

# Self-Corrected Flow Distillation for Consistent One-Step and Few-Step Text-to-Image Generation

Quan Dao<sup>\*1,2†</sup>, Hao Phung<sup>\*1,3†</sup>, Trung Tuan Dao<sup>1</sup>, Dimitris N. Metaxas<sup>2</sup>, Anh Tran<sup>1</sup>

<sup>1</sup>VinAI Research

<sup>2</sup>Rutgers University

<sup>3</sup>Cornell University

quan.dao@rutgers.edu, htp26@cornell.edu, v.trungdt21@vinai.io, dnm@cs.rutgers.edu, v.anhtt152@vinai.io

## Abstract

Flow matching has emerged as a promising framework for training generative models, demonstrating impressive empirical performance while offering relative ease of training compared to diffusion-based models. However, this method still requires numerous function evaluations in the sampling process. To address these limitations, we introduce a self-corrected flow distillation method that effectively integrates consistency models and adversarial training within the flow-matching framework. This work is a pioneer in achieving consistent generation quality in both few-step and one-step sampling. Our extensive experiments validate the effectiveness of our method, yielding superior results both quantitatively and qualitatively on CelebA-HQ and zero-shot benchmarks on the COCO dataset.

**Code** — <https://github.com/VinAIRResearch/SCFlow>

## Introduction

The field of generative modeling has witnessed remarkable progress over the past decade. The modern generative models could create diverse and realistic content across various modalities. Previously, Generative Adversarial Networks (GANs) (Goodfellow et al. 2014; Karras, Laine, and Aila 2019) was dominant in this field by their ability to create realistic images. However, training GAN models is costly in both time and resources due to training instability and mode collapse. The emergence of diffusion models (Ho, Jain, and Abbeel 2020; Song and Ermon 2019; Song et al. 2020) marked a significant focus shift in generative AI. These models, exemplified by groundbreaking works such as DALL-E (Ramesh et al. 2021) and Stable Diffusion (Rombach et al. 2022) have surpassed GANs to become the current state-of-the-art in image synthesis. Diffusion models define a fix forward process which gradually perturbs image to noise and learn a model to perform the reverse process from noise to image. Their success lies in the ability to capture complex distribution of data and produce fidelity and diverse images. This approach has effectively addressed many of the limitations faced by GANs, offering improved

<sup>\*</sup>These authors contributed equally.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>†</sup>Work done while at VinAI.

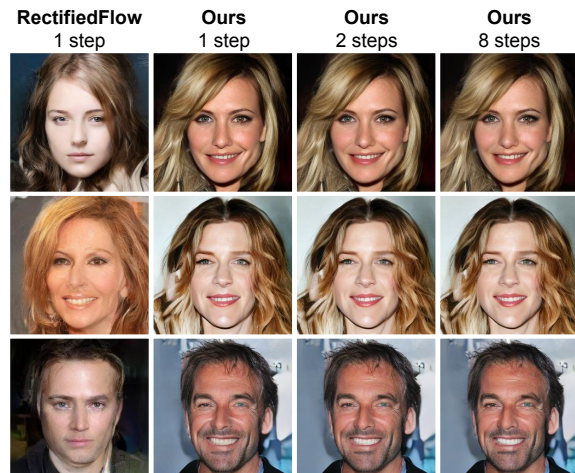


Figure 1: Illustration of consistent one-step and few-step image generation. Our method consistently delivers superior visual quality across different sampling steps, significantly surpassing the performance of the RectifiedFlow counterpart.

stability, diversity, and scalability. However, diffusion training takes long time to converge and requires many NFEs to produce high-quality samples. Recent works (Lipman et al. 2023; Liu 2022; Albergo and Vanden-Eijnden 2022) have introduced a flow matching framework, which is motivated by the continuous normalizing flow. By learning probability flow between noise and data distributions, flow matching models provide a novel perspective on generative modeling. Recent advancements have demonstrated that flow matching can achieve competitive results with diffusion models (Ma et al. 2024) while potentially offering faster sampling.

While flow matching can generate high-quality images with fewer NFEs compared to diffusion models, it still shares the challenge of prolonged sampling times due to its inherently iterative denoising process. This limitation poses a significant barrier to the practical application of both flow matching and diffusion models in real-world scenarios. To address this challenge in diffusion models, recent works (Meng et al. 2023; Luo et al. 2023; Gu et al. 2023; Nguyen and Tran 2024; Dao et al. 2024; Sauer et al. 2023b; Xu et al.

2024) focus on developing timestep distillation technique and show remarkable results. For example, LCM (Luo et al. 2023) utilizes consistency distillation (Song et al. 2023) and yields good results but generates blurry images at one-step sampling. SwiftBrush (Nguyen and Tran 2024) adopts SDS loss for distillation into a one-step generator but sacrifices the ability to perform multi-step sampling. Both UFOGEN (Xu et al. 2024) and SD Turbo (Sauer et al. 2023b) are able to generate high-quality images with one and few-step sampling. However, these methods struggle to maintain consistent results across different sampling schemes. In the context of flow matching, InstaFlow (Liu et al. 2024; Liu 2022) addresses this issue by utilizing rectified flow to produce direct transitions from source to target data. Instaflow has three training stages: collecting data, rectified flow and distillation. The instaflow could produce high-quality one-step generation but fail to perform few-step sampling due to their simple regression distillation at the third stage.

In this paper, we investigate how to distill a latent flow matching teacher into a consistent one and few-step generator. Motivated from consistency model (Luo and Hu 2021; Song et al. 2023), we apply consistency framework into latent flow model. However, we found that naively applying consistency distillation faces two challenges which are blurry one-step and oversaturated few-step generated images<sup>1</sup>. The blurry one-step is also observed in LCM (Luo and Hu 2021). These limitations could be due to discrepancy in statistic of latent compared to pixel space. To deal with blurry one-step generation, we propose to use GAN model for enhancing quality of one-step images. For oversaturated few-step problem, we introduce truncated consistency and reflow loss. These losses effectively mitigate the oversaturation problem, ensuring improved performance in few-step sampling. Besides, we also propose bidirection loss to improve the consistency across different sampling schemes. Our proposed framework is called self-corrected flow distillation. By thorough experiments, we validate the effectiveness of our framework to produce high quality and consistent images in both one and few-step sampling.

