

Entropy-driven Sampling and Training Scheme for Conditional Diffusion Generation

Guangcong Zheng^{1*}, Shengming Li^{1*}, Hui Wang¹, Taiping Yao², Yang Chen²,
Shouhong Ding², and Xi Li^{1,3,4**}

¹ College of Computer Science & Technology, Zhejiang University
{guangcongzheng, shengming22, wanghui_17, xilizju}@zju.edu.cn

² Youtu Lab, Tencent, China

{taipingyao, wizyangchen, ericshding}@tencent.com

³ Shanghai Institute for Advanced Study, Zhejiang University

⁴ Shanghai AI Laboratory

Abstract. Denoising Diffusion Probabilistic Model (DDPM) is able to make flexible conditional image generation from prior noise to real data, by introducing an independent noise-aware classifier to provide conditional gradient guidance at each time step of denoising process. However, due to the ability of the classifier to easily discriminate an incompletely generated image only with high-level structure, the gradient, which is a kind of class information guidance, tends to vanish early, leading to the collapse from conditional generation process into the unconditional process. To address this problem, we propose two simple but effective approaches from two perspectives. For sampling procedure, we introduce the entropy of predicted distribution as the measure of guidance vanishing level and propose an entropy-aware scaling method to adaptively recover the conditional semantic guidance. For the training stage, we propose the entropy-aware optimization objectives to alleviate the overconfident prediction for noisy data. On ImageNet1000 256×256, with our proposed sampling scheme and trained classifier, the pretrained conditional and unconditional DDPM model can achieve 10.89% (4.59 to 4.09) and 43.5% (12 to 6.78) FID improvement, respectively. Code is available at <https://github.com/ZGCTroy/ED-DPM>.

Keywords: Denoising Diffusion Probabilistic Model · Conditional Generation · Distribution Entropy · Gradient Vanishing

1 Introduction

Conditional image generation, usually class conditional, aims to generate the specific class of high-quality images. There are many generative models that are able to make high-quality conditional generation based on the joint training scheme, such as Generative Adversial Networks (GAN) [7,2] or Variational Autoencoder (VAE) [27,38]. However, when the condition requirement is changed, the generative models will be retrained, which is very inconvenient.

* The first two authors contributed equally to this paper.

** The corresponding author is Xi Li.

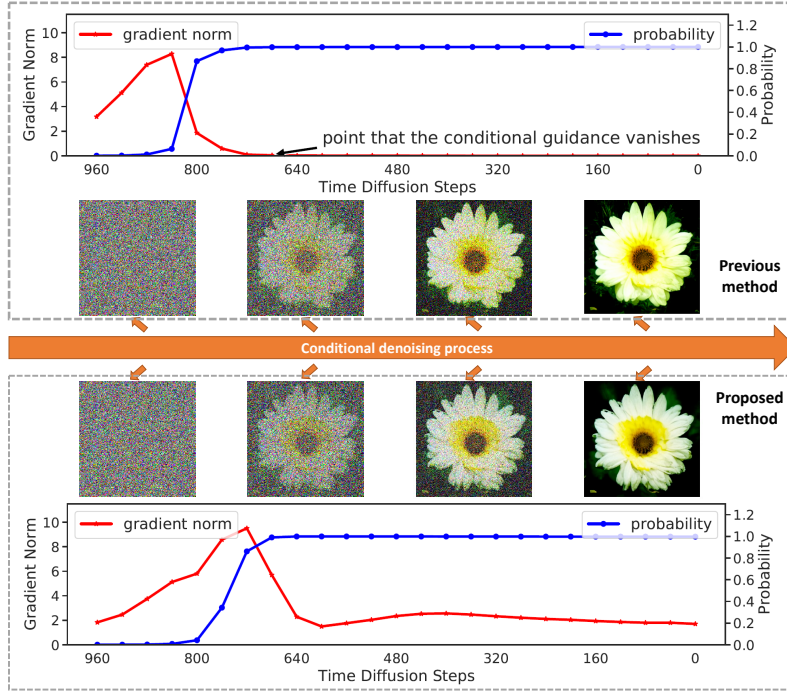


Fig. 1: The visualization of denoising sampling process. The classifier gradient, a kind of class information in conditional generation, quickly converge to 0 in the previous method. It will lead to the collapse from conditional generation to unconditional generation, while our method recovers the gradient guidance and succeed to generate fine-grained features in the subsequent iterations.

Denoising Diffusion Probabilistic Model (DDPM) [10,23,33] is a class of iterative generation models, which has achieved remarkable performance in unconditional image generation recently. The flexibility of DDPM [6,32] is that it can be easily extended to conditional variants by introducing an independent noise-aware classifier. Recent researches modeled the prior denoising distribution by training an unconditional DDPM, following the training scheme of Denoising Score Matching [39], and computed likelihood score by backwaring the classifier gradient. Dhariwal *et al.* [6] further proposed fixed scaling factor to improve the predicted probability of generated samples for DDPM, achieving superior performance than GAN on several image generation benchmarks. In conditional generation process of DDPM, by backwaring the gradient of classification probability to image, the classifier provides high-level semantic information in the early stage of iterations, and gradually strengthens fine-grained features in the subsequent iterations, both of which are indispensable. However, there exists a huge gap between discriminating the class of a image and generating a specific class of image with fine-grained textures. As shown in Fig. 1, the predicted

distribution of the classifier for noisy images tends to quickly converge to the desired class distribution, which is one-hot distribution, leading to the early vanishing of conditional gradient guidance. This is because that the incompletely generated image, which is still a noisy image and lacks fine-grained features, can be easily classified in the middle of denoising process. In this way, the image is considered to have been completely generated and will no longer be guided by classifier gradient containing class information. As a result, the conditional generation process will degrade into an unconditional generation process in the later stage.

