91</b>	<b>4.57</b>	0.77	<b>0.61</b>
<b>CADM-CFG</b>	250	54.97	1.89	4.45	<b>0.77</b>	0.60
<b>CADM-CompCFG</b>	<b>25</b>	<b>29.29</b>	<b>1.84</b>	<b>4.38</b>	<b>0.77</b>	<b>0.61</b>
<b>ImageNet 128x128</b>						
BigGAN	-	-	6.02	7.18	0.86	0.35
LOGAN	-	-	3.36	-	-	-
CADM (No guidance)	0	61.60	6.14	4.96	0.69	0.65
<b>CADM-<math>\bar{G}</math></b>	250	94.06	2.95	5.45	<b>0.81</b>	0.54
<b>CADM-CompG</b>	<b>50</b>	<b>66.19</b>	<b>2.86</b>	<b>5.29</b>	0.79	<b>0.58</b>
<b>ImageNet 256x256</b>						
BigGAN	-	-	7.03	7.29	0.87	0.27
DCTrans	-	-	36.51	8.24	0.36	<b>0.67</b>
VQ-VAE-2	-	-	31.11	17.38	0.36	0.57
IDDPM	-	-	12.26	5.42	0.70	0.62
CADM (No guidance)	0	240.33	10.94	6.02	0.69	0.63
<b>CADM-<math>\bar{G}</math></b>	250	336.05	4.58	<b>5.21</b>	0.81	0.51
<b>CADM-CompG</b>	<b>50</b>	<b>259.25</b>	<b>4.52</b>	5.29	<b>0.82</b>	<b>0.51</b>
<b>DiT-CFG</b>	250	75.04	2.25	<b>4.56</b>	0.82	0.58
<b>DiT-CompCFG</b>	<b>22</b>	<b>42.20</b>	<b>2.19</b>	4.74	<b>0.82</b>	<b>0.60</b>

Table 13: We show full results of the model compared to other models not related to guidance.

the conflict between denoising signals and guidance signals, similar to Wang et al. (2024); Dinh et al. (2023b) identifies adversarial features and mitigates them by reducing uncertainty; and Zheng et al. (2022) tackles the gradient vanishing issue in classifier guidance by adapting the guidance weight. While all of these methods are training-free, our proposed CompG is the first to focus on reducing guidance-related computational costs by identifying and eliminating redundant guidance steps during sampling. Our results demonstrate that CompG achieves the best FID while maintaining the lowest running time. The comparative results are presented in Figure 14.

Model	$ G  (\downarrow)$	GPU hours ( $\downarrow$ )	FID ( $\downarrow$ )	sFID ( $\downarrow$ )	Prec ( $\uparrow$ )	Rec ( $\uparrow$ )
<b>ImageNet 64x64</b>						
CADM (No guidance)	0	26.64	2.07	4.29	0.73	0.63
<b>CADM-<math>\bar{G}</math></b>	250	53.52	2.47	4.88	<b>0.80</b>	0.57
CADM-ProG Dinh et al. (2023b)	250	53.60	1.87	4.33	0.77	0.60
CADM-PxP Dinh et al. (2023a)	250	54.32	1.84	<b>3.97</b>	0.76	0.60
CADM-EDS Zheng et al. (2022)	250	53.23	1.85	4.36	0.76	<b>0.61</b>
<b>CADM-CompG</b>	<b>50</b>	<b>32.22</b>	<b>1.82</b>	4.31	0.76	<b>0.61</b>
<b>CADM-CompCFG</b>	<b>25</b>	<b>29.29</b>	<b>1.84</b>	4.38	0.77	<b>0.61</b>

Table 14: Comparing CompG and CompCFG with other variants Dinh et al. (2023a;b); Zheng et al. (2022) of classifier guidance on conditional diffusion model ADM Dhariwal &amp; Nichol (2021)

Guidance	On-samp.	Off-samp.	Resnet	FID
Vanilla	90.8	62.5	34.17	2.47
Early Stopping	63.05	55.22	33.55	2.21
CompG (ours)	<b>91.2</b>	<b>64.2</b>	<b>34.93</b>	<b>1.82</b>

Table 15: *Model-fitting on ImageNet64x64 samples. ES suffers from the forgetting problem and has low performance. CompG achieves higher both on on-sampling and off-sampling acc.*

## F MATHEMATICAL DETAILS

### Proof of Theorem 1

*Proof.* Given real data  $\mathbf{x}_0$ , at timestep  $t$  we have  $\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$ . On the other hand, the prediction of real data has the form  $\tilde{\mathbf{x}}_0^{(t)} = \frac{\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(\mathbf{x}_t, t)}{\sqrt{\bar{\alpha}_t}}$ , replace  $\mathbf{x}_t$  with  $\mathbf{x}_0$  and  $\epsilon$  we have  $\tilde{\mathbf{x}}_0^{(t)} = \mathbf{x}_0 + \frac{\sqrt{1 - \bar{\alpha}_t} (\epsilon - \epsilon_\theta(\mathbf{x}_t, t))}{\sqrt{\bar{\alpha}_t}}$ . Thus,  $\|\tilde{\mathbf{x}}_0^{(t)} - \mathbf{x}_0\| = \frac{1 - \bar{\alpha}_t \|\epsilon - \epsilon_\theta(\mathbf{x}_t, t)\|}{\bar{\alpha}_t}$   $\square$