Our key contributions are threefold:

- We propose a training framework to effectively address the unique challenges of latent consistency distillation and offers optimal combinations for improved performance, including a truncated consistency loss to mitigate oversaturation, GAN to overcome blurry one-step generation. Additionally, the reflow and bidirection losses are introduced to enhance the consistency of generator across different sampling steps.
- Through extensive experiments on multiple datasets, we demonstrate that our approach significantly outperforms existing methods in both one-step and few-step generation, achieving competitive FID scores while maintaining generation speed. We provide detailed ablation studies to analyze the impact of each component of our method.

<sup>1</sup>“oversaturated” refers to the phenomenon where images generated by the model exhibit excessively vibrant colors and overly high contrast, resulting in a loss of natural color balance and detail.




Figure 2: Qualitative results of our Distilled Text-to-Image diffusion model.

- For the first time, we have achieved consistent, high-quality image generation in both few-step and one-step sampling using flow matching. The model will be publicly released to support further research.

## Related Work

### Flow Matching

Flow matching is emerging as the competitive alternative to diffusion models, as it deterministically finds the mapping between noise and data distribution. The deterministic property is favored in many generative applications, such as image inversion (Pokle et al. 2023) and editing (Hu et al. 2024), as well as in video and beyond (Davtyan, Sameni, and Favaro 2023; Song et al. 2024; Gao et al. 2024), due to its fast generation capability and reduced need for large NFE (Liu, Gong, and Liu 2023; Lipman et al. 2023; Liu 2022; Dao et al. 2023). Recently, some works has linked the connection between diffusion models (known as score-based models) and flow matching (Kingma and Gao 2024; Ma et al. 2024). Given these advantages, SDv3 (Esser et al. 2024) has adopted flow matching as their core framework, combined with a powerful transformer-based architecture (Peebles and Xie 2022), resulting in groundbreaking image generation capabilities. However, the computational complexity of iterative sampling still hinders these models from achieving real-time performance and lags behind GAN counterparts (Kang et al. 2023; Sauer et al. 2023a). Therefore, developing one-step and few-step sampling techniques is crucial to strike a balance between generation quality and sampling speed.

### Distillation Technique

Knowledge distillation (Hinton, Vinyals, and Dean 2015) has gained remarkable success in enhancing the performance of lightweight models under the guidance of a complex teacher model, allowing the student model to match or even surpass the teacher one. In context of diffusion models, instead of reducing model size, there is a line of methods (Luhman and Luhman 2021; Salimans and Ho 2022) that aims to distill a pre-trained diffusion-model teacher for reducing the number of sampling steps. Recently, Consistency model (Song et al. 2023) has achieved promising results in enhancing sampling efficiency of diffusion models. LCM (Luo et al. 2023) is the closest to our work where they directly adopted the consistency distillation objective,

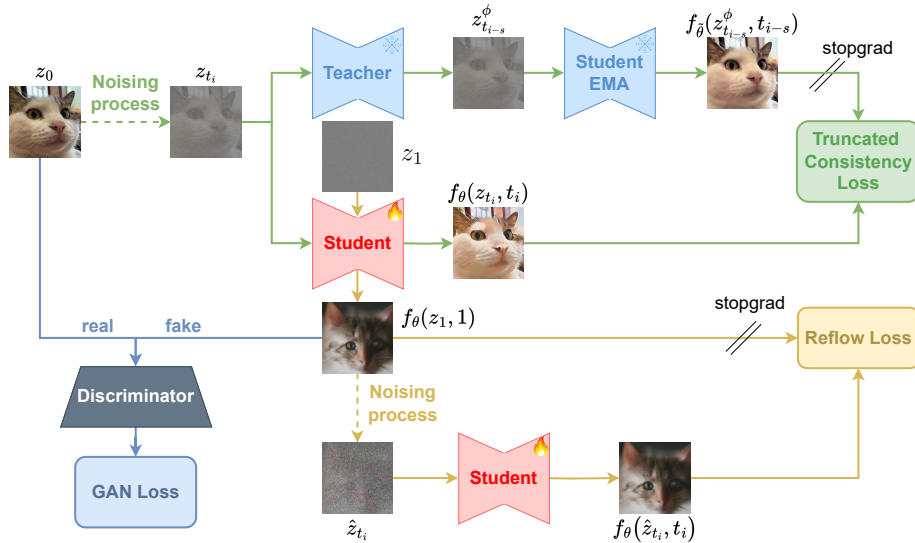


Figure 3: The overview of our Self-Corrected Flow Distillation method. All the latents are inputted as image for easier follow.

allowing Stable Diffusion model<sup>2</sup> to generate an image with just a few steps. Similar to our work, their method is also exploited in latent space of a pre-trained encoder. In contrast, our method proposes a consistency-based distillation method that is well adapted to the flow-matching framework. By combining with adversarial training and reflow objectives, our method can significantly increase the performance of few-step generation for both unconditional and conditional tasks. Unlike RectifiedFlow (Liu 2022), which iteratively fine-tunes the model on generated noise-image pairs using a pretrained flow model, our method eliminates the need for this costly, separate flow-straightening stage for distillation. Instead, it performs flow rectification and distillation concurrently during only one-stage training. While our method integrates adversarial objectives akin to AdversarialDSM (Jolicœur-Martineau et al. 2020), it distinguishes itself by optimizing GAN on latent encoded features instead of coarse pixels. This enhances the generation quality of one-step generation, diverging from a sole focus on improving the fidelity of score-based networks like the former.

## Method

Flow matching exhibits faster training convergence (Dao et al. 2023) and better image generation (Ma et al. 2024) compared to diffusion model. Thanks to these advantage, research community start shifting attention to this framework. Recent work (Esser et al. 2024) has scaled up the flow matching to text-to-image generation with high quality result. In contrast, flow matching still take long time for sampling compared to GAN model. This motivates us take deeper investigation in distillation method for this framework. In this section, we start by revisiting latent flow matching framework (Ma et al. 2024; Dao et al. 2023).

