


# ControlNet++: Improving Conditional Controls with Efficient Consistency Feedback

Project Page: [liming-ai.github.io/ControlNet\\_Plus\\_Plus](https://liming-ai.github.io/ControlNet_Plus_Plus)

Ming Li<sup>1</sup>, Taojiannan Yang<sup>1</sup>, Huafeng Kuang<sup>2</sup>, Jie Wu<sup>2</sup>,  
Zhaoning Wang<sup>1</sup>, Xuefeng Xiao<sup>2</sup>, and Chen Chen<sup>1</sup> 

<sup>1</sup> Center for Research in Computer Vision, University of Central Florida

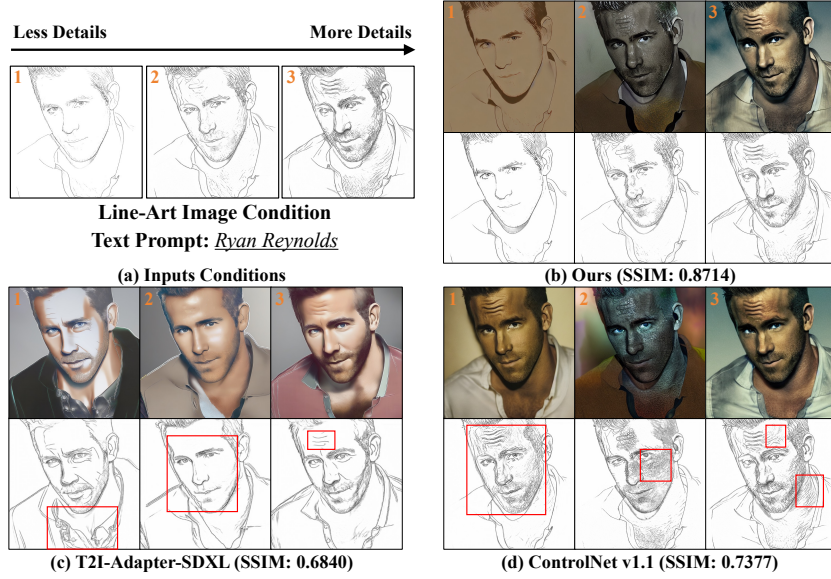
<sup>2</sup> ByteDance

**Abstract.** To enhance the controllability of text-to-image diffusion models, existing efforts like ControlNet incorporated image-based conditional controls. In this paper, we reveal that existing methods still face significant challenges in generating images that align with the image conditional controls. To this end, we propose ControlNet++, a novel approach that improves controllable generation by explicitly optimizing pixel-level cycle consistency between generated images and conditional controls. Specifically, for an input conditional control, we use a pre-trained discriminative reward model to extract the corresponding condition of the generated images, and then optimize the consistency loss between the input conditional control and extracted condition. A straightforward implementation would be generating images from random noises and then calculating the consistency loss, but such an approach requires storing gradients for multiple sampling timesteps, leading to considerable time and memory costs. To address this, we introduce an efficient reward strategy that deliberately disturbs the input images by adding noise, and then uses the single-step denoised images for reward fine-tuning. This avoids the extensive costs associated with image sampling, allowing for more efficient reward fine-tuning. Extensive experiments show that ControlNet++ significantly improves controllability under various conditional controls. For example, it achieves improvements over ControlNet by 11.1% mIoU, 13.4% SSIM, and 7.6% RMSE, respectively, for segmentation mask, line-art edge, and depth conditions. All the code, models, demo and organized data have been open sourced on our [Github Repo](https://github.com/liming-ai/ControlNet_Plus_Plus).

**Keywords:** Controllable Generation · Diffusion Model · ControlNet

## 1 Introduction

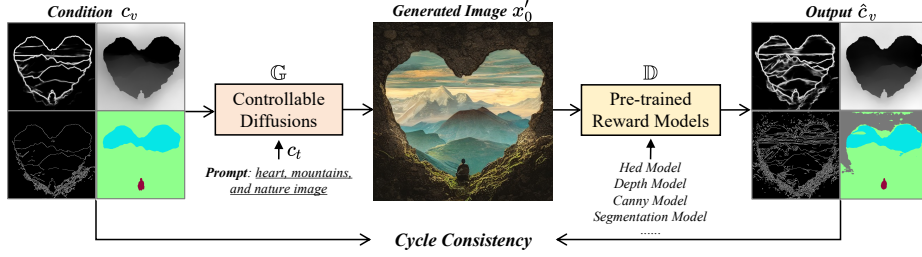
The emergence and improvements of diffusion models [11, 40, 47], along with the introduction of large-scale image-text datasets [45, 46], has catalyzed significant strides in text-to-image generation. Nonetheless, as the proverb “an image is worth a thousand words” conveys, it’s challenging to depict an image accurately and in detail through language alone, and this dilemma also perplexes existing text-to-image diffusion models [40, 43]. To this end, many studies focus on incorporating conditional controls such as segmentation mask into text-to-image diffusion models [21, 29, 35, 59, 60]. Despite the diversity in these methods, the core objective remains to facilitate more accurate and controllable image generation with explicit image-based conditional controls.