Therefore, our motivation is to enable the classifier to continuously give conditional guidance throughout the entire denoising process. We propose two simple but effective schemes from the procedure of sampling and the design of classifier training.

From the perspective of sampling procedure, we focus on how to detect the gradient vanishing and rescale the gradient to avoid the existence of gradient vanishing or recover the gradient when the vanishing does happen. We propose **Entropy-Driven conditional Sampling (EDS)** method, which is able to adaptively measure the level of gradient vanishing and rescale the gradient guidance to a appropriate level. In design of training classifier, we propose **Entropy-Constraint Training (ECT)**, which will penalize the classifier when it gives a overconfident classification probability to a generated noisy image, thus constraining the classifier to provide more gentle guidance.

Our contributions can be summarized as follows:

- We are the first to discover the problem of vanishing gradient guidance for DDPM-based conditional generation methods, and point out that category information guidance should be continuously provided throughout the entire generation process.
- We propose EDS to alleviate the vanishing guidance by dynamically measuring and rescaling gradient guidance. At the classifier training stage, to alleviate the vanishing gradient caused by one-hot label supervision, we utilize discrete uniform distribution to build an entropy-aware optimization term, which is **Entropy-Constraint Training** scheme (ECT).
- We conduct experiments on ImageNet1000 and achieve state-of-the-art FID (Fréchet Inception Distance) results at various resolutions. On ImageNet1000 256×256, with our proposed sampling scheme and trained classifier, the pretrained conditional and unconditional DDPM model can achieve 10.89% (4.59 to 4.09) and 43.5% (12 to 6.78) FID (Fréchet Inception Distance) improvement, respectively.

2 Related Work

2.1 Denoising Diffusion Probabilistic Model

Denoising diffusion probabilistic models (DDPM) is the latest generation model which achieve superior generation performance than traditional generative models

like Generative Adversarial Networks (GAN) [8,7,2], Variational Autoencoders (VAE) [38,25,27] on several benchmarks about unconditional generation. Its key idea is to model the diffusion process based on total T time steps, which adds noise gradually to the clean data, and its reverse process, which denoises the white noise into the clean sample. Accordingly, its diffusion process is modeled as a fixed Markov Chain and its transition kernel is formulated as:

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) := \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}), \quad (1)$$

where β_1, \dots, β_T are the fixed variance parameters, which are not learnable. According to the transition kernel above, when the clean data \mathbf{x}_0 is given, the noisy data \mathbf{x}_t can be sampled with a closed-form distribution:

$$q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t}\mathbf{x}_0, (1 - \bar{\alpha}_t)\mathbf{I}), \quad (2)$$

where $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$. When t is close to T , x_T can be approximated as a Gaussian distribution.

Given a prior diffusion process, DDPM aims to model its reverse process to sample from the data distribution. The optimization objective of the reverse transition can be derived from a variational bound [13]. Thus, DDPM introduced the variational solution [13] and assumed that its reverse transition kernel also subjects to Gaussian distribution, which is the same as the diffusion process. In this way, the generation process parameterized the mean of the Gaussian transition distribution and fixed its variance as follows:

$$\begin{aligned} p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) &= \mathcal{N}(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t), \sigma_t^2\mathbf{I}) \\ \mu_\theta(\mathbf{x}_t) &= \frac{1}{\sqrt{\alpha_t}}(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}}\epsilon_\theta(\mathbf{x}_t)), \end{aligned} \quad (3)$$

where $\epsilon_\theta(\mathbf{x}_t)$ is a noise estimator modeled by a neural network. The variance is designed as the hyperparameters for training diffusion models.

Recently, there were some researches [36,37,34,6] which indicated that the noise estimator can be regarded as an approximation of score function so that the sampling process is equivalent to solving a stochastic differential equation. Based on above, Song et al [34] proposed an effective sampling process that shares the same training objectives as DDPM and its corresponding denoising process, which is also the iteration solution to solve the stochastic differential equation, is designed as follows:

$$\mathbf{x}_{t-1} = \sqrt{\alpha_{t-1}}f_\theta(\mathbf{x}_t, t) + \sqrt{1 - \alpha_{t-1} - \sigma_t^2}\epsilon_\theta(\mathbf{x}_t) + \sigma_t\mathbf{z}, \quad (4)$$

where $f_\theta(\mathbf{x}_t, t)$ is the prediction of clean data x_0 when noisy data \mathbf{x}_t is observed and noise prediction $\epsilon_\theta(\mathbf{x}_t)$ is given. The $f_\theta(\mathbf{x}_t, t)$ can be expressed as:

$$f_\theta(\mathbf{x}_t, t) = \frac{\mathbf{x}_t - \sqrt{1 - \alpha_t}\epsilon_\theta(\mathbf{x}_t)}{\sqrt{\alpha_t}}, \quad (5)$$

When the variance σ_t is set to 0, the sampling process becomes deterministic. At the same time, the non-Markovian diffusion process [34] allows the generation quality to remain unchanged within fewer denoising steps, which is called the DDIM sampling process.

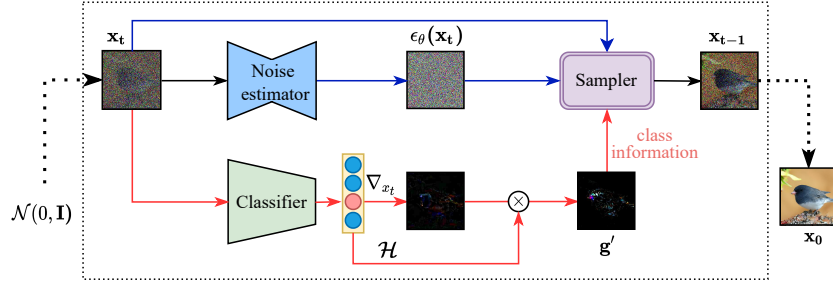


Fig. 2: Pipeline for Entropy-driven Sampling process. Sampler represents a class of iteration method (DDPM or DDIM), which is non-parametric. All models are pretrained without gradient updating in the sampling process.