<sup>2</sup><https://github.com/Stability-AI/stablediffusion.git>

Next, we detail technical aspects of our proposed distillation framework named Self-Corrected Flow Distillation. Our distillation method is motivated from (Luo et al. 2023; Song et al. 2023). We show that straightly applying consistency distillation on latent flow matching framework yield low quality generation with both one-step and few-step sampling scheme. This behaviour also appeared in LCM (Luo et al. 2023), which remains unsolved until now. By utilizing GAN and Rectified technique (Liu 2022), we could mitigate the drawbacks of latent consistency distillation.

## Preliminary

Given the training dataset  $\mathbf{D}$ , we draw a sample  $\mathbf{x}_0 \in \mathbf{R}^d$  from the dataset. Denote that  $\mathcal{E}$  and  $\mathcal{D}$  are encoder and decoder of a pretrained VAE, we obtain the latent  $\mathbf{z}_0 = \mathcal{E}(\mathbf{x}_0) \in \mathbf{R}^{d/h}$ , where  $h$  represents the compressed rate of VAE model. The training objective of latent flow matching is to approximate a probabilistic path from a random noise  $\mathbf{z}_1 \sim \mathcal{N}(0, \mathbf{I}^{d/h})$  to the training dataset distribution  $\mathbf{z}_0$ . Previous works (Ma et al. 2024; Dao et al. 2023; Liu 2022; Lipman et al. 2023) use the following velocity loss to train flow matching framework:

$$\hat{\theta} = \operatorname{argmin}_{\theta} \mathbf{E}_{t, \mathbf{z}_t} [\|\mathbf{z}_1 - \mathbf{z}_0 - v_{\theta}(\mathbf{z}_t, t)\|_2^2]. \quad (1)$$

To enable the conditional generation, the condition information  $c$  is injected into the flow matching framework as below (Liu et al. 2024; Ma et al. 2024; Dao et al. 2023):

$$\hat{\theta} = \operatorname{argmin}_{\theta} \mathbf{E}_{t, \mathbf{z}_t} [\|\mathbf{z}_1 - \mathbf{z}_0 - v_{\theta}(\mathbf{z}_t, c, t)\|_2^2]. \quad (2)$$

Conditioning information  $c$  can be images, text, or class labels, with different conditional mechanisms like AdaIN (Peebles and Xie 2022) or cross-attention (Wang et al. 2018).

---

**Algorithm 1: Self-Corrected Flow Distillation**


---

**Data:** data  $p_0$ , Encoder  $\mathcal{E}$ , distilled model  $v_\theta$ , pretrained model  $v_\phi$ , lr  $\eta$ , ema decay  $\mu$ , and  $\lambda_{GAN}, \lambda_{RF}, \lambda_{BI}$  are weight terms

```

 $\theta \leftarrow \phi$ ;
for  $iter \in \{1, \dots, N\}$  do
   $\mathbf{x}_0 \sim p_0, \mathbf{z}_0 \leftarrow \mathcal{E}(\mathbf{x}_0)$ ;
   $\mathbf{z}_1 \leftarrow \mathcal{N}(0, \mathbf{I}), i \sim \mathcal{U}[1, N]$ ;
   $\mathbf{z}_{t_i} \leftarrow (1 - t_i)\mathbf{z}_0 + t_i\mathbf{z}_1$ ;
   $\mathcal{L}_{distill} = \mathcal{L}_{CD}$  (using eq. (8))
  if  $iter \geq N_{GAN}$  then
     $\mathcal{L}_{distill} \leftarrow \mathcal{L}_{distill} + \lambda_{GAN} * \mathcal{L}_{GAN}$  (using eq. (7))
  end
  if  $iter \geq N_{RF}$  then
     $\mathcal{L}_{distill} \leftarrow \mathcal{L}_{distill} + \lambda_{RF} * \mathcal{L}_{RF}$  (using eq. (9))
  end
  if  $iter \geq N_{BI}$  then
     $\mathcal{L}_{distill} \leftarrow \mathcal{L}_{distill} + \lambda_{BI} * \mathcal{L}_{BI}$  (using eq. (10))
  end
   $\theta \leftarrow \theta - \eta \nabla_{\theta} \mathcal{L}_{distill}$ ;
   $\tilde{\theta} = \text{sg}(\mu \tilde{\theta} + (1 - \mu)\theta)$ 
end

```

---

To better control the diversity and quality of generation, previous works (Liu et al. 2024; Dao et al. 2023) adopt classifier-free guidance sampling algorithm similar to (Ho and Salimans 2022):

$$\tilde{v}_\theta(\mathbf{x}_t, \mathbf{c}, t) \approx \gamma v_\theta(\mathbf{x}_t, \mathbf{c}, t) + (1 - \gamma)v_\theta(\mathbf{x}_t, \mathbf{c} = \emptyset, t), \quad (3)$$

where  $v_\theta(\mathbf{x}_t, \mathbf{c} = \emptyset, t)$  represents the unconditional velocity trained with null token  $\mathbf{c}$ . Hyperparameter  $\gamma$  controls the generation of flow matching framework. While smaller values of  $\gamma$  promote diverse outputs, larger  $\gamma$  values tend to yield higher fidelity images at the cost of reduced diversity.