**Fig. 1:** (a) Given the same input image condition and text prompt, (b) the extracted conditions of our generated images are more consistent with the inputs, (c,d) while other methods fail to achieve accurate controllable generation. SSIM scores measure the similarity between all input edge conditions and the extracted edge conditions. All the line edges are extracted by the same line detection model used by ControlNet [60].

Achieving controllable image generation could involve retraining diffusion models from scratch [35, 40], but this comes with high computational demands and a scarcity of large public datasets [60]. In light of this, a feasible strategy is fine-tuning pre-trained text-to-image models [22, 58] or introducing trainable modules [29, 59, 60] like ControlNet [60]. However, despite these studies have explored the feasibility of controllability [29, 59, 60] in text-to-image diffusion models and expanded various applications [21, 22, 35], a significant gap remains in achieving precise and fine-grained control. As shown in Fig. 1, existing methods of controllable generation (e.g., ControlNet [60] and T2I-Adapter [29]) still struggle to accurately generate images that are consistent with the input image condition. For example, T2I-Adapter-SDXL consistently produced incorrect wrinkles on the forehead in all generated images, while ControlNet v1.1 introduced many wrong details. Regrettably, current efforts lack specific methods for improving controllability, which impedes progress in this research area.

To address this issue, we model image-based controllable generation as an image translation task [16] from input conditional controls to output generated images. Inspired by CycleGAN [68], we propose to employ pre-trained discriminative models to extract the condition from the generated images and directly optimize the cycle consistency loss for better controllability. The idea is that if we translate images from one domain to the other (condition  $c_v \rightarrow$  generated image  $x'_0$ ), and back again (generated image  $x'_0 \rightarrow$  condition  $c'_v$ ) we should arrive where we started ( $c'_v = c_v$ ), as shown in Fig. 2. For example, given a segmentation mask as a conditional control, we can employ existing methods such as Control-




**Fig. 2: Illustration of the cycle consistency.** We first prompt the diffusion model  $\mathbb{G}$  to generate an image  $x'_0$  based on the given image condition  $c_v$  and text prompt  $c_t$ , then extract the corresponding image condition  $\hat{c}_v$  from the generated image  $x'_0$  using pre-trained discriminative models  $\mathbb{D}$ . The cycle consistency is defined as the similarity between the extracted condition  $\hat{c}_v$  and input condition  $c_v$ .

Net [60] to generate corresponding images. Then the predicted segmentation masks of these generated images can be obtained by a pre-trained segmentation model. Ideally, the predicted segmentation masks and the input segmentation masks should be consistent. Hence, the cycle consistency loss can be formulated as the per-pixel classification loss between the input and predicted segmentation masks. Unlike existing related works [26, 29, 35, 60, 62] that implicitly achieve controllability by introducing conditional controls into the latent-space denoising process, our method **explicitly** optimizes controllability at the pixel-space for better performance, as demonstrated in Fig. 3.

To implement pixel-level loss within the context of diffusion models, an intuitive approach involves executing the diffusion model’s inference process, starting from random Gaussian noise and performing multiple sampling steps to obtain the final generated images, following recent works focusing on improving image quality with human feedback [10, 34, 57]. However, multiple samplings can lead to efficiency issues, and require the storage of gradients at every timestep and thus significant time and GPU memory consumption. We demonstrate that initiating sampling from random Gaussian noise is unnecessary. Instead, by directly adding noise to training images to disturb their consistency with input conditional controls and then using single-step denoised images to reconstruct the consistency, we can conduct more efficient reward fine-tuning. Our contributions are summarized as:

- *New Insight*: We reveal that existing efforts in controllable generation still perform poorly in terms of controllability, with generated images significantly deviating from input conditions and lacking a clear strategy for improvement.
- *Consistency Reward Feedback*: We show that pre-trained discriminative models can serve as powerful visual reward models to improve the controllability of controllable diffusion models in a cycle-consistency manner.
- *Efficient Reward Fine-tuning*: We disrupt the consistency between input images and conditions, and enable the single-step denoising for efficient reward fine-tuning, avoiding time and memory overheads caused by image sampling.
- *Evaluation and Promising Results*: We provide a unified and public evaluation of controllability under various conditional controls, and demonstrate that ControlNet++ comprehensively outperforms existing methods.