2.2 Conditional Image Generation

Conditional image generation aims to generate samples with desired condition information. The condition can be extended to multi-modal information, such as class [20,25,27], text [41,26,40], and low-resolution image [3]. Most previous work modeled this by the joint training scheme with both condition and random noise, utilizing generative models like GAN or VAE.

Considering that Denoising Diffusion Probabilistic Model (DDPM) has made remarkable progress in recent years [34,12,17,11,5,35,21,24], there raised many researches [6,14,4,30,31,16,18] which applied DDPM to conditional generation. Due to the ideal theoretical properties of DDPM, it can be extended flexibly to conditional variants using the Bayes theorem without retraining, which is similar to score-based generative models [36,37,19]. In this paper, we focus on the class-conditional generation task, in which the condition is represented by a class discrete distribution, and design a more effective sampling and training scheme to improve the generation quality for DDPM, further exploring its potential in image generation aspects.

3 Proposed Methods

We start by introducing the conditional generation process for diffusion models with classifier guidance (Sect. 3.1). For ease of description, we will firstly introduce our dynamic scaling technology to recover gradient guidance adaptively in the sampling process and its motivation (Sect. 3.2). Then, we describe the entropy-aware optimization loss for alleviating vanishing conditional guidance from training perspective (Sect. 3.3), utilizing the uniform distribution, which is a more dense distribution.

3.1 Conditional Diffusion Generation

The goal of conditional image generation is to model probability density $p(\mathbf{x}|\mathbf{y})$. To be specific, in class-conditional image generation, \mathbf{y} represents the desired class

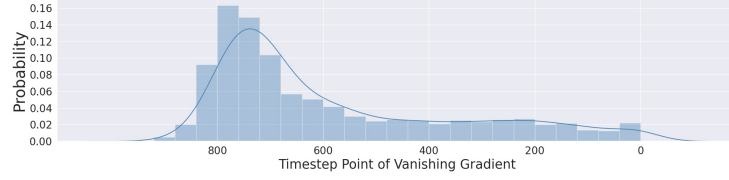


Fig. 3: The various gradient vanishing points for different samples. We randomly generate 5000 samples and define the time point, whose gradient norm that is smaller than 0.15 as the gradient vanishing point. Premature vanishing of gradients occurs in almost every sample.

label, which tends to be one-hot distribution, and \mathbf{x} represents the image sample. For diffusion models, it can be implemented by introducing an independent classifier, as shown in Fig. 2. Specifically, the goal is converted into modeling the conditional transition distribution $p(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{y})$, derived from Markov chain sampling scheme:

$$p_\varphi(\mathbf{x}_0|\mathbf{y}) = \int p_\varphi(\mathbf{x}_{0:T}|\mathbf{y})d\mathbf{x}_{1:T}, \quad (6)$$

$$p_\varphi(\mathbf{x}_{0:T}|\mathbf{y}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\varphi(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{y}),$$

where φ represents the model and T is the total length of Markov chain, which tends to be large. Then, we decompose the conditional transition distribution into two independent terms:

$$p_\varphi(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{y}) = Z p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) p_\phi(\mathbf{y}|\mathbf{x}_t), \quad (7)$$

where Z is a normalizing constant independent from \mathbf{x}_{t-1} and φ can be seen as the combination of models θ and ϕ , which is proven theoretically [32,6]. Furthermore, the log density of Eq. (7) can be approximated as a Gaussian distribution [6]:

$$\log(p_\varphi(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{y})) \approx \log p(\mathbf{z}) + \log Z$$

$$\mathbf{z} \sim \mathcal{N}\left(\frac{1}{\sqrt{\alpha_t}}(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}}\boldsymbol{\epsilon}_\theta(\mathbf{x}_t)) + \sigma_t^2 \mathbf{g}, \sigma_t^2 \mathbf{I}\right), \quad (8)$$

where $\mathbf{g} = s \nabla_{\mathbf{x}_t} \log p_\phi(\mathbf{y}|\mathbf{x}_t)$ and s is the gradient scale. \mathbf{g} can be derived by backwarding the gradient of a pretrained classifier on noisy data \mathbf{x}_t . Usually, s is set to a constant [6] to improve the predicted probability. $\boldsymbol{\epsilon}_\theta(\mathbf{x}_t)$ is the pretrained noise estimator for diffusion models and σ_t^2 is the hyperparameters for unconditional DDPM to control the variance.

In this way, unconditional diffusion models with parameterized noise estimator $\boldsymbol{\epsilon}_\theta$ can be extended to conditional generative models by introducing the condition-aware guidance $\nabla_{\mathbf{x}_t} \log p_\phi(\mathbf{y}|\mathbf{x}_t)$, as shown in Fig. 2. Specifically, we start with the prior noise $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$, and utilize DDPM-based (Eq. 8) sampler to make

transition iteratively to generate the samples conditioned on desired class \mathbf{y} . The sampler can also be extended into DDIM iteration method (Eq. 3). The extension details for conditional process with classifier guidance can refer to Dhariwal *et al.*[37,6].

