# Schedule On the Fly: Diffusion Time Prediction for Faster and Better Image Generation

Zilyu Ye<sup>1,2,†</sup> Zhiyang Chen<sup>1,4,\*</sup>

<sup>\*</sup> Tiancheng Li<sup>1,2</sup> Zemin Huang<sup>1</sup> Weijian Luo<sup>1,3</sup> Guo-Jun Qi<sup>1,4,\*</sup> <sup>1</sup>MAPLE Lab, Westlake University

<sup>2</sup>South China University of Technology

<sup>3</sup>Peking University

<sup>4</sup>Institute of Advanced Technology, Westlake Institute for Advanced Study

zilyuye@foxmail.com, {chenzhiyang, litiancheng, huangzemin}@westlake.edu.cn

luoweijian@stu.pku.edu.cn, guojunq@gmail.com

#### Abstract

Diffusion and flow models have achieved remarkable successes in various applications such as text-to-image generation. However, these models typically rely on the same predetermined denoising schedules during inference for each prompt, which potentially limits the inference efficiency as well as the flexibility when handling different prompts. In this paper, we argue that the optimal noise schedule should adapt to each inference instance, and introduce the Time Prediction Diffusion Model (TPDM) to accomplish this. TPDM employs a plug-and-play Time Prediction Module (TPM) that predicts the next noise level based on current latent features at each denoising step. We train the TPM using reinforcement learning, aiming to maximize a reward that discounts the final image quality by the number of denoising steps. With such an adaptive scheduler, TPDM not only generates high-quality images that are aligned closely with human preferences but also adjusts the number of denoising steps and time on the fly, enhancing both performance and efficiency. We train TPDMs on multiple diffusion model benchmarks. With Stable Diffusion 3 Medium architecture, TPDM achieves an aesthetic score of 5.44 and a human preference score (HPS) of 29.59, while using around 50% fewer denoising steps to achieve better performance. We will release our best model alongside this paper.

## 1. Introduction

In recent years, deep generative models, including diffusion models [10, 44, 46] have achieved extraordinary per-

formance across a variety of tasks, including image synthesis [14, 15, 38, 39, 41], video generation [3, 11, 55], and others [18, 31, 35, 43]. As a multi-step denoising framework, diffusion models progressively refine random noise into coherent data through iterative sampling, which underlies their impressive capabilities in generating high-quality, diverse outputs.

However, inference with a diffusion model always involves manually selecting a noise scheduler, e.g. how the noise level changes step by step when denoising a clean image from Gaussian noise. This requires the user to delicately adjust parameters like the number of steps, resulting in a high threshold for customers to use. Numerous works attempt to find the optimal scheduler to liberate the users from these tedious tasks.

The leading flow-based models, Stable Diffusion 3 and Flux, provide recommended schedulers that adjust noise levels only based on the targets' resolution. Sabour [40] and Xia [52] explore ways to fine-tune the schedule for a model, improving either its efficiency or overall performance. In addition, some new one-step generators [13, 28–30] also achieve impressive performance. Despite the excellent performance they achieve, most of these works hold the assumption that there exists a universally applicable schedule that is optimal for all prompts and images, which is doubtful.

Take several images in Fig. 4 as an example. Images on the right are rich in content, requiring more denoising steps to capture finer details. In contrast, the images on the left are relatively simple and can be generated with fewer steps without compromising quality. Moreover, Karras [16] also proves that different noise schedules also greatly affect the generation quality. Thus, we may ask: *Is it possible to adaptively adjust both the number of denoising steps and the noise level at each step during inference, without any* 

The project was initiated and supported by MAPLE Lab at Westlake University.

<sup>†</sup> The work was done during an internship at MAPLE lab.

<sup>\*</sup> Corresponding author.



Figure 1. Samples generated by TPDM-FLUX.1-dev showcase stunning visual effects while adaptively adjusting inference steps based on the target output. The number in the lower right corner of each image indicates the inference steps used.

#### manual intervention from the user?

In this paper, we propose Time Prediction Diffusion Models(TPDMs) that can adaptively adjust both the number of steps and the denoising strength during inference. This is achieved by implementing a plug-and-play Time-Prediction Module (TPM) that can predict the next diffusion time conditioned on the latent features at the current step, so that the noise schedule can be adjusted on the fly.

TPM is trained with reinforcement learning. We regard the multi-step denoising process as an entire trajectory, and the quality of the generated image discounted by the number of steps as the reward. The image quality is measured by a reward model aligned with human preference [54].