### Self-Corrected Flow Distillation

Given pretrained latent flow matching model  $v_\phi$ , we would like to distill from that teacher model to a student  $v_\theta$  that is capable of both one or many step sampling. Therefore, we firstly apply consistency distillation (Luo et al. 2023; Song et al. 2023) for pretrained teacher over  $N$  discrete times  $0 = t_1 < t_2 < \dots < t_N = 1$  as follows:

$$\mathcal{L}_{CD} = \mathbf{E}_{t_i, \mathbf{z}_{t_i}} \left[ \|f_{\tilde{\theta}}(\mathbf{z}_{t_{i-s}}^\phi, t_{i-s}) - f_\theta(\mathbf{z}_{t_i}, t_i)\|_2^2 \right], \quad (4)$$

where  $s$  is skipping timesteps and  $\tilde{\theta} = \text{sg}(\mu \tilde{\theta} + (1 - \mu)\theta)$  is the exponential moving average (EMA) of  $\theta$  model with a decay rate  $\mu \in [0, 1]$  with stop-grad operator  $\text{sg}$ . The terms  $\mathbf{z}_{t_{i-s}}^\phi$  and  $f_\theta(\mathbf{z}_{t_i}, t_i)$  are defined as follow:

$$f_\theta(\mathbf{z}_{t_i}, t_i) = \mathbf{z}_{t_i} - t_i * v_\theta(\mathbf{z}_{t_i}, t_i), \quad (5)$$

$$\mathbf{z}_{t_{i-s}}^\phi = \mathbf{z}_{t_i} - (t_i - t_{i-s}) * v_\phi(\mathbf{z}_{t_i}, t_i). \quad (6)$$

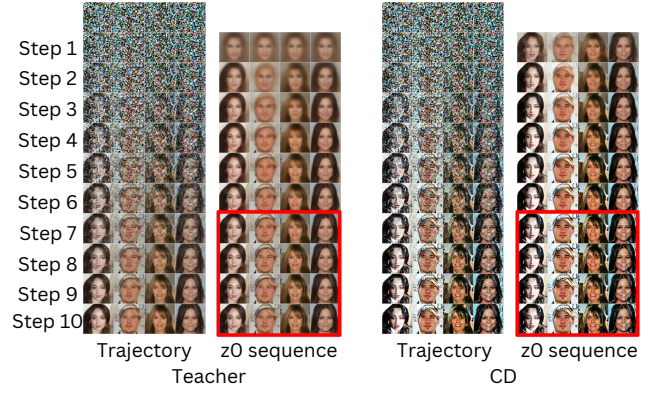


Figure 4: Trajectory of 10 NFEs Euler sampling of vanilla flow matching (teacher model) and CD model.

Solely applying consistency distillation (Song et al. 2023) on latent space presents two challenges: (1) one-step synthesis produces blurry images, which would significantly degrade the FID metric - this observation aligns with findings in (Luo et al. 2023); (2) when sampling with few-step, the student model generates oversaturated images, as illustrated in fig. 6. These limitations could due to the statistical difference between latent and pixel space. To address these limitations, we propose to use GAN and Reflow techniques.

**Blurry outputs of one-step generation.** We realize that one-step images produced by student model is blurry as seen in first row of fig. 6. Since single-step image still contain coarse structure information, we propose to apply GAN to further boost the sharpness of one-step images. To ensure that the one-step images already contain coarse structure, we start applying GAN loss after several iterations of consistency distillation. This is similar to the warm-up technique of VQGAN (Esser, Rombach, and Ommer 2021) to reduce the training instability. The proposed GAN is as follow:

$$\mathcal{L}_{GAN} = \mathcal{D}_{adv}(f_\theta(\mathbf{z}_1, 1), z_0), \quad (7)$$

where  $f_\theta(\mathbf{z}_1, 1)$  is the student’s one-step generated image. **Oversaturated outputs of few-step generation.** As shown in fig. 4, we realize that  $f_\theta(\mathbf{z}_{t_i}, t_i)$  become over-saturated as  $t_i \leq 0.4$ . Besides, we observe that when  $t_i$  is small,  $f_\theta(\mathbf{z}_{t_i}, t_i)$  can be well approximated by  $z_0$ , which is not hold for large  $t_i$ . Therefore, we can use diffusion loss  $\|f_\theta(x_{t_i}, t_i) - x_0\|_2^2$  instead of  $\sum_{t_i < 0.4} \|f_{\tilde{\theta}}(x_{t_{i-s}}, t_{i-s}) - f_\theta(x_{t_i}, t_i)\|_2^2$ . Furthermore, the reason for not using consistency loss for small  $t_i$  is that the value of  $f_{\tilde{\theta}}(x_{t_{i-s}}, t_{i-s}) - f_\theta(x_{t_i}, t_i) \approx 0$ , therefore the update gradient for it is minimal leading to subliminal gradient update on small  $t_i$ . By using truncated consistency loss (eq. (8)), we observe less oversaturated synthesis for a few-step generation.

$$\mathcal{L}_{CD} = \begin{cases} \mathbf{E}_{t_i, \mathbf{z}_{t_i}} \left[ \|f_{\tilde{\theta}}(\mathbf{z}_{t_{i-s}}^\phi, t_{i-s}) - f_\theta(\mathbf{z}_{t_i}, t_i)\|_2^2 \right] & \text{if } t_i > 0.4 \\ \mathbf{E}_{t_i, \mathbf{z}_{t_i}} \left[ \|z_0 - f_\theta(\mathbf{z}_{t_i}, t_i)\|_2^2 \right] & \text{if } t_i \leq 0.4 \end{cases} \quad (8)$$

However, truncated  $\mathcal{L}_{CD}$  cannot fully eliminate oversaturated limitation as shown in fig. 6. In addition, since the proposed GAN loss only enhance the quality of one-step image,

| Method         | NFE↓ | FID30K↓      | CLIP↑        | P↑          | R↓          |
|----------------|------|--------------|--------------|-------------|-------------|
| 2-RF (teacher) | 25   | 11.08        | -            | -           | -           |
| Guided Distill | 1    | 37.30        | 0.270        | -           | -           |
| UFOGen         | 1    | 12.78        | -            | -           | -           |
| SD Turbo       | 1    | 16.10        | <b>0.330</b> | <b>0.65</b> | 0.35        |
| LCM            | 1    | 35.56        | 0.24         | -           | -           |
| InstaFlow-0.9B | 1    | 13.10        | 0.280        | 0.53        | 0.45        |
| InstaFlow-1.7B | 1    | 11.83        | -            | -           | -           |
| 2-RF           | 1    | 36.83        | 0.282        | 0.30        | 0.33        |
| Ours-0.9B      | 1    | <b>11.91</b> | <u>0.312</u> | <u>0.54</u> | <b>0.47</b> |
| Ours-0.9B      | 2    | <b>11.46</b> | <u>0.315</u> | <u>0.57</u> | <b>0.46</b> |

Table 1: Text-to-image results on zero-shot COCO2014.

the inconsistency still exists between one-step and few-step images as shown in second row of fig. 6.