3.2 Entropy-driven Conditional Sampling

During the conditional generation process, we observe that the guidance provided from noise-aware classifier tends to vanish prematurely, which could be attributed to the discrepancy between generative pattern and discriminative pattern. For example, the noisy samples with high-level semantic information, such as contour or the color, may be guided iteratively by the classifier with nearly one-hot distribution, in which situation the gradient guidance tends to be weak or vanish, while the samples still lack the condition-aware semantic details. In this way, the condition-aware textures of generated samples are guided by unconditional denoising process (Eq. 3).

An intuitive solution is to manually select a time step during the sampling process, after which the semantic details tend to vanish in most instances, and rescale the conditional gradient by an empirical constant after the selected time step. However, since the stochasticity in the generation process of diffusion models [4], the denoising trajectories for generated samples would differ from each other. It will lead to the various initial vanishing points, as shown in Fig. 3. At the same time, considering the learning bias of classifier for different conditional classes, the level of recovery factor for each class may also differ. In summary, the effective scaling factor can be related to the current time step, the class condition, and the stochasticity of generation process. The experiment design and results about more intuitive approaches can be seen in Sect. 4.3.

Algorithm 1 Entropy-driven sampling scheme (DDPM/DDIM)

Require: a pretrained diffusion model $\epsilon_\theta(\mathbf{x}_t)$, classifier $p_\phi(\hat{\mathbf{y}}|\mathbf{x}_t)$, and desired class condition \mathbf{y}

```

1:  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
2: for  $t = T, \dots, 1$  do
3:    $s \leftarrow \gamma * \frac{\mathcal{H}(\mathcal{U}(\hat{\mathbf{y}}))}{\mathcal{H}(p_\phi(\hat{\mathbf{y}}|\mathbf{x}_t))}$  if EDS, else  $s \leftarrow \gamma$ 
4:   if use DDPM then
5:      $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  if  $t > 1$ , else  $\mathbf{z} \leftarrow \mathbf{0}$ 
6:      $\mathbf{g} \leftarrow s \cdot \nabla_{\mathbf{x}_t} \log p_\phi(\mathbf{y}|\mathbf{x}_t)$ ,
7:      $\mathbf{x}_{t-1} \leftarrow \frac{1}{\sqrt{\alpha_t}} \left( \mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\alpha_t}} \epsilon_\theta(\mathbf{x}_t) \right) + \sigma_t^2 \mathbf{g} + \sigma_t \mathbf{z}$ 
8:   else if use DDIM then
9:      $\hat{\epsilon} \leftarrow \epsilon_\theta(\mathbf{x}_t) - s \cdot \nabla_{\mathbf{x}_t} \log p_\phi(\mathbf{y}|\mathbf{x}_t)$ ,
10:     $\mathbf{x}_{t-1} \leftarrow \sqrt{\alpha_{t-1}} \left( \frac{\mathbf{x}_t - \sqrt{1-\alpha_t} \hat{\epsilon}}{\sqrt{\alpha_t}} \right) + \sqrt{1-\alpha_{t-1}} \hat{\epsilon}$ 
11:   end if
12: end for
13: return  $\mathbf{x}_0$ 
```

Motivated from above, we propose an dynamic scaling technology for the conditional diffusion generation to recover the semantic details adaptively for each sample. We noticed that, when time step is close to T , the predicted distribution tends to be dense, which can be nearly approximated as uniform distribution. The reason is that the noisy data derived from Eq. 1 can be approximated as random noise, $\mathcal{N}(0, \mathbf{I})$, in which state the gradient guidance is obvious. As time step declines to 0 inversely, the noise hidden in sample will be removed gradually and the predicted distribution tends to be close to one-hot, in which case the classifier gradient is invalid to provide semantic details for generation. Statistically, entropy can represent the sparsity of the predicted distribution, inspiring us to take it into consideration:

$$\begin{aligned}\mathcal{H}(p_\phi(\tilde{\mathbf{y}}|\mathbf{x}_t)) &= -\mathbb{E}_{\tilde{\mathbf{y}}|\mathbf{x}_t} \log p_\phi(\tilde{\mathbf{y}}|x_t) \\ &= -\sum_{i=1}^{|Y|} p_\phi(\tilde{\mathbf{y}}_i|\mathbf{x}_t) \log p_\phi(\tilde{\mathbf{y}}_i|\mathbf{x}_t),\end{aligned}\tag{9}$$

where $p_\phi(\tilde{\mathbf{y}}|\mathbf{x}_t)$ represents the predicted distribution from classifier and Y represents the set of all class conditions.

In this paper, we utilize $\mathcal{H}(p_\phi(\tilde{\mathbf{y}}|\mathbf{x}_t))$ to adaptively fit various gradient vanishing time step. Furthermore, the entropy $\mathcal{H}(p_\phi(\tilde{\mathbf{y}}|\mathbf{x}_t))$ can also capture bias caused by different conditions, due to its sample-aware rescaling effect. Thus, when we sample from a pretrained noise estimator ϵ_θ in Eq. 8 conditionally, we reformulate the gradient term \mathbf{g} as following, which is shown in Fig. 2:

$$\begin{aligned}\mathbf{g}' &= s(x_t, \phi) * \nabla_{\mathbf{x}_t} \log p_\phi(\mathbf{y}|\mathbf{x}_t), \\ s(x_t, \phi) &= \gamma * \frac{\mathcal{H}(\mathcal{U}(\tilde{\mathbf{y}}))}{\mathcal{H}(p_\phi(\tilde{\mathbf{y}}|\mathbf{x}_t))}\end{aligned}\tag{10}$$

where γ is a hyper-parameter to balance the guiding gradient and entropy-aware scaling factor $s(\mathbf{x}_t, \phi)$. In order to maintain the numerical range, we renormalize the entropy by its theoretical upper bound $\mathcal{H}(\mathcal{U}(\tilde{\mathbf{y}}))$, where $\mathcal{U}(\tilde{\mathbf{y}})$ represents the uniform distribution of class variable, so that the gradients are almost not rescaled when t is close to T .