If we further assume that  $q(\mathbf{x}_0)$  has a form of Normal Distribution, we would have the final objective as  $D_{KL}(q(\mathbf{x}_0) || p_\theta(\tilde{\mathbf{x}}_0 | \mathbf{x}_t))$ . Since  $q(\mathbf{x}_0)$  has the form of Gaussian, we can have the minimization of  $\|\tilde{\mathbf{x}}_0^{(t)} - \mathbf{x}_0\|$  would result in the minimization of  $\|q(\tilde{\mathbf{x}}_0) - q(\mathbf{x}_0)\| = \|\frac{q(\tilde{\mathbf{x}}_0)q(\mathbf{x}_t | \tilde{\mathbf{x}}_0)}{q(\mathbf{x}_t)} - q(\mathbf{x}_0)\|$  since  $\tilde{\mathbf{x}}_0 \sim p_\theta(\tilde{\mathbf{x}}_0 | \mathbf{x}_t)$  with a deterministic forward of  $\mathbf{x}_t$  to  $\epsilon_\theta$ , we have  $q(\tilde{\mathbf{x}}_0) \approx \frac{q(\tilde{\mathbf{x}}_0)q(\mathbf{x}_t | \tilde{\mathbf{x}}_0)}{q(\mathbf{x}_t)} = p_\theta(\tilde{\mathbf{x}}_0 | \mathbf{x}_t)$ . Assume we have two density functions:  $p(\mathbf{x})$  and  $q(\mathbf{x})$ . The KL divergence between these two has the form:

$$\int_0^1 p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})} = \int_0^1 p(\mathbf{x}) \log(p(\mathbf{x})) - p(\mathbf{x}) \log(q(\mathbf{x})) d\mathbf{x} \quad (22)$$

$$= \int_0^1 p(\mathbf{x}) \log(p(\mathbf{x})) d\mathbf{x} - \quad (23)$$

$$\int_0^1 p(\mathbf{x}) \log(p(\mathbf{x})) + p(\mathbf{x}) \log\left(\left(\frac{p(\mathbf{x})}{q(\mathbf{x})} - 1\right) + 1\right) d\mathbf{x} \quad (24)$$

$$= \int_0^1 -p(\mathbf{x}) \log\left(\left(\frac{q(\mathbf{x})}{p(\mathbf{x})} - 1\right) + 1\right) d\mathbf{x} \quad (25)$$

$$= \int_0^1 -(q(\mathbf{x}) - p(\mathbf{x})) + (q(\mathbf{x}) - p(\mathbf{x}))^2 \left(\frac{1}{p(\mathbf{x})} - \frac{1}{q(\mathbf{x})}\right) d\mathbf{x} \quad (26)$$

$$\leq \int_0^1 (q(\mathbf{x}) - p(\mathbf{x}))^2 \left(\frac{1}{p(\mathbf{x})} - \frac{1}{q(\mathbf{x})}\right) d\mathbf{x} \quad (26)$$

$$\leq \int_0^1 (q(\mathbf{x}) - p(\mathbf{x}))^2 \left(\frac{1}{a} - \frac{1}{b}\right) d\mathbf{x} = \frac{b-a}{ab} \|p - q\| \quad (27)$$

Thus  $D_{KL}(p(\mathbf{x}) || q(\mathbf{x})) \leq \frac{b-a}{ab} \|p - q\|$

Base on this bound we would have the minimization of  $\|p_\theta(\tilde{\mathbf{x}}_0 | \mathbf{x}_t) - q(\mathbf{x}_0)\|$  is equivalent to the minimization of  $D_{KL}(q(\mathbf{x}_0) || p_\theta(\tilde{\mathbf{x}}_0 | \mathbf{x}_t))$ .

### Proof of first property of eq. 13




Figure 10: Gradient magnitude difference measured at two consecutive steps

*Proof.* Let  $k_1 < k_2$  and  $k_1, k_2 \in [1; +\infty]$ , with  $\frac{T}{|G|^k} i^k = T(\frac{i}{|G|})^k$  and  $\frac{i}{|G|} < 1$ , we have:

$$(\frac{i}{|G|})^{k_1} \geq (\frac{i}{|G|})^{k_2} \quad (28)$$

$$\Leftrightarrow T(\frac{i}{|G|})^{k_1} \geq T(\frac{i}{|G|})^{k_2} \quad (29)$$

$$\Leftrightarrow \lfloor T(\frac{i}{|G|})^{k_1} \rfloor \geq \lfloor T(\frac{i}{|G|})^{k_2} \rfloor \quad (30)$$

$$\Leftrightarrow T - \lfloor T(\frac{i}{|G|})^{k_1} \rfloor \leq T - \lfloor T(\frac{i}{|G|})^{k_2} \rfloor \quad (31)$$