To better reduce saturated effect and improve the consistency between one-step and few-step synthesized images, we propose a reflow loss motivated from RectifiedFlow (Liu 2022; Liu et al. 2024). This loss directly self-corrects the flow estimation, given a reliable estimate of one-step source  $f_\theta(z_1, 1)$  as below:

$$\mathcal{L}_{RF} = \mathbf{E}_{t_i, \hat{z}_{t_i}} [\| \text{sg}(f_\theta(\mathbf{z}_1, 1)) - f_\theta(\hat{z}_{t_i}, t_i) \|_2^2], \quad (9)$$

with  $\hat{z}_{t_i} = (1 - t_i) * \text{sg}(f_\theta(\mathbf{z}_1, 1)) + t_i * \mathbf{z}_1$ .

Rectified Flow (Liu 2022) requires to sample  $\mathbf{z}_0$  by using multi-step generation before applying rectified flow technique. This process costs both time and memory for generating high-quality images from teacher model. Unlike RectifiedFlow, we directly use one-step image  $f_\theta(\mathbf{z}_1, 1)$  for rectified flow technique instead of multi-step image. Interestingly, one-step images are not oversaturated and are high quality due to GAN loss as seen in second row of fig. 6. Consequently, our proposed reflow loss effectively addresses the oversaturation issue in few-step sampling while enhancing consistency between few-step and one-step generation through the straightness penalty of the rectified flow loss, as illustrated in the third row of fig. 6.

**Bi-directional Consistency Distillation.** In consistency distillation, the  $L_{CD}$  objective forces the output  $f_\theta(z_{t_i}, t_i)$  to close to  $f_\theta(z_0, 0) \approx x_0$ , the high-quality source at the end of the denoising process (due to  $\|f_\theta(\mathbf{z}_{t_i}, t_i) - f_\theta(\mathbf{z}_0, 0)\|_2^2 \leq \sum_i \|f_\theta(\mathbf{z}_{t_i}, t_i) - f_\theta(\mathbf{z}_{t_{i-1}}, t_{i-1})\|_2^2 = L_{CD}$ ). However, thanks to the GAN objective (eq. (7)), we can generate high-quality one-step samples  $f_\theta(z_1, 1)$ , at the start of the denoising process. Therefore, incorporating the bidirectional loss ensures that  $f_\theta(z_{t_i}, t_i)$  receives quality signals from both endpoints of the denoising process, thus enhancing the consistency at both directions. The bi-directional objective is written below:

$$\mathcal{L}_{BI} = \mathbf{E}_{t_i, \mathbf{z}_{t_i}} [\|f_{\hat{\theta}}(\mathbf{z}_{t_{i+s}}^\phi, t_{i+s}) - f_\theta(\mathbf{z}_{t_i}, t_i)\|_2^2]. \quad (10)$$

Importantly, it is activated only when high-quality 1-NFE outputs is ensured by GAN loss, allowing beneficial signals to guide student training and avoiding poor-quality information in early training stages.

| Method          | NFE↓ | FID5K↓       | CLIP↑        | Time (s) |
|-----------------|------|--------------|--------------|----------|
| 2-RF (teacher)  | 25   | 21.50        | 0.315        | 0.88     |
| Guided Distill. | 1    | 37.20        | 0.275        | 0.09     |
| 2-RF            | 1    | 47.00        | 0.271        | 0.09     |
| UFOGen          | 1    | 22.50        | 0.311        | 0.09     |
| InstaFlow-0.9B  | 1    | 23.40        | 0.304        | 0.09     |
| InstaFlow-1.7B  | 1    | 22.40        | 0.309        | 0.12     |
| Ours-0.9B       | 1    | <b>22.09</b> | <b>0.313</b> | 0.09     |
| Guided Distill. | 2    | 26.00        | 0.297        | 0.13     |
| 2-RF            | 2    | 31.30        | 0.296        | 0.13     |
| Ours-0.9B       | 2    | <b>21.20</b> | <b>0.317</b> | 0.13     |

Table 2: Text-to-image results on zero-shot COCO2017.

| Model                    | NFE↓ | FID↓   |
|--------------------------|------|--------|
| One-Step                 |      |        |
| LFM                      | 1    | 200.13 |
| LFM+ Rectified           | 1    | 18.03  |
| LFM+ Rectified + Distill | 1    | 12.95  |
| LFM+ CD                  | 1    | 41.34  |
| Ours                     | 1    | 8.06   |
| Multi-Step               |      |        |
| LFM+ Rectified + Distill | 2    | 30.85  |
| LFM+ CD                  | 2    | 23.56  |
| Ours                     | 2    | 7.67   |

Table 3: Quantitative results on CelebA-HQ 256.

With the proposed loss terms, our distill student could be able to generate high quality images in both one and few step setting, refer to fig. 5. Our overall distillation framework is briefly described by the algorithm 1.

## Experiment

### Self-Corrected Flow Distillation For Unconditional Generation

**Training details.** Our experiments are conducted on CelebA-HQ 256 for pretrained latent flow matching model from LFM (Dao et al. 2023). We modify and use the discriminator architecture from (Phung, Dao, and Tran 2023). Our distillation procedure uses 200 training epochs with learning rate  $1e-5$  for both discriminator and student. The default ema rate  $\mu$  is 0.9,  $t_{trunc}$  is 0.4 and  $t_{skip}$  is 0.1. The loss weight ( $\lambda_{GAN, RF, BI}$ ) and warm-up iteration ( $N_{GAN, RF, BI}$ ) are set to (0.1, 0.1, 0.1) and (0, 1000, 1000). For sampling process, we use for Euler solver by default.