3.3 Training Noise-aware Classifier with Entropy Constraint

From the training perspective, the vanishing gradient can be partly attributed to the label supervision pattern for noise-aware classifier. Since one-hot distribution is very sparse and is utilized to supervise the noisy data, the predicted distributions are inclined to converge to one-hot under noisy samples in the sampling process, so that the gradient guidances are too weak to generate condition-aware semantic details at sampling stage.

Specifically, given the dataset $(\mathbf{x}_0, \mathbf{y}) \sim \mathbb{D}$ and a prior diffusion process (Eq. 1), the classifier will be trained under noisy data \mathbf{x}_t to build the gradient field in Eq. 8 of each time step. To alleviate the weak guidance caused by the sparsity,

Algorithm 2 Entropy-constraint training scheme.

Require: training set \mathbb{D} , a neural classifier ϕ , training set \mathbb{D} , total time steps T

- 1: **repeat**
- 2: $(\mathbf{x}_0, \mathbf{y}) \leftarrow \text{sample from } \mathbb{D}$
- 3: $t \sim \mathcal{U}(\{1, \dots, T\})$
- 4: $\mathbf{x}_t \sim q(\mathbf{x}_t | \mathbf{x}_0)$ (Eq. 2)
- 5: $L_{CE} \leftarrow \mathbf{y} \log p_\phi(\tilde{\mathbf{y}} | \mathbf{x}_t)$
- 6: **if** use ECT **then**
- 7: $L_{ECT} \leftarrow -\mathcal{H}(p_\phi(\tilde{\mathbf{y}} | \mathbf{x}_t))$
- 8: Take gradient descent step on $\nabla_\phi(L_{CE} + \eta L_{ECT})$
- 9: **else**
- 10: Take gradient descent step on $\nabla_\phi L_{CE}$
- 11: **end if**
- 12: **until** converged

we utilize discrete uniform distribution, which is a dense distribution and has maximum entropy, as a perturbing distribution and introduce the optimization term at the training stage of the classifier to constrain the predicted distribution from the classifier as followed:

$$\begin{aligned}
\mathcal{L}_{ECT}(\mathbf{x}_t, \mathbf{y}) &= D_{KL}(p_\phi(\tilde{\mathbf{y}} | \mathbf{x}_t) || \mathcal{U}(\tilde{\mathbf{y}})) \\
&= \mathbb{E}_{\tilde{\mathbf{y}} | \mathbf{x}_t} \log p_\phi(\tilde{\mathbf{y}} | \mathbf{x}_t) - \mathbb{E}_{\tilde{\mathbf{y}} | \mathbf{x}_t} \log \mathcal{U}(\tilde{\mathbf{y}}) \\
&= -\mathcal{H}(p_\phi(\tilde{\mathbf{y}} | \mathbf{x}_t)) + \mathbf{C},
\end{aligned} \tag{11}$$

where \mathbf{C} is a constant term independent from the parameter ϕ . This loss term is equivalent to maximizing entropy of the predicted distribution $p(\tilde{\mathbf{y}} | \mathbf{x}_t)$. The whole training loss of guiding classifier is composed of the normal cross-entropy loss and entropy constraint training loss (ECT), which is formally given by:

$$\mathcal{L}_{tot}(\mathbf{x}_t, \mathbf{y}) = \mathcal{L}_{CE}(\mathbf{x}_t, \mathbf{y}) + \eta \mathcal{L}_{ECT}(\mathbf{x}_t, \mathbf{y}), \tag{12}$$

where η is a hyper-parameter to adjust the divergence about predicted label distribution and the uniform distribution.

Different from Entropy-driven Sampling, the proposed training scheme tries to alleviate the vanishing guidance by adjusting the gradient direction in sampling process, instead of the gradient scale. Thus, entropy-constrain training scheme can be complementary with entropy-driven sampling.

4 Experiments

In this section, we present experiments to verify the effectiveness and motivation of our proposed schemes. More visualization of generated samples and ablation experiments about hyperparameters can be found in supplementary materials.

4.1 Experiment Setup

Dataset We perform our experiments mainly on ImageNet dataset [28] at 256×256 resolutions. ImageNet contains 14,197,122 images with 1000 classes in total, which is a very challenging benchmark for conditional image generation.

Implementation Details. For verifying the effect of proposed schemes, we apply the neural network architecture, ablated diffusion model (ADM), proposed by Dhariwal et al [6]. ADM is mainly based on the UNet, with increased depth versus width, the number of attention heads, and rescaling residual connections with $\frac{1}{\sqrt{2}}$. It is also possible to train a conditional diffusion models. We call the conditional diffusion architecture [6] as CADM for short. Correspondingly, UADM means unconditional diffusion architecture. UADM-G and CADM-G additionally use noise-aware classifier guidance to perform conditional generation, separately.

Our training hyperparameters of noise-aware classifier, including batch size, total number of iterations, and decay rate, are kept the same as those of Dhariwal et al.[6] for fair comparison. We adopt the fixed linear variance schedule β_1, \dots, β_T [23,10] for prior noising process Eq. 1 and choose T as 1000. In this paper, we set η to 0.2 to keep a slight disturbance during training stage of classifier in all experiments. The selection details of γ can be seen in supplementary materials. All the experiments are conducted on 16 NVIDIA 3090s.