TPM can be easily integrated into any diffusion models with marginal additional computation, and allow them to automatically adjust hyper-parameters like sample steps and the noise level at each step, reaching a balance between image quality and efficiency without human interventions. Moreover, during training, the model rolls out the diffusion process the same as in inference, directly optimizing the inference performance and reducing denoising steps.

We implement TPDM on several start-of-the-art models, including Stable Diffusion and Flux. With an adaptive noise schedule, our model can also generate 50% fewer steps on average with image quality on par or slightly better (0.322 CLIP-T, 5.445 Aesthetic Score, 22.33 Pick Score, 29.59 HPSv2.1) than Stable Diffusion 3. These results demonstrate that TPDM has the potential to either pursue high-quality image generation or improve model efficiency.

Our contributions are summarized below:

- We introduce the Time Prediction Diffusion Model (TPDM) that can adaptively adjust the noise schedule during inference, achieving a balance between image quality and model efficiency.
- To train TPDM, we maximize the generated image quality discounted by the number of steps with reinforcement

learning, directly optimizing final performance and efficiency.

• Our model demonstrates improved performance across multiple evaluation benchmarks, achieving better results with fewer inference steps.

## 2. Related Works

## 2.1. Diffusion Models

Diffusion Probability Models (DPMs) [10, 44] recover the original data from pure Gaussian noise by learning the data distribution across varying noise levels. With their strong adaptability to complex data distributions, diffusion models have achieved remarkable performance in diverse fields such as images [12, 17, 34, 35, 37], videos [3, 11, 55], and others [7, 50, 56, 57], significantly advancing the capabilities of Artificial Intelligence Generated Content.

#### 2.2. Noise Schedule

In order to generate an image, the model must determine the diffusion time for each step. This can be achieved using either discrete-time schedulers [10, 45] or continuous-time schedulers [22, 25] depending on the model. Typically, the diffusion time reflects the noise strength at each step, and most existing approaches rely on pre-determined schedules. Currently, the leading flow-based models, Stable Diffusion 3[8] and FLUX[19], provide recommended schedulers that adjust noise levels only based on the target resolution.

There exist some methods fine-tuning the scheduler to speed up sampling or improve image quality. Xia et al. [52] predict a new denoising time for each step to find a more accurate integration direction, Sabour et al. [40] uses stochastic calculus to find the optimal sampling plan for different schedulers and different models. Wang et al. [49] leverages reinforcement learning to automatically search for an optimal sampling scheduler. Some one-step generators[13, 29, 30] with diffusion distillation [26] also achieve impressive performance.

All of the above noise schedulers not only have high thresholds that require users to adjust many parameters, but also use the same denoising schedule for all prompts and images. On the contrary, TPDM can adaptively adjust the noise schedule during inference, and dynamically select the optimal sampling plan with decreased number of sampling steps for all prompts, achieving a balance between image quality and model efficiency. We will introduce the detailed definitions and practical algorithms of TPDM in Sec. 3.

### 2.3. Reinforcement Learning and Learning from Human Feedback

Reinforcement Learning from Human Feedback (RLHF) has recently gained significant attention in the field of large language models (LLMs) [1, 42, 47] and is gradually ex-

panding into other domains. Advances in diffusion models have increasingly incorporated reward models, trained on human preferences, to enhance alignment with human values [21, 51, 54]. Some recent works have also studied the RLHF for few-step generative models [27, 32]. However, most approaches only employ reward models as substitutes for diffusion loss or in conjunction with diffusion processes to align with human preferences, which significantly differs from the Proximal Policy Optimization (PPO) paradigm [42] commonly used in RLHF for LLMs. This divergence arises primarily because the multi-step denoising process in DPMs does not naturally yield likelihoods in the same manner as language models, posing challenges for the direct application of PPO. In this paper, however, we introduce an approach that, to the best of our knowledge, is the first to treat each denoising step as a distribution analogous to an action in reinforcement learning. This enables the use of well-established optimization techniques from LLMbased RLHF within the context of diffusion models.

## 3. The Proposed Approach

In this section, we first provide a brief review of the fundamental principles of diffusion models, followed by an introduction to the Time Prediction Module (TPM). Finally, we detail the training algorithm for the TPM.

#### 3.1. Preliminary

Diffusion Models learn to generate images with a reverse process that gradually removes noise from a sample. The leading paradigm for implementing this reverse process is flow matching [8, 9, 23]. Thus we introduce how flow-matching models work, and the detailed structure inside current start-of-the-art models here.