**Fig. 3:** (a) Existing methods achieve implicit controllability by introducing image-based conditional control  $c_v$  into the denoising process of diffusion models, with the guidance of latent-space denoising loss. (b) We utilize discriminative reward models  $\mathbb{D}$  to explicitly optimize the controllability of  $\mathbb{G}$  via pixel-level cycle consistency loss.

## 2 Related Work

### 2.1 Diffusion-based Generative Models

The diffusion probabilistic model presented in [47] has undergone substantial advancements [11, 18, 24], thanks to iterative refinements in training and sampling strategies [17, 48, 49]. To alleviate the computational demands for training diffusion models, Latent Diffusion [40] maps the pixel space diffusion process into the latent feature space. In the realm of text-to-image synthesis, diffusion models [30, 33, 38–40, 43] integrate cross-attention mechanisms between UNet [41] denoisers and text embeddings from pre-trained language models like CLIP [36] and T5 [37] to facilitate reasonable text-to-image generation. Furthermore, diffusion models are employed across image editing tasks [2, 13, 23, 28] by manipulating inputs [38], editing cross-attentions [15], and fine-tuning models [42]. Despite the astonishing capabilities of diffusion models, language is a sparse and highly semantic representation, unsuitable for describing dense, low-semantic images. Furthermore, existing methods [33, 40] still struggle to understand detailed text prompts, posing a severe challenge to the controllable generation [60].

### 2.2 Controllable Text-to-Image Diffusion Models

To achieve conditional control in pre-trained text-to-image diffusion models, ControlNet [60] and T2I-Adapter [29] introduce additional trainable modules for guided image generation. Furthermore, recent research employs various prompt engineering [26, 58, 61] and cross-attention constraints [5, 22, 55] for a more regulated generation. Some methods also explore multi-condition or multi-modal generation within a single diffusion model [20, 35, 62] or focus on the instance-based controllable generation [51, 66]. However, despite these methods exploring feasibility and applications, there still lacks a clear approach to enhance controllability under various controls. Furthermore, existing works implicitly learn controllability by the denoising process of diffusion models, while our ControlNet++ achieves this in an explicit cycle-consistency manner, as shown in Fig. 3.

### 2.3 Linguistic and Visual Reward Models

The reward model is trained to evaluate how well the results of generative models align with human expectations, and its quantified results will be used to



facilitate generative models for better and more controllable generation. It is usually trained with reinforcement learning from human feedback (RLHF) in NLP tasks [9, 31, 50], and has recently extended into the vision domain to improve the image quality for text-to-image diffusion models [1, 10, 12, 34, 53, 57]. However, image quality is an exceedingly subjective metric, fraught with individual preferences, and requires the creation of new datasets with human preferences [25, 52, 53, 57] and the training of reward models [34, 52, 57]. Diverging from the pursuit of global image quality with subjective human preference in current research, we target the more fine-grained and objective goal of controllability. Also, it's more cost-effective to obtain AI feedback compared to human feedback.

### 3 Method

In this section, we first introduce the background of diffusion models in Sec. 3.1. In Sec. 3.2, we discuss how to design the cycle consistency loss for controllable diffusion models to enhance the controllability. Finally, in Sec. 3.3, we examine the efficiency issues with the straightforward solution and correspondingly propose an efficient reward strategy that utilizes the single-step denoised images for consistency loss, instead of sampling images from random noise.

#### 3.1 Preliminary

The diffusion models [17] define a Markovian chain of diffusion forward process  $q(x_t|x_0)$  by gradually adding noise to input data  $x_0$ :

$$x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, I), \quad (1)$$

where  $\epsilon$  is a noise map sampled from a Gaussian distribution, and  $\bar{\alpha}_t := \prod_{s=0}^t \alpha_s$ .  $\alpha_t = 1 - \beta_t$  is a differentiable function of timestep  $t$ , which is determined by the denoising sampler such as DDPM [17]. To this end, the diffusion training loss can be represented by:

$$\mathcal{L}(\epsilon_\theta) = \sum_{t=1}^T \mathbb{E}_{x_0 \sim q(x_0), \epsilon \sim \mathcal{N}(\mathbf{0}, I)} \left[ \|\epsilon_\theta(\sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon) - \epsilon\|_2^2 \right]. \quad (2)$$

In the context of controllable generation [29, 60], with given image condition  $c_v$  and text prompt  $c_t$ , the diffusion training loss at timestep  $t$  can be re-written as:

$$\mathcal{L}_{\text{train}} = \mathbb{E}_{x_0, t, c_t, c_v, \epsilon \sim \mathcal{N}(0, I)} \left[ \|\epsilon_\theta(x_t, t, c_t, c_v) - \epsilon\|_2^2 \right]. \quad (3)$$