Evaluation Metrics. We select FID (Fréchet Inception Distance) [9] as our default evaluation metric, which is the most widely used metric for generation evaluation. FID [9] measures the KL divergence of two Gaussian distributions, which is computed by the real reference samples and the generated samples, in the feature space of Inception-V3. To capture more spatial relationships, sFID are proposed as a variant of FID, which is more sensitive to the consistent image distribution with high-level structures.

In addition, we apply several other metrics for more comprehensive evaluations. Inception Score (IS) was proposed [1], to measure the mutual information between input sample and the predicted class. Improved Precision and Recall metrics are proposed [15] for further evaluation of generative models. Precision is computed by estimating the proportion of generated samples that fall into real data manifold, measuring the sample fidelity. By contraries, recall is computed by estimating the proportion of real samples which fall into generated data manifold, measuring the sample diversity. Following Dhariwal et al. [6], we randomly generated **50k** images to compute all the above metrics based on **10k** real images when compared with previous methods. For consistent comparisons, we use the evaluation metrics for all methods based on the same codebase as Dhariwal et al.[6].

4.2 Comparison with State-of-the-art Methods

In this section, we show the comparison results of our proposed schemes with other SOTA methods. based on UADM and CADM architectures [6]. All results of other previous methods are cited from Dhariwal et al.[6]. As shown in Table 1,

Table 1: Comparison results with state-of-the-art generative models on ImageNet1000 256×256. Annotation '(25)' means the DDIM [34] sampling method with 25 steps. Otherwise, it means DDPM sampling method with 250 steps.

Method	FID ↓	sFID ↓	IS ↑	Prec ↑	Rec ↑	Prec+Rec ↑
DCTransformer [22]	36.51	8.24	-	0.36	0.67	1.03
VQ-VAE-2 [27]	31.11	17.38	-	0.36	0.57	0.93
IDDPM [23]	12.26	5.42	-	0.70	0.62	1.32
SR3 [29]	11.30	-	-	-	-	-
BigGAN-deep [2]	6.95	7.36	-	0.87	0.28	1.15
UADM-G(25 steps) [6]	14.21	8.53	83	0.7	0.46	1.16
UADM-G(25 steps)+ EDS+ECT	8.28	6.37	163.17	0.76	0.44	1.20
CADM-G(25 steps) [6]	5.44	5.32	194.48	0.81	0.49	1.30
CADM-G(25 steps)+ EDS+ECT	4.67	5.12	235.24	0.82	0.47	1.29
UADM-G [6]	12	10.4	95.41	0.76	0.44	1.20
UADM-G+ EDS+ECT	6.78	6.56	168.78	0.81	0.45	1.26
CADM [6]	10.94	6.02	100.98	0.69	0.63	1.32
CADM-G [6]	4.59	5.25	186.70	0.82	0.52	1.34
CADM-G+ EDS+ECT	4.09	5.07	221.57	0.83	0.50	1.33

we achieve the best results in terms of FID metric. Compared to UADM-G in ImageNet 256×256, our methods achieve relatively about 40% improvement on FID metric (from 14.21 to 8.28, from 12.0 to 6.78) based on both DDPM and DDIM sampling iteration methods, with comparable or even better precision and recall. For the CADM architecture, our proposed schemes still maintain a significant improvement margin on FID metric (from 5.44 to 4.67, from 4.59 to 4.09), and comparable precision and recall.

It is worth mentioning that our method based on the UADM architecture outperformed BigGAN-deep in terms of FID by 0.17 margin (6.78 vs 6.95), with no dependency on conditional architecture CADM, which has not been achieved in previous work [6,10,23].

4.3 Ablation Study

Effect of proposed schemes. To further verify the contribution of each component of proposed schemes, we conduct ablation experiments on both UADM and CADM architectures, as shown in Table 2. It can be concluded that EDS and ECT both improve the generation quality. Combining with two schemes, the generation results can be further improved, with more semantic details achieved from the improved direction (ECT) and scale (EDS) aspects of the guidance gradient.

Table 2: Ablation study of our proposed methods EDS and ECT on ADM-G under DDIM 25 steps and DDPM 250 steps on ImageNet1000 256×256.

Method	+ECT	+EDS	FID ↓	sFID ↓	IS ↑	Precision ↑	Recall ↑
UADM-G (25 steps)	✗	✗	14.21	8.53	83	0.7	0.46
	✗	✓	10.09	6.86	133.71	0.73	0.45
	✓	✗	12.21	8.14	100.95	0.74	0.44
	✓	✓	8.28	6.37	163.17	0.76	0.44
CADM-G (25 steps)	✗	✗	5.46	5.32	194.48	0.81	0.48
	✗	✓	4.82	5.04	218.97	0.80	0.50
	✓	✗	5.34	5.3	196.8	0.81	0.49
	✓	✓	4.67	5.12	235.24	0.82	0.48
UADM-G	✗	✗	12.0	10.4	95.41	0.76	0.44
	✗	✓	7.98	10.61	178.73	0.82	0.40
	✓	✗	10.79	10.49	117.56	0.79	0.41
	✓	✓	6.78	6.56	168.78	0.81	0.45
CADM-G	✗	✗	4.59	5.25	186.7	0.82	0.52
	✗	✓	4.01	5.15	217.25	0.82	0.52
	✓	✗	4.62	5.16	182.48	0.81	0.53
	✓	✓	4.09	5.07	221.57	0.83	0.50

Effect of entropy-driven sampling. In this part, we design various intuitive scaling methods and compare them under optimal hyperparameters with EDS. We select the scaling method with constant recovery factor for all time range [6] as our baseline. Intuitively, we can manually select vanishing time point through observing Fig. 3 and Fig. 1 i.e., 700, and finetune the constant rescaling factor to adjust the weak gradient for all generated samples, which we called Constant (range 0-700). The constant scale can be further adjusted to time-aware form, such like $T - t$. We call this method as Timestep-aware, which can dynamically adjust the scaling factor according to time step in sampling process. In this way, the sample-aware vanishing characteristic is ignored and the scale cannot fit the various vanishing level.