We consider a generative model that establishes a mapping between samples  $x_1$  drawn from a noise distribution  $p_1$  and samples  $x_0$  from a data distribution  $p_0$ . The Flow Matching objective aims to directly regress a vector field  $v_t$ , which generates a probability flow, enabling the transition from  $p_1$  to  $p_0$ .

$$\mathcal{L}_{FM}(\theta) = \mathbb{E}_{t,p_t(x)} \left\| v_{\theta}(x_t, t) - u(x_t, t) \right\|^2 \tag{1}$$

The flow-matching model with parameters  $\theta$  aims to predict the noise prediction function  $v_{\theta}(x_t, t)$ , which approximates the true velocity field  $u(x_t, t)$  that guides the diffusion process from the noise distribution to the clean distribution of generated samples. Thus, we can get the diffusion ODE:

$$\frac{dx_t}{dt} = v_\theta(x_t, t) \tag{2}$$

During inference, suppose we generate an image with N steps, each step has a time  $t_n$  corresponding to its noise



Figure 2. The Inference Process of TPDMs: The horizontal axis represents diffusion time, ranging from 1 to 0. The image starts from random noise  $x_{t_0}$  and is progressively denoised until a clean image  $x_{t_N}$ . Meanwhile, the reward is calculated for the final image and discounted by  $\gamma$  to influence previous steps.



Figure 3. The architecture of TPDM involves a frozen Diffusion Models, and a plug-and-play Time Prediction Module.

strength. Then, the n-th generation step can be formulated as:

$$x_{t_{n+1}} = x_{t_n} + (t_{n+1} - t_n) \cdot v_{\theta}(x_{t_n}, t_n)$$
(3)

Usually, in a typical flow matching algorithm, there is a predetermined schedule that arranges  $t_n$  every step.

Currently, many state-of-the-art diffusion models are built upon the DiT architecture [6, 8, 20], which only focus on conditional image generation and employ a modulation mechanism to condition the network on both the diffusion time during the diffusion denoising step and the text prompts.

With such a model, we can perform a single denoising step given the latent feature and current noise level.

#### 3.2. Time Prediction Diffusion Model (TPDM)

As aforementioned, we need a series of denoising steps to generate an image with the trained diffusion model. Usually, a fixed noise schedule is applied in this process for all prompts, assigning a pre-determined noise level for each step.

On the contrary, to enable the model to adjust the noise schedule on the fly, the TPDM predicts how noise strength t decays between adjacent steps. This ensures the noise strength is monotonously decreasing to reflect the denoising progress, avoiding backward progression. To be compatible with training in Sec. 3.3, TPM estimates the distribution over the denoising schedule, rather than predicting an exact value. Suppose we are performing the n-th denoising step. Besides the denoising outputs at the current step, TPDM predicts the distribution of the decay rate  $r_n$ as well. This distribution can be parameterized as a Beta Distribution over (0, 1), where the TPDM needs to estimate two parameters  $\alpha$  and  $\beta$ . We notice that when  $\alpha > 1$  and  $\beta > 1$  it can ensure a unimodal distribution. Therefore, we reparameterize that TPDM predicts two real numbers a and b, and determines the distributions with Eq. 4. Thus the decay rate  $r_n$  and the next noise level  $t_n$  can be sampled as in Eq. 5-Eq. 6.

$$\alpha = 1 + e^a, \qquad \beta = 1 + e^b, \tag{4}$$

$$r \sim \text{Beta}(\alpha, \beta)$$
 (5)

$$t_n = r_n \cdot t_{n-1} \tag{6}$$

TPDM requires only a minimal modification to the original diffusion model: adding a lightweight Time Prediction Module (TPM) as shown in Fig. 3. This module takes the latents before and after the transformer blocks as inputs, so that both the original noisy inputs and the predicted results at this step are taken into consideration. After several convolution layers, TPM pools the latents into a single feature vector and predicts a and b with two linear layers. We also use an adaptive normalization layer [36] in TPM so that the model is aware of the current time embedding.

During training, we freeze the original diffusion model and only update the newly introduced TPM. Thus, the model learns to predict the next diffusion time while preserving the original capacity for image generation.



#### A: A bright yellow sunflower in a field on a sunny day.

B: A steaming cup of coffee on a wooden table, surrounded by open books and a flickering candle.

C: A medieval knight in shining armor standing on a cliff, overlooking a vast battlefield at dawn, with a storm approaching in the distance. D: A majestic dragon soaring through a stormy sky, with lightning flashing around it, and a burning castle below, as mythical creatures run through the smoke and chaos.



A: A single red apple. B: A vintage pocket watch lying on an old leather-bound book, with soft sunlight casting gentle shadows. C: A futuristic cityscape at dusk, with sleek flying cars, towering skyscrapers, and neon signs glowing in the mist. D: An astronaut floating in space, with Earth visible in the background, surrounded by swirling galaxies and distant stars, while a come streaks across the sky.