During the inference, given a random noise  $x_T \sim \mathcal{N}(\mathbf{0}, I)$ , we can predict final denoised image  $x_0$  with the step-by-step denoising process [17]:

$$x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left( x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(\mathbf{x}_t, t) \right) + \sigma_t \epsilon, \quad (4)$$

where  $\epsilon_\theta$  refers to the predicted noise at timestep  $t$  by U-Net [41] with parameters  $\theta$ , and  $\sigma_t = \frac{1 - \bar{\alpha}_t - 1}{1 - \bar{\alpha}_t} \beta_t$  is the variance of posterior Gaussian distribution  $p_\theta(x_0)$ .

### 3.2 Reward Controllability with Consistency Feedback

As we model controllability as the consistency between input conditions and the generated images, we can naturally quantify this outcome through the discriminative reward models. Once we quantify the results of the generative model, we can perform further optimization for more controllable generation based on these quantified results in a unified manner for various conditional controls.

To be more specific, we minimize the consistency loss between the input condition  $c_v$  and the corresponding output condition  $\hat{c}_v$  of the generated image  $x'_0$ , as depicted in Fig. 2. The reward consistency loss can be formulated as:

$$\begin{aligned}\mathcal{L}_{\text{reward}} &= \mathcal{L}(c_v, \hat{c}_v) \\ &= \mathcal{L}(c_v, \mathbb{D}(x'_0)) \\ &= \mathcal{L}(c_v, \mathbb{D}[\mathbb{G}^T(c_t, c_v, x_T, t)]),\end{aligned}\tag{5}$$

where  $\mathbb{G}^T(c_t, c_v, x_T, t)$  denotes the process that the model performs  $T$  **denoising steps** to generate the image  $x'_0$  from random noise  $x_T$ , as shown in the Fig. 4 (a). Here,  $\mathcal{L}$  is an abstract metric function that can take on different concrete forms for different visual conditions. For example, in the context of using segmentation mask as the input conditional control,  $\mathcal{L}$  could be the per-pixel cross-entropy loss. The reward model  $\mathcal{D}$  is also dependent on the condition, and we use the UperNet [54] for segmentation mask conditions. The details of loss functions and reward models are summarized in the supplementary material.

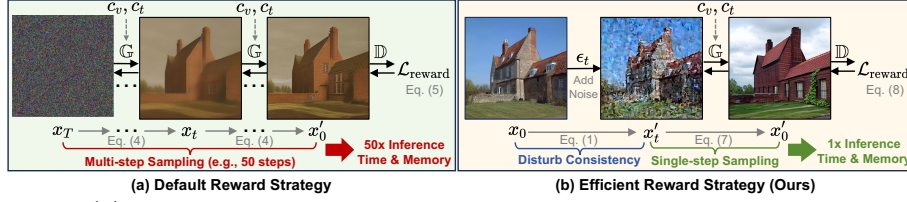
In addition to the reward loss, we also employ diffusion training loss in Eq. 3 to ensure that the original image generation capability is not compromised since they have different optimization goals. Finally, the total loss is the combination of  $\mathcal{L}_{\text{train}}$  and  $\mathcal{L}_{\text{reward}}$ :

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{train}} + \lambda \cdot \mathcal{L}_{\text{reward}},\tag{6}$$

where  $\lambda$  is a hyper-parameter to adjust the weight of the reward loss. Through this approach, the consistency loss can guide the diffusion model on how to sample at different timesteps to obtain images more consistent with the input controls, thereby enhancing controllability. Nonetheless, directly applying such reward consistency still poses challenges in efficiency in real-world settings.

### 3.3 Efficient Reward Fine-tuning

To achieve the pixel-space consistency loss  $\mathcal{L}_{\text{reward}}$ , it requires  $x_0$ , the final diffused image, to calculate the reward consistency from the reward models. As modern diffusion models, such as Stable Diffusion [40], require multiple steps, e.g., 50 steps, to render a full image, directly using such a solution is impractical in realistic settings: (1) multiple time-consuming samplings are required to derive images from random noise. (2) to enable gradient backpropagation, we have to store gradients at each timestep, meaning the GPU memory usage will increase linearly with the number of time-steps. Taking ControlNet as an example,



**Fig. 4:** (a) Pipeline of default reward fine-tuning strategy. Reward fine-tuning requires sampling all the way to the full image. Such a method needs to keep all gradients for each timestep and the memory required is unbearable by current GPUs. (b) Pipeline of our efficient reward strategy. We add a small noise ( $t \leq t_{thre}$ ) to disturb the consistency between input images and conditions, then the single-step denoised image can be directly used for efficient reward fine-tuning.

when the batch size is 1 with FP16 mixed precision, the GPU memory required for a single denoising step and storing all training gradients is approximately 6.8GB. If we use the 50-step inference with the DDIM [48] scheduler, approximately 340GB of memory is needed to perform reward fine-tuning on a single sample, which is nearly impossible to achieve with current hardware capabilities. Although GPU memory consumption can be reduced by employing techniques such as Low-Rank Adaptation (LoRA) [10, 19], gradient checkpointing [6, 10], or stop-gradient [57], the efficiency degradation caused by the number of sampling steps required to generate images remains significant and cannot be overlooked. Therefore, an efficient reward fine-tuning approach is necessary.