To verify that entropy-aware scaling could better fit the initial time steps of vanishing guidance, we design another approach which is based on norm of gradient map. Specifically, we empirically select a norm bound M for gradient. When the norm of gradient map is smaller than the threshold vanishing norm bound M , we regard that the gradient guidance is weak and need to be rescaled. Thus, s in Eq. 10 is rewritten as followed:

$$s = \begin{cases} 1, & \|\nabla_{\mathbf{x}_t} \log p_\phi(\mathbf{y}|\mathbf{x}_t)\|_2 < M \\ C, & \text{Otherwise} \end{cases} \quad (13)$$

The rescaling factor C is a large constant. Compared to the methods above, it can be experimentally verified that EDS not only fits the sample-aware initial

Table 3: Sample quality comparison with several intuitive sampling schemes. Generate 5,000 samples for more efficient comparisons.

Method	FID	IS	Precision	Recall
Baseline	20.17	84.53	0.71	0.59
Constant (range 0-700)	19.84	88.42	0.70	0.61
Timestep-aware	19.42	87.42	0.70	0.63
Gradient Norm	19.08	93.14	0.71	0.63
Entropy-driven	16.56	133.26	0.73	0.66

time steps of weak guidance, but also provides a reasonable rescaling factor for recovery, as shown in Table 3.

4.4 Qualitative Results

Gradient Map Visualization. We collected several gradient maps \mathbf{g} derived from the classifier in the previous sampling process and our EDS process, with an equal time interval. From Fig. 4, it can be observed that the classifier provides high-level semantic guidance at the beginning. Gradually, the classifier will provide condition-aware texture guidance. Compared with EDS, which can maintain guidance of semantic details throughout the denoising process for refined generation results, sampling scheme based on the fixed scaling factor lost a lot of condition-aware details (first line) or introduce unnatural details (third line) at the later sampling stage.

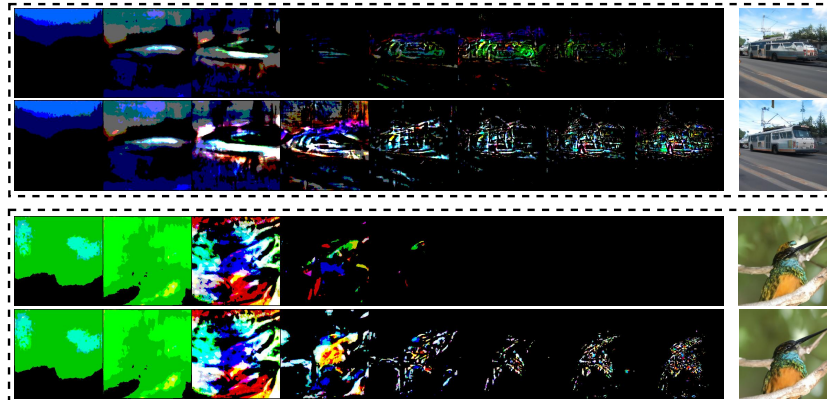


Fig. 4: Gradient map comparison between UADM-G [6] (the first line of each image box) and our UADM-G + EDS (the second line of each image box) on ImageNet1000 256×256, DDIM 25 steps.

Generation Results Visualization. We visualized various generated images conditioned on different classes and compared our final generated results with that of previous SOTA method [6] on UADM architecture. DDIM is adopted as the deterministic generation process to generate samples with the same initial noise. From Fig. 5, previous SOTA method cannot generate condition-aware semantic details such like beaks of birds Fig. 5b or petal texture Fig. 5i, while our method can generate more refined textures, with similar high-level structure.

5 Conclusion

In this paper, we proposed an entropy-aware scaling technology for the guiding classifier in sampling and constrain the predicted distribution in training. Experiments demonstrate that our methods can recover textures in generated samples, achieving state-of-the-art generation results.

Acknowledgements This work is supported in part by National Key Research and Development Program of China under Grant 2020AAA0107400, Zhejiang Provincial Natural Science Foundation of China under Grant LR19F020004, National Natural Science Foundation of China under Grant U20A20222.

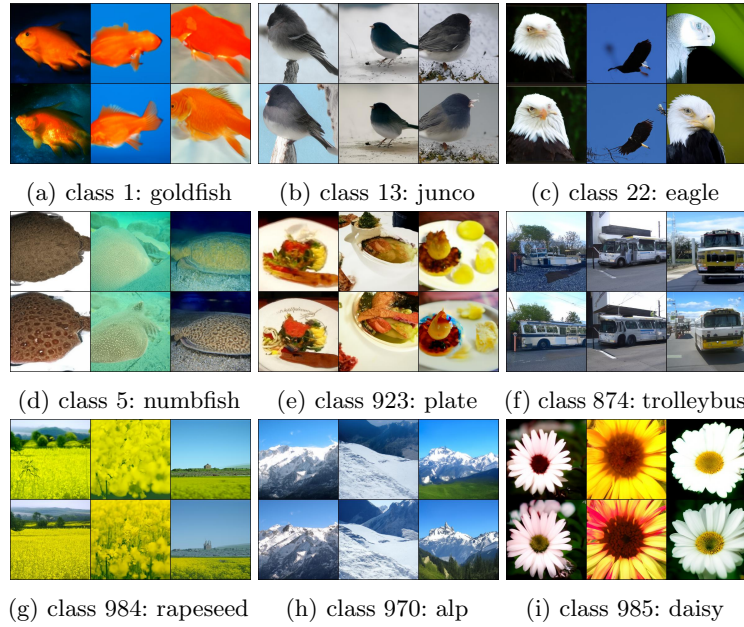


Fig. 5: Sample quality comparison between UADM-G [6] (the first line of each box) and our UADM-G+ECT+EDS (the second line of each box) on ImageNet1000 256×256 , DDIM 25 steps.