Figure 4. From left to right, the images generated by TPDM-FLUX1.0-dev progress from simple to complex. Our Time Prediction Module adaptively adjusts the generation schedule to suit the complexity of each generation target.



**E.** A table topped with a cake covered in berries next to a plate of sandwiches.

F. a remote control holder attached to living room furniture that is full of controllers

Figure 5. Our TPDM-SD3-Medium, when compared to the SD3-Medium with the recommended and equivalent number of steps, demonstrates superior detail processing ability and generation accuracy.



Figure 6. The effect of  $\gamma$  on the average number of inference steps.

#### **3.3. Training Algorithms**

To train TPM, we need to roll out at least two denoising steps: predicting the next diffusion time with the first step, and denoising with this time at the next step. A naive method provides a noised image as input to the first step and trains the model with the reconstruction loss calculated at the second step. The gradients would backpropagate through the predicted  $t_n$  to update TPM in the first step.

However, we found that the trained model tends to complete the denoising process in very few steps during inference, leading to poor image quality. We hypothesize that by supervising the loss calculated after two steps, the model learns to generate a fully denoised image after two steps, and stop intermediately to minimize the loss function. However, what matters is the final image generated after the whole diffusion reverse process, which is ignored in this method.

Therefore, we optimize TPM to maximize the image quality generated after the whole denoising process to achieve precise time prediction. The quality is measured by an image reward model. Considering the computation graph of whole inference is too deep for gradients to backprop, we train the model with Proximal Policy Optimization (PPO) [42], whose loss function is formulated as

$$\mathcal{L}(\theta) = -\left[\frac{\pi_{\theta}(y|s)}{\pi_{\text{old}}(y|s)}\hat{A}(s,y) - \lambda \text{KL}[\pi_{\text{ref}}(\cdot|s), \pi_{\theta}(\cdot|s)]\right]$$
(7)

Under our problem definition,  $s = (c, \epsilon)$  denotes initial states including the input prompt c and gaussian noise  $\epsilon \sim \mathcal{N}(0, 1)$ ;  $\pi_{\theta}$  denotes the policy network, e.g. our TPDM model.  $\pi_{\text{old}}$  denotes the old policy for sampling trajectories;  $\pi_{\text{ref}}$  denotes a reference policy for regularization; y denotes the action our policy takes, i.e., the scheduled time; and  $\hat{A}(s, y)$  denotes the advantage the action y has given the state s.

We will specify the above elements in the PPO including the action and advantage we use in the following.

**Treat the whole schedule as an action** Usually, when the model makes a sequence of predictions, PPO regards every single prediction as an action and optimizes with them as a batch. Recently, RLOO [1] claims that when the reward signal only appears at the end of the sequence and the environment dynamics is fully deterministic, we can regard the whole sequence as an action with no harm to the performance. Thus for simplicity, the *entire generation*, including all predicted time in the schedule, is considered as a single action to optimize.

Thus, when computing expectation in Eq. 7, we consider the whole trajectory as a training sample in optimization. However, TPDM only outputs the distribution of each single time prediction, which we denote by  $\pi_{\theta}^{(i)}$ . By the chain rule, the probability of the entire generation can be calculated as the product of each prediction as in Eq. 8,

$$\pi_{\theta}(t_1, ..., t_N | s) = \prod_{i=1}^{N} \pi_{\theta}^{(i)}(t_i | s, t_1, ..., t_{i-1}), \quad (8)$$

where N denotes the total number of the generation steps. In our model, each factor of the probabilistic policy network  $\pi_{\theta}^{(i)}$  is computed through the TPM that outputs a distribution over the possible diffusion time  $t_i$  for the next step.

Image reward discounted by the number of steps When generating samples for policy gradient training, we obtain trajectories by generating images from Gaussian noise with the current TPDM policy. Since there exists no ground truth for these generated images, we choose ImageReward [54] to assign a reward score based solely on the final image in the last time. Considering we do not want a trivially large number of too many steps during generation, we apply a discount factor  $\gamma < 1$  to discount reward to intermediate diffusion steps and calculate the average over the total number of denoising steps N. The final reward for this trajectory is shown in Eq. 9 below,

$$R(s,y) = \frac{1}{N} \sum_{k=1}^{N} \gamma^{k-1} \operatorname{IR}(y,s),$$
(9)

where IR denotes the image reward model.