In contrast to diffusing from random noise  $x_T$  to obtain the final image  $x_0$ , as illustrated in Fig. 4 (a), we instead propose an *one-step* efficient reward strategy. Specifically, instead of randomly sampling from noise, we add noise to the training images  $x_0$ , thereby explicitly disturbing the consistency between the diffusion inputs  $x_t'$  and their conditional controls  $c_v$ , by performing diffusion forward process  $q(x_t|x_0)$  in Eq. 1. We demonstrate this process as the **Disturb Consistency** in Fig. 4 (b), which is the same procedure as the standard diffusion training process. When the added noise  $\epsilon$  is relatively small, we can predict the original image  $x_0'$  by performing single-step sampling<sup>3</sup> on disturbed image  $x_t'$  [17]:

$$x_0 \approx x_0' = \frac{x_t' - \sqrt{1 - \alpha_t} \epsilon_\theta(x_t', c_v, c_t, t)}{\sqrt{\alpha_t}}, \quad (7)$$

and then we directly utilize the denoised image  $x_0'$  to perform reward fine-tuning:

$$\mathcal{L}_{reward} = \mathcal{L}(c_v, \hat{c}_v) = \mathcal{L}(c_v, \mathbb{D}(x_0')) = \mathcal{L}(c_v, \mathbb{D}[G(c_t, c_v, x_t', t)]). \quad (8)$$

Essentially, the process of adding noise destroys the consistency between the input image and its condition. Then the reward fine-tuning in Eq. 8 instructs the diffusion model to generate images that can reconstruct the consistency, thus enhancing its ability to follow conditions during generation.

<sup>3</sup> We provide a more detailed proof in the supplementary material.

Please note that here we avoid the sampling process in Eq. 5. Finally, the loss is the combination of diffusion training loss and the reward loss:

$$\mathcal{L}_{\text{total}} = \begin{cases} \mathcal{L}_{\text{train}} + \lambda \cdot \mathcal{L}_{\text{reward}}, & \text{if } t \leq t_{\text{thre}}, \\ \mathcal{L}_{\text{train}}, & \text{otherwise,} \end{cases} \quad (9)$$

where  $t_{\text{thre}}$  denotes the timestep threshold, which is a hyper-parameter used to determine whether a noised image  $x_t$  should be utilized for reward fine-tuning. We note that a small noise  $\epsilon$  (i.e., a relatively small timestep  $t$ ) can disturb the consistency and lead to effective reward fine-tuning. When the timestep  $t$  is large,  $x_t$  is closer to the random noise  $x_T$ , and predicting  $x'_0$  directly from  $x_t$  results in severe image distortion. The advantage of our efficient rewarding is that  $x_t$  can be employed both to train and reward the diffusion model without the need for time and GPU memory costs caused by multiple sampling, thereby significantly improving the efficiency during the reward fine-tuning stage.

During the reward fine-tuning phases, we freeze the pre-trained discriminative reward model and text-to-image model, and only update the ControlNet following its original implementation, which ensures the generative capabilities are not compromised. We also observe that using only the reward loss will lead to image distortion, aligning with the conclusions drawn in previous studies [57].

## 4 Experiments

### 4.1 Experimental Setup

*Condition Controls and Datasets.* Given that existing text-image paired datasets for generative models are unable to provide accurate conditional control data pairs [45, 46], such as image-segmentation pairs, we endeavor to select specific datasets for different tasks that can offer more precise image-label data pairs. More specifically, ADE20K [64, 65] and COCOStuff [3] are used for the segmentation mask condition following ControlNet [60]. For the canny edge map, hed edge map, lineart map, and depth map condition, we utilize the MultiGen-20M dataset proposed by UniControl [35], which is a subset of LAION-Aesthetics [45]. For the datasets without text caption such as ADE20K, we utilize MiniGPT-4 [67] to generate the image caption with the instruction “*Please briefly describe this image in one sentence*”. The training and inference resolution is  $512 \times 512$  for all datasets and methods. Details are provided in the supplementary material.