References

1. Barratt, S., Sharma, R.: A note on the inception score. arXiv:1801.01973 (2018)
2. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. arXiv:1809.11096 (2018)
3. Bulat, A., Yang, J., Tzimiropoulos, G.: To learn image super-resolution, use a gan to learn how to do image degradation first. In: Proceedings of the European conference on computer vision (ECCV). pp. 185–200 (2018)
4. Choi, J., Kim, S., Jeong, Y., Gwon, Y., Yoon, S.: Ilvr: Conditioning method for denoising diffusion probabilistic models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 14367–14376 (October 2021)
5. De Bortoli, V., Thornton, J., Heng, J., Doucet, A.: Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems* **34** (2021)
6. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems* **34** (2021)
7. Donahue, J., Simonyan, K.: Large scale adversarial representation learning. arXiv:1907.02544 (2019)
8. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. arXiv:1406.2661 (2014)
9. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in Neural Information Processing Systems* 30 (NIPS 2017) (2017)
10. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: *NeurIPS* (2020)
11. Huang, C.W., Lim, J.H., Courville, A.C.: A variational perspective on diffusion-based generative models and score matching. *Advances in Neural Information Processing Systems* **34** (2021)
12. Kingma, D.P., Salimans, T., Poole, B., Ho, J.: Variational diffusion models. arXiv preprint arXiv:2107.00630 (2021)
13. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *stat* **1050**, 1 (2014)
14. Kong, Z., Ping, W., Huang, J., Zhao, K., Catanzaro, B.: Diffwave: A versatile diffusion model for audio synthesis. arXiv:2009.09761 (2020)
15. Kynkäänniemi, T., Karras, T., Laine, S., Lehtinen, J., Aila, T.: Improved precision and recall metric for assessing generative models. *Advances in Neural Information Processing Systems* **32** (2019)
16. Liu, J., Li, C., Ren, Y., Chen, F., Liu, P., Zhao, Z.: Diffsinger: Singing voice synthesis via shallow diffusion mechanism. arXiv preprint arXiv:2105.02446 (2021)
17. Luhman, E., Luhman, T.: Knowledge distillation in iterative generative models for improved sampling speed. arXiv:2101.02388 (2021)
18. Lyu, Z., Kong, Z., Xu, X., Pan, L., Lin, D.: A conditional point diffusion-refinement paradigm for 3d point cloud completion. arXiv preprint arXiv:2112.03530 (2021)
19. Meng, C., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: Sedit: Image synthesis and editing with stochastic differential equations. arXiv preprint arXiv:2108.01073 (2021)
20. Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv:1411.1784 (2014)

21. Nachmani, E., Roman, R.S., Wolf, L.: Non gaussian denoising diffusion models. arXiv preprint arXiv:2106.07582 (2021)
22. Nash, C., Menick, J., Dieleman, S., Battaglia, P.W.: Generating images with sparse representations. arXiv:2103.03841 (2021)
23. Nichol, A., Dhariwal, P.: Improved denoising diffusion probabilistic models. arXiv preprint arXiv:2102.09672 (2021)
24. Nie, W., Vahdat, A., Anandkumar, A.: Controllable and compositional generation with latent-space energy-based models. *Advances in Neural Information Processing Systems* **34** (2021)
25. van den Oord, A., Vinyals, O., Kavukcuoglu, K.: Neural discrete representation learning. arXiv:1711.00937 (2017)
26. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. arXiv preprint arXiv:2102.12092 (2021)
27. Razavi, A., van den Oord, A., Vinyals, O.: Generating diverse high-fidelity images with VQ-VAE-2. arXiv:1906.00446 (2019)
28. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: Imagenet large scale visual recognition challenge. arXiv:1409.0575 (2014)
29. Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image super-resolution via iterative refinement. arXiv:arXiv:2104.07636 (2021)
30. Sasaki, H., Willcocks, C.G., Breckon, T.P.: Unit-ddpm: Unpaired image translation with denoising diffusion probabilistic models. arXiv preprint arXiv:2104.05358 (2021)
31. Sinha, A., Song, J., Meng, C., Ermon, S.: D2c: Diffusion-decoding models for few-shot conditional generation. *Advances in Neural Information Processing Systems* **34** (2021)
32. Sohl-Dickstein, J., Weiss, E.A., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. arXiv:1503.03585 (2015)
33. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: *International Conference on Learning Representations* (2020)
34. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv:2010.02502 (2020)
35. Song, Y., Durkan, C., Murray, I., Ermon, S.: Maximum likelihood training of score-based diffusion models. *Advances in Neural Information Processing Systems* **34** (2021)
36. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. arXiv:arXiv:1907.05600 (2020)
37. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. arXiv:2011.13456 (2020)
38. Vahdat, A., Kautz, J.: Nvae: A deep hierarchical variational autoencoder. arXiv:2007.03898 (2020)
39. Vincent, P.: A connection between score matching and denoising autoencoders. *Neural computation* **23**(7), 1661–1674 (2011)
40. Wang, H., Lin, G., Hoi, S.C.H., Miao, C.: Cycle-consistent inverse gan for text-to-image synthesis (2021)
41. Xia, W., Yang, Y., Xue, J.H., Wu, B.: Tedigan: Text-guided diverse face image generation and manipulation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2256–2265 (2021)