We refer readers to RLOO [1] on how to compute the advantage  $\hat{A}(s, y)$  directly from a batch R(s, y) without a value model for optimization. This reward function would

encourage TPDM to generate images with high quality, and reduce the number of diffusion steps for model efficiency as well. By adjusting  $\gamma$ , one can control how many diffusion steps tend to be applied to generate high-quality images. The smaller the value of  $\gamma$ , the fewer steps would tend to be allowed. We will elaborate more in Sec. 4.2.

## 4. Experiments

#### 4.1. Implementation details

**Dataset** We collect prompts to train our model. These prompts are generated with Florence-2 [53] and Llava-Next [24] to generate captions for the Laion-Art [48] and COYO-700M [4] datasets, and utilize these prompts to constitute our training set. We will elaborate on this in Appendix.

**Training Configurations** We use the AdamW optimizer with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.99$ , a constant learning rate of 1e-5, and a maximum gradient norm of 1.0. Our TPM module typically only requires around 200 training steps with a batch size of 256.

#### 4.2. Main Results

**Dynamic Schedule for Different Images** In Fig. 4, we present images generated with different prompts and their corresponding schedules predicted by TPDM. When prompting TPDM with shorter and simpler prompts, there exists fewer objects and details in the generated image, thus the diffusion time decreases faster and reaches 0 in relatively fewer steps. On the contrary, when longer and more complex prompts are provided, the model needs to generate more visual details. Therefore, the diffusion time decreases slower to generate delicate details. In this case TPDM requires more denoising steps in the generation process.

Adjusting  $\gamma$  for Different Number of Steps  $\gamma$  in Eq. 9 controls how the image reward discounts with more generation steps, influencing how fast diffusion time decays during the denoising process, thereby affecting the average number of denoising steps of our model.

As shown in Fig. 6, as we decrease  $\gamma$  from 0.97 to 0.85, TPDM tends to decrease the diffusion time more rapidly, leading to fewer sampling steps, from 15.0 to 7.5. Additionally, when compared to the baseline (yellow line), TPDM (purple line) consistently achieves a significantly higher aesthetic score with the same number of inference steps, achieving a good balance between model efficiency and generation performance.

**Visual Comparison** Our method achieves superior ability in generating fine-grained details. Images generated by TPDM have a more realistic laptop keyboard than both images from SD3-Medium in Fig. 5C, and the results in

Models	Inference Steps	FID	CLIP-T	Aesthetic v2	Pick Score	HPSv2.1
SD3-Medium [8]	28	25.00	0.322	<u>5.433</u>	22.12	29.12
SD3-Medium	15	24.72	0.321	5.426	22.30	28.52
TPDM-SD3-Medium	15.28	25.26	0.322	5.445	22.33	29.59
SD3.5-Large [2]	28	23.29	0.318	5.487	22.81	30.85
SD3.5-Large	15	<u>23.35</u>	0.323	5.475	22.62	30.02
TPDM-SD3.5-Large	15.22	24.48	<u>0.322</u>	5.525	22.81	<u>30.64</u>
FLUX.1-dev [19]	28	29.09	0.308	5.622	23.03	31.94
FLUX.1-dev	15	29.10	0.306	5.613	22.86	31.14
TPDM-FLUX.1-dev	13.57	28.98	0.314	5.685	<u>22.94</u>	30.77

Table 1. Evaluation of models across various benchmarks, best, second best

Fig.5D also demonstrate that TPDM can generate images without errors. This may due to that TPDM can adjust the diffusion schedule according to the denoising quality.

## 4.3. Quantitative Results

We apply TPM on several state-of-the-art diffusion models, including Stable Diffusion 3 Medium, Stable Diffusion 3.5 Large, and Flux 1.0 dev, and demonstrate how TPM could enhance their performance. We primarily evaluate metrics from two kinds: The first is objective metrics including FID, alignment with the given prompts (CLIP-T), and human preference scores (Aesthetic Score v2, and HPSv2.1). The second is directly comparing images generated from different models via user study.

**Quantitative Metrics** We compared TPDM and the above-mentioned models in Tab. 1. While maintaining competitive performance. All these models can generate images in about half of their recommended steps on average.

Moreover, the metrics representing human preference improve the most. For example, by generating with just 15.28 steps on average, TPDM-SD3-Medium gets a 29.59 HPS score, +1.07 higher than Stable Diffusion 3 with similar steps, also +0.47 higher than the original 28-step results. This may attributes to the reward model we utilized to calculate advantage in optimization, leading to aesthetically pleasing images aligned with human preferences.

**User Study** To better reflect human attitudes towards these models, we conduct a user study by inviting volunteers to compare images generated from different models and select the one they prefer.