*Evaluation and Metrics.* We train ControlNet++ on the training set of each corresponding dataset and evaluate all methods on the validation dataset. All the experiments are evaluated under  $512 \times 512$  resolution for fair comparison. For each condition, we evaluate the controllability by measuring the similarity between the input conditions and the extracted conditions from generated images of diffusion models. For semantic segmentation and depth map controls, we use mIoU and RMSE as evaluation metrics respectively, which is a common practice in related research fields. For the edge task, we use F1-Score for hard edges (canny

**Table 1:** Controllability comparison with state-of-the-art methods under different conditional controls and datasets.  $\uparrow$  denotes higher result is better, while  $\downarrow$  means lower is better. ControlNet++ achieves significant controllability improvements. ‘-’ indicates that the method does not provide a public model for testing. We generate four groups of images in png format and report the average result to reduce random errors.

Condition (Metric)	T2I Model	Seg. Mask (mIoU $\uparrow$ )		Canny Edge (F1 Score $\uparrow$ )	Hed Edge (SSIM $\uparrow$ )	LineArt Edge (SSIM $\uparrow$ )	Depth Map (RMSE $\downarrow$ )
Dataset		ADE20K	COCO-Stuff	MultiGen-20M	MultiGen-20M	MultiGen-20M	MultiGen-20M
ControlNet	SDXL	-	-	-	-	-	40.00
T2I-Adapter	SDXL	-	-	28.01	-	0.6394	39.75
T2I-Adapter	SD1.5	12.61	-	23.65	-	-	48.40
Gligen	SD1.4	23.78	-	26.94	0.5634	-	38.83
Uni-ControlNet	SD1.5	19.39	-	27.32	0.6910	-	40.65
UniControl	SD1.5	25.44	-	30.82	0.7969	-	39.18
ControlNet	SD1.5	32.55	27.46	34.65	0.7621	0.7054	35.90
<b>Ours</b>	SD1.5	<b>43.64</b>	<b>34.56</b>	<b>37.04</b>	<b>0.8097</b>	<b>0.8399</b>	<b>28.32</b>

edge) because it can be regarded as a binary classification problem of 0 (non-edge) and 1 (edge) and has a serious long-tail distribution, following the standard evaluation in edge detection [56]. The threshold used for evaluation is (100, 200) for OpenCV, and (0.1, 0.2) for Kornia implementation. The SSIM metric is used for the soft edges conditional controls (i.e., hed edge & lineart edge) following previous works [62]. For ControlNet++, we use the UniPC [63] sampler with 20 denoising steps to generate images with the original text prompt following ControlNet v1.1 [60], without any negative prompts. For other methods beyond ControlNet and ours, we utilized their open-source code to generate images and conducted fair evaluations under the same data, without changing their inference configures such as the number of inference steps or denoising sampler.

*Baselines.* Our evaluation primarily focuses on T2I-Adapter [29], ControlNet v1.1 [60], GLIGEN [26], Uni-ControlNet [62], and UniControl [35], as these methods are pioneering in the realm of controllable text-to-image diffusion models and offer public model weights for various image conditions. To ensure fairness of evaluation, all methods use the same image conditions and text prompts. While most methods employ the user-friendly SD1.5 as their text-to-image model for controllable generation, we have observed that recently there are a few models based on SDXL [33]. Therefore, we also report the controllability results for ControlNet-SDXL and T2I-Adapter-SDXL. Please note that ControlNet-SDXL mentioned here is not an officially released model as in ControlNet [60].

## 4.2 Experimental Results

*Comparison of Controllability with State-of-the-art Methods.* The experimental results are shown in Tab. 1, which can be summarized as the following observations: (1) Existing methods still underperform in terms of controllability, struggling to achieve precise controlled generation. For instance, current methods (i.e., ControlNet) achieve only a 32.55 mIoU for images generated under the condition of segmentation masks, which is far from its performance on real datasets with a 50.7 mIoU, under the same evaluation from Mask2Former segmentation model [7]. (2) Our ControlNet++ significantly outperforms existing

**Table 2:** FID ( $\downarrow$ ) comparison with state-of-the-art methods under different conditional controls and datasets. All the results are conducted on  $512 \times 512$  image resolution with Clean-FID implementation [32] for fair comparisons. ‘-’ indicates that the method does not provide a public model for testing. We generate four groups of images in png format and report the average result to reduce random errors.