**Experimental results.** We compare our method with 2 baselines Rectified Flow (Liu 2022; Liu et al. 2024) and Consistency Distillation (Song et al. 2023; Luo et al. 2023). The reason for choosing these baseline is that the technique allows both one and few-step sampling scheme. For rectified flow, we follow rectified distillation framework (Liu 2022) which comprises of three stage: data generation, rectified flow and distillation. We firstly create a set of 50,000 pairs



Figure 5: Varying NFEs on CelebA-HQ. Increasing NFEs accentuates details and sharpness in generated faces without oversaturation issues.

$(z_1, z_0)$  using 500 steps Euler solver, where  $z_1, z_0$  are random noise and synthesized image correspondingly. We then train rectified flow for 50 epochs on the synthesized set. Finally, we perform distillation stage by directly mapping from  $z_1$  to  $z_0$  in 10 epochs. For Consistency Distillation, we follow (Song et al. 2023) implementation and distill model for 50 epochs on CelebA-HQ 256 real dataset.

The experiment result is reported in Table 3. For one-step generation, our approach achieves 8.06 FID which outperforms all the baselines. For two-step sampling, the same observation is also hold. Our method’s FID is 7.67 compared to Rectified Flow 30.85 and Consistency Distillation 23.56. Notably, 2-step FID of Rectified Distillation is higher than 1-step counterpart because the third stage mapping from  $z_1$  to  $z_0$  hurts the multistep sampling ability. Refer to fig. 1 for quality comparison between Rectified Distillation framework and our proposed framework. For the quality comparison with Consistency Distillation, please check the first and last row of fig. 6. These results underscore the efficacy of self-corrected flow distillation which not only produce high quality one-step generation but also maintaining high-quality with multiple-step sampling. Furthermore, fig. 5 and fig. 1 demonstrate the consistency generation of our framework with both few-step and one-step sampling produces same image given same noise input.

### Self-Corrected Flow Distillation For Text-to-Image Generation

**Evaluation metrics.** We evaluate our text-to-image model using a “zero-shot” framework, wherein the model is trained on one dataset and tested on another, ensuring a robust assessment of generalization capabilities. Our evaluation encompasses three critical dimensions: image quality, diversity, and textual fidelity. The primary metric for assessing image quality is the Fréchet Inception Distance (FID) (Heusel et al. 2017). In addition to FID, we employ precision and recall (Kynkäänniemi et al. 2019) as a complementary metric to assess image quality and diversity. For measuring the alignment between generated images and their

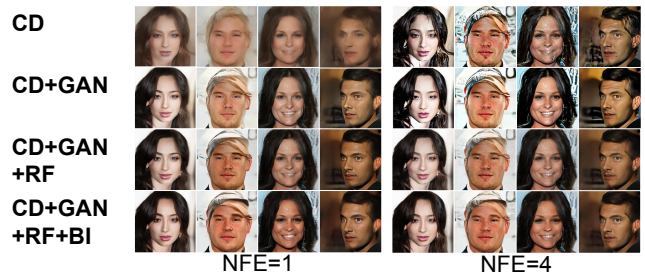


Figure 6: Qualitative of loss choice: 1 NFE vs 4 NFE

corresponding text prompts, we use the CLIP score (Radford et al. 2021). Following (Liu et al. 2024; Gu et al. 2023; Sauer et al. 2023b; Luo et al. 2023; Sauer et al. 2023a; Kang et al. 2023), we employ the MS COCO-2014 validation set and MSCOCO-2017 as our standard zero-shot text-to-image benchmarks. For MSCOCO-2014, we generate samples from the first 30,000 prompts, while for MSCOCO-2017, we use the first 5,000 prompts.

**Training details.** For our experiments, we employ a two-stage rectified flow (2-RF) model as the teacher, ensuring a fair comparison with InstaFlow. Our training process utilizes 2 million samples from the LAION dataset with an aesthetic score larger than 6.25. The architecture of our discriminator is based on a UNet-encoder design, augmented with an additional head, drawing inspiration from the UFOGen model. The model undergoes training for 18,000 iterations, with a consistent learning rate of  $1e-5$  applied to both the generator and discriminator components. To maintain consistency, all other hyperparameters are aligned with the configuration used in our CelebA-HQ experiments.

**Experimental results.** Table 1 presents our zero-shot text-to-image generation results on COCO2014, comparing our method with state-of-the-art approaches like 2-RF (Liu et al. 2024), Guided Distillation (Meng et al. 2023), UFOGen (Xu et al. 2024), SD Turbo (Sauer et al. 2023b), LCM (Luo et al. 2023), and InstaFlow (Liu et al. 2024).

Model Ours-0.9B achieves the best FID score of 11.91 with just one step, surpassing all other methods, including the larger InstaFlow-1.7B. This demonstrates our approach’s efficiency in generating high-quality images with minimal computation. For text-image alignment, our model’s CLIP score of 0.312 is second only to SD Turbo (0.330), indicating high relevance to input prompts. It also shows a good balance between precision (0.54) and recall (0.47) for one-step generation, suggesting diverse yet accurate image generation. Increasing to two steps further improves performance, with an FID of 11.46 and a CLIP score of 0.315. This showcases our approach’s scalability and ability to leverage additional computational steps for enhanced quality.

On the other hand, table 2 shows the results of our zero-shot evaluation on the COCO2017 dataset. Here, we observe similar trends to the COCO2014 results, with our model outperforming other methods in both one-step and two-step generation scenarios. For one-step generation, our model achieves the best FID score of 22.09 and the highest CLIP score of 0.313 among all compared methods. This perfor-

| Loss                 | NFE=1 | NFE=2 | NFE=4 | NFE=8 | NFE=16 |
|----------------------|-------|-------|-------|-------|--------|
| $\mathcal{L}_{CD}$   | 41.34 | 23.56 | 38.75 | 51.34 | 53.69  |
| $+\mathcal{L}_{GAN}$ | 11.55 | 19.82 | 38.29 | 49.45 | 54.56  |
| $+\mathcal{L}_{RF}$  | 10.62 | 11.07 | 11.8  | 12.68 | 12.66  |
| $+\mathcal{L}_{BI}$  | 8.06  | 7.67  | 7.68  | 7.63  | 7.24   |

Table 4: Ablation of our Self-Corrected Flow Distillation. FID is used for all experiments (Lower is better).