As a result,  $G_i^{(k_1)} \leq G_i^{(k_2)} \forall k_1, k_2 \geq 1$  and  $k_1 < k_2$ . With  $k_2 \rightarrow +\infty$ ,  $G_i^{(k_2)}$  is bounded by T. This means that larger  $k$  values would result in the distribution of the timesteps toward the early stage of the sampling process.  $\square$

### Proof of first property of eq. 13

*Proof.* Similar to previous proof we have  $G_i^{(k_1)} \leq G_i^{(k_2)} \forall k_1, k_2 \geq 1$  and  $k_1 < k_2$ . This also mean that  $G_i^{(k_1)} > G_i^{(1)}$ ,  $\forall 0 \leq k_1 < 1$  and if  $k_1 \rightarrow 0$  then  $G_i^{(k_1)} \rightarrow 0$ , hence all the  $g_i \in G^{(k_1)}$  is bounded by 0. As a result, by adjusting  $k$  toward 0, we would have the distribution of guidance steps toward the later stage of the sampling process.  $\square$

## G GRADIENT MAGNITUDE DIFFERENCE BETWEEN TWO CONSECUTIVE SAMPLING STEPS

In this section, we analyze the variation in the classification gradient throughout the sampling process, particularly its significant fluctuations during the early stages. To investigate this, we generate 32 images from the ImageNet64 dataset using ADM-G (Dhariwal & Nichol (2021)). The guidance classifier employed in this process is the noise-aware classifier trained within ADM-G. Our observations, illustrated in Figure 10, highlight how the classification gradient behaves over time, providing insights into its impact on the sampling process and model performance.

## H ADDITIONAL QUALITATIVE RESULTS

Due to space limitations in the main paper, we present qualitative results in this supplementary material. Figures 11, 13, 14, 15, and 16 provide additional comparisons with the vanilla baseline, while Figures 17 and 18 showcase high-quality images generated by DiT models combined with CompG.




Figure 11: Qualitative results on ImageNet256x256. Left: Vanilla guidance applied at all timesteps. Right: Compact Guidance applied at 50 of 250 timesteps. Compact Guidance corrects misclassification by the on-sampling classifier, preventing out-of-class image generation and restoring accurate class information. More qualitative results are shown in Appendix H




Figure 12: Trade-off: Running time versus performance. We measure the compact rate as  $\frac{T}{|G|}$ . In (a), IS decreases with increasing compact rate, while FID and Recall improve. However, when the rate exceeds 10, FID begins to rise. This suggests that increased diversity from more features initially enhances Recall and FID, but excessive diversity degrades image quality.

1080  
1081  
1082  
1083  
1084  
1085  
1086  
1087  
1088  
1089  
1090  
1091  
1092  
1093  
1094  
1095  
1096  
1097  
1098  
1099  
1100  
1101  
1102  
1103

Quiet forest  
path surrounded  
by tall trees.

Quiet forest  
path surrounded  
by tall trees.




Figure 13: *Stable Diffusion with classifier-free guidance. The left figure is the vanilla classifier-free guidance with application on all 50 timesteps. Our proposed Compress Guidance method is the right figure, where we only apply guidance on 10 over 50 steps. The output shows our methods' superiority over classifier-free guidance regarding image quality, quantitative performance and efficiency.*

1134

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

Serene mountain landscape with a clear sky

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

1162

A white plate with breakfast foods on it

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172


1173

1174

1175

1176

1177



StableDiffusion

(ours)

Figure 14: *Stable Diffusion with classifier-free guidance*. The left figure is the vanilla classifier-free guidance with application on all 50 timesteps. Our proposed Compress Guidance method is the right figure, where we only apply guidance on 10 over 50 steps. The output shows our methods' superiority over classifier-free guidance regarding image quality, quantitative performance and efficiency.

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

1188

1189

1190

1191

1192

1193

1194

1195

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

1208

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

Flowers are  
arranged in a vase  
sitting on a table.

A plate with food  
on it, a fork and  
some kind of drink



StableDiffusion

(ours)

Figure 15: *Stable Diffusion with classifier-free guidance. The left figure is the vanilla classifier-free guidance with application on all 50 timesteps. Our proposed Compress Guidance method is the right figure, where we only apply guidance on 10 over 50 steps. The output shows our methods' superiority over classifier-free guidance regarding image quality, quantitative performance and efficiency.*

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241




Figure 16: *Qualitative comparison between ADM-G and ADM-CompG. The images generated by ADM-G and ADM-CompG are put side by side. On the left side is ADM-G and on the right side is ADM-CompG.*




Figure 17: *Images generated by DiT-CompCFG. From top to bottom classes goldfish, Welsh springer spaniel, Pembroke Welsh corgi, Cardigan Welsh corgi.*




Figure 18: *Images generated by DiT-CompCFG. From top to bottom classes redfox, kitfox, Arctic fox, tabby cat.*

1394  
1395  
1396  
1397  
1398  
1399  
1400  
1401  
1402  
1403