Method	T2I Model	Seg. Mask		Canny Edge	Hed Edge	LineArt Edge	Depth Map
		ADE20K	COCO	MultiGen-20M	MultiGen-20M	MultiGen-20M	MultiGen-20M
Gligen	SD1.4	33.02	-	18.89	-	-	18.36
T2I-Adapter	SD1.5	39.15	-	15.96	-	-	22.52
UniControlNet	SD1.5	39.70	-	17.14	17.08	-	20.27
UniControl	SD1.5	46.34	-	19.94	15.99	-	18.66
ControlNet	SD1.5	33.28	21.33	<b>14.73</b>	15.41	17.44	17.76
<b>Ours</b>	SD1.5	<b>29.49</b>	<b>19.29</b>	18.23	<b>15.01</b>	<b>13.88</b>	<b>16.66</b>

**Table 3:** CLIP-score ( $\uparrow$ ) comparison with state-of-the-art methods under different conditional controls and datasets. ‘-’ indicates that the method does not provide a public model for testing. We generate four groups of images in png format and report the average result to reduce random errors.


Method	T2I Model	Seg. Mask		Canny Edge	Hed Edge	LineArt Edge	Depth Map
		ADE20K	COCO	MultiGen-20M	MultiGen-20M	MultiGen-20M	MultiGen-20M
Gligen	SD1.4	31.12	-	31.77	-	-	31.75
T2I-Adapter	SD1.5	30.65	-	31.71	-	-	31.46
UniControlNet	SD1.5	30.59	-	31.84	31.94	-	31.66
UniControl	SD1.5	30.92	-	31.97	32.02	-	<b>32.45</b>
ControlNet	SD1.5	31.53	<b>13.31</b>	<b>32.15</b>	<b>32.33</b>	<b>32.46</b>	<b>32.45</b>
<b>Ours</b>	SD1.5	<b>31.96</b>	13.13	31.87	32.05	31.95	32.09

works in terms of controllability across various conditional controls. For example, it achieves 11.1% RMSE improvements against previous state-of-the-art methods for the depth map condition; (3) For controllable diffusion models, the strength of the text-to-image backbone does not affect its controllability. As shown in the table, although SDXL-based [33] ControlNet and T2I-Adapter have better controllability on some specific tasks, the improvement is not large and is not significantly better than the counterparts with SD 1.5 [40].

*Comparison of Image Quality with State-of-the-art Methods.* To verify whether improving controllability leads to a decline in image quality, we reported the FID (Fréchet Inception Distance) metrics of different methods under various conditional generation tasks in Tab. 2. We discovered that, compared to existing methods, ControlNet++ generally exhibits superior FID values in most cases, indicating that our approach, while enhancing the controllability of conditional controls, does not result in a decrease in image quality. This can also be observed in Fig. 6. We provide more visual examples in the supplementary material.

*Comparison of CLIP score with State-of-the-art Methods.* Our ControlNet++ aims to improve the controllability of diffusion models using image-based conditions. Concerned about the potential adverse effects on text controllability, we evaluated various methods using CLIP-Score metrics across different datasets to measure the similarity between generated images and input text. As indicated in Tab. 3, ControlNet++ achieved comparable or superior CLIP-Score outcomes on several datasets relative to existing approaches. This suggests that our method not only markedly enhances conditional controllability but also preserves the original model’s text-to-image generation proficiency.






**Fig. 5:** Training DeepLabv3 (MobileNetv2) from scratch with different images, including ground truth images from ADE20K, and the generated images from ControlNet and ours. All the labels (i.e., segmentation masks) are ground truth labels in ADE20K. **Please note improvements here are non-trivial for semantic segmentation.**

*Effectiveness of Generated Images.* To further validate our improvements in controllability and their impact, we use the generated images along with real human-annotated labels to create a new dataset for training discriminative models from scratch. Please note that the only difference from the original dataset used to train the discriminative model is that we have replaced the images with those generated by the controllable diffusion model while keeping the labels unchanged. If the generative model exhibits good controllability, the quality of the constructed dataset will be higher, thereby enabling to train a stronger model.


Specifically, we conduct experiments on the ADE20K [64, 65] dataset on DeepLabv3 with MobileNetv2 backbone [4]. We use the standard training dataset (20210 training samples) to train the discriminative model and the validation dataset (5000 evaluation samples) for evaluation. We show the experimental results in Fig. 5, the segmentation model trained on our images outperforms the baseline results (ControlNet) by 1.19 mIoU. **Please note that this improvement is significant in segmentation tasks.** For instance, Mask2Former [7] improves previous SOTA MaskFormer [8] with around 1.1 mIoU in semantic segmentation. In addition to conducting experiments solely on the generated dataset, we also combined generated data with real data to train the segmentation model. The experimental results indicate that augmenting real ground truth data with data generated by ControlNet does not yield additional performance improvements (34.11 v.s. 34.08). In contrast, augmenting real data with our generated data results in significant performance enhancements (+1.76 mIoU).