| $(\lambda_{GAN,RF,BI})$ | NFE=1 | NFE=2 | NFE=4 | NFE=8 | NFE=16 |
|-------------------------|-------|-------|-------|-------|--------|
| (1.0, 1.0, 1.0)         | 30.26 | 29.71 | 29.22 | 30.12 | 35.57  |
| (0.4, 0.4, 0.4)         | 12.22 | 12.07 | 11.66 | 11.06 | 9.95   |
| (0.2, 0.2, 0.2)         | 9.83  | 9.65  | 9.43  | 9.33  | 8.82   |
| (0.1, 0.1, 0.1)         | 8.06  | 7.67  | 7.68  | 7.63  | 7.24   |
| (1.0, 0.1, 0.1)         | 11.2  | 11.38 | 11.26 | 11.25 | 12.55  |
| (0.1, 1.0, 1.0)         | 35.22 | 31.68 | 30.8  | 31.44 | 33.12  |

Table 5: Ablation of weighting loss terms. By default, we use our best hyper-params like  $t_{trunc} = 0.4$ ,  $\mu = 0.9$ , and  $t_{skip} = 0.1$  in case one of those is not explicitly mentioned in the table.

mance is particularly impressive considering that our model matches or exceeds the quality of models with more parameters (e.g., InstaFlow-1.7B). When increasing to two steps, our model further improves its performance, achieving an FID of 21.20 and a CLIP score of 0.317. This not only outperforms other two-step methods but also surpasses the quality of models using many more steps, such as 2-RF with NFE=25. Importantly, our model maintains competitive inference times, with 0.09 seconds for one-step and 0.13 seconds for two-step generation, which is comparable to other efficient methods and significantly faster than multi-step approaches. For qualitative result, please refer to fig. 2.

### Ablation Studies For Self-Corrected Flow Distillation

We conduct extensive ablation studies on our distilled model, with results presented in table 6 and table 5. Specifically, Table 6 demonstrates that model performance is mostly influenced by three key parameters: the time-truncated threshold  $t_{trunc}$  in  $\mathcal{L}_{CD}$ , the EMA decay  $\mu$ , and the time-skip threshold  $t_{skip}$ . Our findings indicate that optimal results are achieved when  $t_{trunc}$  is within the range [0.2, 0.5) and  $t_{skip}$  is approximately 0.1. Notably, an EMA decay of  $\mu = 0.9$  consistently yields superior performance, particularly for many-step generation.

Table 5 explores the impact of weights on the GAN loss, reflow loss, and bidirectional consistency loss. Our results indicate that lower weights (around 0.1 to 0.2) for these components lead to optimal performance, highlighting the critical role of precise loss balancing in our framework.

Qualitative results across various NFEs are presented in fig. 5. Furthermore, fig. 6 illustrates the progressive improvements achieved by each component of our method:

- Consistency distillation alone (row 1) can lead to increased contrast and statistical shift at higher NFEs, ex-

|  | NFE=1 | NFE=2 | NFE=4 | NFE=8 | NFE=16 |
|--|-------|-------|-------|-------|--------|
| Time-truncated threshold in $\mathcal{L}_{CD}$ |       |       |       |       |        |
| $t_{trunc} = 0.$                               | 8.50  | 8.89  | 9.42  | 9.86  | 10.97  |
| $t_{trunc} = 0.1$                              | 8.44  | 8.80  | 9.36  | 9.87  | 11.03  |
| $t_{trunc} = 0.2$                              | 7.50  | 7.56  | 7.75  | 7.75  | 8.22   |
| $t_{trunc} = 0.4$                              | 8.06  | 7.67  | 7.68  | 7.63  | 7.24   |
| $t_{trunc} = 0.5$                              | 9.55  | 9.26  | 9.24  | 9.12  | 8.28   |
| EMA decay                                      |       |       |       |       |        |
| $\mu = 0.9$                                    | 8.06  | 7.67  | 7.68  | 7.63  | 7.24   |
| $\mu = 0.95$                                   | 7.79  | 7.90  | 8.07  | 8.00  | 7.93   |
| Time-skip threshold                            |       |       |       |       |        |
| $t_{skip} = 0.1$                               | 8.06  | 7.67  | 7.68  | 7.63  | 7.24   |
| $t_{skip} = 0.2$                               | 11.86 | 11.53 | 11.42 | 11.27 | 10.97  |

Table 6: Ablation of time-truncated threshold, EMA decay, and time-skip threshold. By default, we use our best hyper-params like  $t_{trunc} = 0.4$ ,  $\mu = 0.9$ , and  $t_{skip} = 0.1$  in case one of those is not explicitly mentioned in the table.

plaining the higher FID scores observed in table 4.

- The addition of GAN and TCD losses (row 2) address the blurriness in one-step generation but does not fully resolve oversaturation in multistep outputs.
- Reflow loss (row 3) enforces consistency between one-step and many-step generations, mitigating the oversaturation issue.
- The bidirectional term (row 4) further enhances the consistency generation across one and few-step sampling.

These observations underscore the statistical discrepancies between pixel and latent spaces, which manifest as blurriness in one-step generation and oversaturation in few-step generation. Our proposed method effectively mitigates these issues, as evidenced by fig. 5, where high-quality images are produced across various NFEs without oversaturation.

## Conclusion

This work presents Self-Corrected Flow Distillation ensuring consistent, high-quality generation in both one-step and few-step sampling. Our method successfully mitigates the limitation of latent consistency distillation, including blurry single-step and oversaturated multi-step samples. Our extensive experiments on CelebA-HQ and text-to-image generation tasks demonstrate substantial improvements over existing methods, achieving superior FID and visual quality for both one and few-steps sampling.

## Acknowledgements

Research partially funded by research grants to Prof. Dimitris Metaxas from NSF: 2310966, 2235405, 2212301, 2003874, 1951890, AFOSR 23RT0630, and NIH 2R01HL127661.