Specifically, for each prompt, we provide two images from SD3-Medium generated with 15 and 28 steps separately, and image generated from TPDM-SD3-Medium. We invited 15 volunteers to evaluate images generated from 50 prompts. The result is shown in Tab. 2, indicating that our model can generate images that better align human preferences.

Models	Inference Steps	Win Rate	
SD3-Medium	28	26.58%	
SD3-Medium	15	16.40%	
TPDM-SD3-Medium	15.28	47.25%	

Table 2. User study based on Stable Diffusion 3 architecture

TPM inputs	Steps	FID	CLIP-T	PickScore	Aes v2
First 2 Layers	16.20	24.81	0.321	22.18	5.400
Last 2 Layers	19.30	23.19	0.322	22.16	5.356
First & Last Layers	15.28	25.26	0.322	22.33	5.445

Table 3. Ablation Study of inputs hidden states

#### 4.4. Ablation on Module Architectures

In this section, we ablate on the choice of the inputs of TPM. As shown in Tab. 3, taking features from both the first and last layer into TPM performs better than only taking either of them.

## 5. Conclusions and Limitations

In this paper, we introduce the Time Prediction Diffusion Model (TPDM), a text-to-image diffusion model that has flexible denoising schedulers for different individual prompts. By introducing a Time Prediction Module, we effectively train TPDM by reinforcement learning guided by a reward model. Based on the current leading diffusion model architecture proposed in Stable Diffusion 3 Medium, we train a strong MM-DiT-based TPDM that shows competitive quantitative performances on multiple text-to-image generation benchmarks.

Despite its promising performances, TPDM has some limitations. For instance, in this paper, we only design a relatively simple architecture for TPM. Whether and how to improve such a module to obtain better performances remains unexplored. Second, we froze the parameters of the original model and employed our training method to update the parameters of the diffusion model, leading to improved results. This approach warrants further exploration.

#### References

- [1] Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin, Ahmet Üstün, and Sara Hooker. Back to Basics: Revisiting REINFORCE Style Optimization for Learning from Human Feedback in LLMs, 2024. 3, 7
- Stability AI. Stable diffusion 3.5 large. https://huggingface.co/stabilityai/stable-diffusion-3.5-large, 2024. 8
- [3] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, Varun Jampani, and Robin Rombach. Stable video diffusion: Scaling latent video diffusion models to large datasets, 2023. 1, 3
- [4] Minwoo Byeon, Beomhee Park, Haecheon Kim, Sungjun Lee, Woonhyuk Baek, and Saehoon Kim. Coyo-700m: Image-text pair dataset. https://github.com/ kakaobrain/coyo-dataset, 2022. 7
- [5] CaptionEmporium. coyo-hd-11m-llavanext. https://huggingface.co/datasets/CaptionEmporium/coyohd-11m-llavanext, 2024. 1
- [6] Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Yue Wu, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, and Zhenguo Li. PixArt- $\alpha$ : Fast Training of Diffusion Transformer for Photorealistic Text-to-Image Synthesis, 2023. 4
- [7] Wei Deng, Weijian Luo, Yixin Tan, Marin Biloš, Yu Chen, Yuriy Nevmyvaka, and Ricky TQ Chen. Variational schr\" odinger diffusion models. *arXiv preprint arXiv:2405.04795*, 2024. 3
- [8] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, Kyle Lacey, Alex Goodwin, Yannik Marek, and Robin Rombach. Scaling Rectified Flow Transformers for High-Resolution Image Synthesis, 2024. 3, 4, 8
- [9] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, Kyle Lacey, Alex Goodwin, Yannik Marek, and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis, 2024. 3
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models, 2020. 1, 3
- [11] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. arXiv preprint arXiv:2204.03458, 2022. 1, 3
- [12] Emiel Hoogeboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3d. In *International Conference on Machine Learning*, pages 8867–8887. PMLR, 2022. 3
- [13] Zemin Huang, Zhengyang Geng, Weijian Luo, and Guojun Qi. Flow generator matching. arXiv preprint arXiv:2410.19310, 2024. 1, 3

- [14] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 8110–8119, 2020. 1
- [15] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Proc. NeurIPS*, 2022. 1
- [16] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the Design Space of Diffusion-Based Generative Models, 2022. 1
- [17] Heeseung Kim, Sungwon Kim, and Sungroh Yoon. Guidedtts: A diffusion model for text-to-speech via classifier guidance. In *International Conference on Machine Learning*, pages 11119–11133. PMLR, 2022. 3
- [18] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. In *International Conference on Learning Representations*, 2020. 1
- [19] Black Forest Lab. Announcing Black Forest Labs. https://blackforestlabs.ai/announcing-black-forest-labs/, 2024. 3, 8
- [20] Black Forest Labs. Flux.1, 2023. 4
- [21] Youwei Liang, Junfeng He, Gang Li, Peizhao Li, Arseniy Klimovskiy, Nicholas Carolan, Jiao Sun, Jordi Pont-Tuset, Sarah Young, Feng Yang, Junjie Ke, Krishnamurthy Dj Dvijotham, Katie Collins, Yiwen Luo, Yang Li, Kai J. Kohlhoff, Deepak Ramachandran, and Vidhya Navalpakkam. Rich Human Feedback for Text-to-Image Generation, 2024. 3
- [22] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling, 2023. 3
- [23] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow Matching for Generative Modeling, 2023. 3
- [24] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge, 2024. 7
- [25] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow, 2022. 3
- [26] Weijian Luo. A comprehensive survey on knowledge distillation of diffusion models. arXiv preprint arXiv:2304.04262, 2023. 3
- [27] Weijian Luo. Diff-instruct++: Training one-step text-toimage generator model to align with human preferences. arXiv preprint arXiv:2410.18881, 2024. 3
- [28] Weijian Luo, Tianyang Hu, Shifeng Zhang, Jiacheng Sun, Zhenguo Li, and Zhihua Zhang. Diff-instruct: A universal approach for transferring knowledge from pre-trained diffusion models. *ArXiv*, abs/2305.18455, 2023. 1
- [29] Weijian Luo, Tianyang Hu, Shifeng Zhang, Jiacheng Sun, Zhenguo Li, and Zhihua Zhang. Diff-instruct: A universal approach for transferring knowledge from pre-trained diffusion models. Advances in Neural Information Processing Systems, 36, 2024. 3

- [30] Weijian Luo, Zemin Huang, Zhengyang Geng, J Zico Kolter, and Guo-jun Qi. One-step diffusion distillation through score implicit matching. *arXiv preprint arXiv:2410.16794*, 2024. 1, 3
- [31] Weijian Luo, Boya Zhang, and Zhihua Zhang. Entropybased training methods for scalable neural implicit samplers. *Advances in Neural Information Processing Systems*, 36, 2024. 1
- [32] Weijian Luo, Colin Zhang, Debing Zhang, and Zhengyang Geng. Diff-instruct\*: Towards human-preferred onestep text-to-image generative models. arXiv preprint arXiv:2410.20898, 2024. 3
- [33] George A Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41, 1995. 1
- [34] Alex Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. *arXiv preprint arXiv:2102.09672*, 2021. 3
- [35] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016. 1, 3
- [36] Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer, 2017. 4
- [37] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. arXiv preprint arXiv:2209.14988, 2022. 3
- [38] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021.
- [39] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 1
- [40] Amirmojtaba Sabour, Sanja Fidler, and Karsten Kreis. Align Your Steps: Optimizing Sampling Schedules in Diffusion Models, 2024. 1, 3
- [41] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. arXiv preprint arXiv:2205.11487, 2022. 1
- [42] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy Optimization Algorithms, 2017. 3, 6
- [43] Shuai Shen, Wenliang Zhao, Zibin Meng, Wanhua Li, Zheng Zhu, Jie Zhou, and Jiwen Lu. Difftalk: Crafting diffusion models for generalized audio-driven portraits animation, 2023. 1
- [44] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics, 2015. 1, 3
- [45] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models, 2022. 3