## References

- Albergo, M. S.; and Vanden-Eijnden, E. 2022. Building normalizing flows with stochastic interpolants. *ICLR 2023*, arXiv preprint arXiv:2209.15571.
- Dao, Q.; Phung, H.; Nguyen, B.; and Tran, A. 2023. Flow matching in latent space. arXiv preprint arXiv:2307.08698.
- Dao, T. T.; Nguyen, T. H.; Le, T.; Vu, D. H.; Nguyen, K.; Pham, C.; and Tran, A. T. 2024. SwiftBrush V2: Make Your One-Step Diffusion Model Better Than Its Teacher. In *ECCV (82)*.
- Davtyan, A.; Sameni, S.; and Favaro, P. 2023. Efficient video prediction via sparsely conditioned flow matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 23263–23274.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*.
- Esser, P.; Rombach, R.; and Ommer, B. 2021. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12873–12883.
- Gao, P.; Zhuo, L.; Liu, C.; ; Du, R.; Luo, X.; Qiu, L.; Zhang, Y.; et al. 2024. Lumina-T2X: Transforming Text into Any Modality, Resolution, and Duration via Flow-based Large Diffusion Transformers. arXiv preprint arXiv:2405.05945.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative Adversarial Nets. In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N.; and Weinberger, K. Q., eds., *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.
- Gu, J.; Zhai, S.; Zhang, Y.; Liu, L.; and Susskind, J. M. 2023. Boot: Data-free distillation of denoising diffusion models with bootstrapping. In *ICML 2023 Workshop on Structured Probabilistic Inference & Generative Modeling*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33: 6840–6851.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598.
- Hu, T.; Zhang, D. W.; Mettes, P.; Tang, M.; Zhao, D.; and Snoek, C. G. 2024. Latent Space Editing in Transformer-based Flow Matching. In *AAAI*.
- Jolicœur-Martineau, A.; Piché-Taillefer, R.; Combes, R. T. d.; and Mitliagkas, I. 2020. Adversarial score matching and improved sampling for image generation. arXiv preprint arXiv:2009.05475.
- Kang, M.; Zhu, J.-Y.; Zhang, R.; Park, J.; Shechtman, E.; Paris, S.; and Park, T. 2023. Scaling up GANs for Text-to-Image Synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4401–4410.
- Kingma, D.; and Gao, R. 2024. Understanding diffusion objectives as the elbo with simple data augmentation. *Advances in Neural Information Processing Systems*, 36.
- Kynkäänniemi, T.; Karras, T.; Laine, S.; Lehtinen, J.; and Aila, T. 2019. Improved precision and recall metric for assessing generative models. *Advances in Neural Information Processing Systems*, 32.
- Lipman, Y.; Chen, R. T. Q.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2023. Flow Matching for Generative Modeling. In *The Eleventh International Conference on Learning Representations*.
- Liu, Q. 2022. Rectified Flow: A Marginal Preserving Approach to Optimal Transport. ArXiv:2209.14577 [cs, stat].
- Liu, X.; Gong, C.; and Liu, Q. 2023. Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow.
- Liu, X.; Zhang, X.; Ma, J.; Peng, J.; and Liu, Q. 2024. InstafLOW: One step is enough for high-quality diffusion-based text-to-image generation. In *International Conference on Learning Representations*.
- Luhman, E.; and Luhman, T. 2021. Knowledge distillation in iterative generative models for improved sampling speed. arXiv preprint arXiv:2101.02388.
- Luo, S.; and Hu, W. 2021. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2837–2845.
- Luo, S.; Tan, Y.; Huang, L.; Li, J.; and Zhao, H. 2023. Latent Consistency Models: Synthesizing High-Resolution Images with Few-Step Inference. arXiv:2310.04378.
- Ma, N.; Goldstein, M.; Albergo, M. S.; Boffi, N. M.; Vanden-Eijnden, E.; and Xie, S. 2024. Sit: Exploring flow and diffusion-based generative models with scalable interpolant transformers. arXiv preprint arXiv:2401.08740.
- Meng, C.; Rombach, R.; Gao, R.; Kingma, D.; Ermon, S.; Ho, J.; and Salimans, T. 2023. On distillation of guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14297–14306.
- Nguyen, T. H.; and Tran, A. 2024. Swiftbrush: One-step text-to-image diffusion model with variational score distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7807–7816.
- Peebles, W.; and Xie, S. 2022. Scalable Diffusion Models with Transformers. arXiv preprint arXiv:2212.09748.

Phung, H.; Dao, Q.; and Tran, A. 2023. Wavelet diffusion models are fast and scalable image generators. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10199–10208.

Pokle, A.; Muckley, M. J.; Chen, R. T.; and Karrer, B. 2023. Training-free linear image inversion via flows. *arXiv preprint arXiv:2310.04432*.

Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.

Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-shot text-to-image generation. In *International conference on machine learning*, 8821–8831. Pmlr.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10684–10695.

Salimans, T.; and Ho, J. 2022. Progressive Distillation for Fast Sampling of Diffusion Models. In *International Conference on Learning Representations*.

Sauer, A.; Karras, T.; Laine, S.; Geiger, A.; and Aila, T. 2023a. Stylegan-t: Unlocking the power of gans for fast large-scale text-to-image synthesis. In *International conference on machine learning*, 30105–30118. PMLR.

Sauer, A.; Lorenz, D.; Blattmann, A.; and Rombach, R. 2023b. Adversarial diffusion distillation. *arXiv preprint arXiv:2311.17042*.

Song, Y.; Dhariwal, P.; Chen, M.; and Sutskever, I. 2023. Consistency Models. *arXiv:2303.01469*.

Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32.

Song, Y.; Gong, J.; Xu, M.; Cao, Z.; Lan, Y.; Ermon, S.; Zhou, H.; and Ma, W.-Y. 2024. Equivariant flow matching with hybrid probability transport for 3d molecule generation. *Advances in Neural Information Processing Systems*, 36.

Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.

Wang, T.-C.; Liu, M.-Y.; Zhu, J.-Y.; Tao, A.; Kautz, J.; and Catanzaro, B. 2018. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8798–8807.

Xu, Y.; Zhao, Y.; Xiao, Z.; and Hou, T. 2024. Ufogen: You forward once large scale text-to-image generation via diffusion gans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8196–8206.