- [46] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations, 2021. 1
- [47] Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. Learning to summarize with human feedback. In Advances in Neural Information Processing Systems, pages 3008–3021. Curran Associates, Inc., 2020. 3
- [48] Laion Team. Laion-art. https://huggingface.co/datasets/laion/laion-art, 2024. 7, 1
- [49] Yunke Wang, Xiyu Wang, Anh-Dung Dinh, Bo Du, and Charles Xu. Learning to schedule in diffusion probabilistic models. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, page 2478–2488, New York, NY, USA, 2023. Association for Computing Machinery. 3
- [50] Yifei Wang, Weimin Bai, Weijian Luo, Wenzheng Chen, and He Sun. Integrating amortized inference with diffusion models for learning clean distribution from corrupted images. arXiv preprint arXiv:2407.11162, 2024. 3
- [51] Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li. Human Preference Score v2: A Solid Benchmark for Evaluating Human Preferences of Text-to-Image Synthesis, 2023. 3
- [52] Mengfei Xia, Yujun Shen, Changsong Lei, Yu Zhou, Ran Yi, Deli Zhao, Wenping Wang, and Yong-jin Liu. Towards More Accurate Diffusion Model Acceleration with A Timestep Aligner, 2023. 1, 3
- [53] Bin Xiao, Haiping Wu, Weijian Xu, Xiyang Dai, Houdong Hu, Yumao Lu, Michael Zeng, Ce Liu, and Lu Yuan. Florence-2: Advancing a unified representation for a variety of vision tasks. *arXiv preprint arXiv:2311.06242*, 2023. 7, 1
- [54] Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. ImageReward: Learning and Evaluating Human Preferences for Text-to-Image Generation, 2023. 2, 3, 7
- [55] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, Da Yin, Xiaotao Gu, Yuxuan Zhang, Weihan Wang, Yean Cheng, Ting Liu, Bin Xu, Yuxiao Dong, and Jie Tang. Cogvideox: Text-to-video diffusion models with an expert transformer, 2024. 1, 3
- [56] Zilyu Ye, Jinxiu Liu, Ruotian Peng, Jinjin Cao, Zhiyang Chen, Yiyang Zhang, Ziwei Xuan, Mingyuan Zhou, Xiaoqian Shen, Mohamed Elhoseiny, Qi Liu, and Guo-Jun Qi. Openstory++: A large-scale dataset and benchmark for instance-aware open-domain visual storytelling, 2024. 3
- [57] Boya Zhang, Weijian Luo, and Zhihua Zhang. Enhancing adversarial robustness via score-based optimization. Advances in Neural Information Processing Systems, 36:51810–51829, 2023. 3

# Schedule On the Fly: Diffusion Time Prediction for Faster and Better Image Generation

Supplementary Material

## **A. Dataset Details**

In Section 4, we provide a brief overview of our training dataset. Below, we present a more detailed description of the dataset.

## A.1. Dataset Source

We curated prompts from image captions in two highquality datasets, Laion-art [48]and COYO-11M [5]. The images in Laion-art are recaptioned with Florence-2 [53], while COYO-11M retained its original captions.

To ensure diversity during training, we design a filtering pipeline to select a diverse subset of these prompts.

## **A.2. Filtering Pipeline**



Figure 7. The filtering pipeline used to construct the training set.

Our filtering pipeline, depicted in Figure 7, aims to select diverse and high-quality prompts by analyzing key linguistic features, such as nouns, prepositions, and adjectives. To assess diversity, we utilize *WordNet* [33] to count the number of valid nouns, adjectives, and prepositions in each prompt.

Prompts are ranked based on these counts. Those exhibiting the highest scores are included in the training set. The selection process prioritizes **noun-diversity**, ensuring a balanced representation across categories such as *Person*, *Animal*, *Plant*, *Artifacts* (*Large/Small Objects*), and *Natural Views*. Next, prompts with high **preposition-diversity** are selected, emphasizing those that contain spatial and relational terms (e.g., near, on). Finally, prompts are evaluated for **adjective-diversity**, with particular focus on adjectives describing *color* and *shape*, to enhance descriptive richness.

The process is iterative, selecting the top prompts in order of their noun, adjective, and preposition diversity scores until the desired data quantity (51,200 prompts, as specified in Section 4.1) and diversity are achieved. As shown in Figure 8, the diversity of the training set has significantly enhanced after filtering.

## **B.** Prompts of Figure

## **B.1. Figure 1**

- 1. 8k uhd A man looks up at the starry sky, lonely and ethereal, Minimalism, Chaotic composition Op Art.
- 2. A deep-sea exploration vessel descending into the pitchblack ocean, its powerful lights illuminating the glowing, alien creatures that inhabit the abyss. A massive, ancient sea creature with bioluminescent patterns drifts into view, its eyes glowing as it watches the explorers from the shadows of an underwater cave.
- 3. Half human, half robot, repaired human.
- 4. A baby painter trying to draw very simple picture, white background.
- 5. an astronaut sitting in a diner, eating fries, cinematic, analog film.
- 6. Van Gogh painting of a teacup on the desk.
- 7. A galaxy scene with stars, planets, and nebula clouds.
- 8. A hidden, forgotten city deep in a jungle, with crumbling stone temples overgrown by thick vines. In the heart of the city, a mysterious glowing artifact lies on an ancient pedestal, surrounded by an eerie mist. Strange symbols shimmer faintly on the stone walls, waiting to be uncovered.
- 9. A lone astronaut stranded on a desolate planet, gazing up at the sky. The planet's surface is cracked and barren, with glowing, unearthly ruins scattered across the horizon. In the distance, a massive, alien ship slowly descends, casting an eerie shadow over the landscape.



Figure 8. The statistics of prompts from the re-captioned high-resolution Laion-Art dataset, before and after filtering, highlight an improvement in diversity. Our dataset demonstrates greater variety compared to the original.