*Qualitative Comparison.* Figs. 6 and 7 provide a qualitative comparison between our ControlNet++ and previous state-of-the-art methods across different conditional controls. When given the same input text prompts and image-based conditional controls, we observe that existing methods often generate areas



**Fig. 6:** Visualization comparison results in different conditional controls.


inconsistent with the image conditions. For instance, in the segmentation mask generation task, other methods often produce extraneous picture frames on the walls, resulting in a mismatch between the segmentation masks extracted from the generated image and the inputs. A similar situation occurs under depth conditions, where other methods fail to accurately represent the depth of different fingers. In contrast, images generated by ControlNet++ maintain good consistency with the input depth map.



**Fig. 7:** Comparison on Line-Art Edge.

### 4.3 Ablation Study

*Loss Settings.* In Fig. 8, we find that maintaining the original diffusion training process is crucial for preserving the quality and controllability of generated images. Relying solely on pixel-level consistency loss leads to severe image distortion, whereas training the model with both this loss and the diffusion training loss can enhance controllability without affecting image quality.



**Fig. 8:** Ablation study on different loss settings during training. Using only pixel-level consistency loss leads to severe image distortion, affecting both image quality and controllability. However, when combined with diffusion training loss, it is possible to gradually improve controllability without compromising image quality.

*Generalizability of Efficient Reward Fine-tuning.* Although the reward fine-tuning is used in a small subset of timesteps, it updates all the parameters of the ControlNet and therefore helps more timesteps to improve controllability during sampling. To prove this, we divide the sampling process into two parts: the unoptimized timesteps  $[T, t_{thre}]$ , and the optimized timesteps  $[t_{thre}, 1]$  and use ControlNet and our model for inference crossly, with 20-step sampling following ControlNet. Table 4 shows that our reward finetuning performed on a small number of timesteps  $[t_{thre}, 1]$  can be generalized to larger timesteps  $[T, t_{thre}]$ .

**Table 4:** The impact of efficient reward fine-tuning on different timesteps.

Unoptimized $[T, t_{thre}]$	Optimized $[t_{thre}, 1]$	ADE20K mIoU ( $\uparrow$ )
ControlNet	ControlNet	32.55
ControlNet	Ours	38.03
Ours	ControlNet	41.46
<b>Ours</b>	<b>Ours</b>	<b>43.64</b>

**Table 5:** Stronger reward model (UperNet-R50) leads to better controllability than the weaker reward model (DeepLabv3-MBv2).

Reward Model (RM)	RM mIoU $\uparrow$	Eval mIoU $\uparrow$
-	-	32.55
DeepLabv3-MBv2	34.02	31.96
FCN-R101	39.91	40.44
UperNet-R50	<b>42.05</b>	<b>43.64</b>

*Choice of Different Reward Models.* We demonstrate the effectiveness of different reward models in Tab. 5, all the evaluation results (i.e., Eval mIoU in the table) are evaluated by the most powerful segmentation model Mask2Former [7] with 56.01 mIoU, on ADE20K dataset. We experiment with three different reward models, including DeepLabv3 [4] with MobileNetv2 [44] backbone (DeepLabv3-MBv2), FCN [27] with ResNet-101 [14] backbone (FCN-R101) and UperNet [54] with ResNet-50 backbone. The results demonstrate that a more powerful reward model leads to better controllability for controllable diffusion models.

## 5 Discussion

How to make Hed/LineArt Edge extraction methods differentiable? The Hed and LineArt Edge extraction models are neural networks without non-differentiable operations. Differentiability can be achieved by simply modifying forward code. Some conditions (e.g., Box/Sketch/Pose) are not available. Our reward fine-tuning leverages a pre-trained ControlNet and a differentiable reward model. Currently, pre-trained ControlNet for object bounding boxes and differentiable reward models for sketches are lacking. In existing pose models, there are non-differentiable operations such as the NMS and keypoints grouping. We leave the question of how to extend consistency reward to more conditions to future work. Influence of Text Prompt. We discuss how different types of text prompts (No Prompt, Conflicting Prompt, and Perfect Prompt) affect the final results. As shown in Fig. 9, when the text prompt is empty or there is a semantic conflict with the image conditional control, ControlNet often struggles to generate accurate content. In contrast, our ControlNet++ manages to generate images that comply with the input conditional controls under various text prompt scenarios.



**Fig. 9:** When the input text prompt is empty or conflicts with the image-based conditional controls (the segmentation map in the top left corner), ControlNet struggles to generate correct content (red boxes), whereas our method manages to generate it well.

## 6 Conclusion

In this paper, we demonstrate from both quantitative and qualitative perspectives that existing works focusing on controllable generation still fail to achieve precise conditional control, leading to inconsistency between generated images and input conditions. To address this issue, we introduce ControlNet++, it explicitly optimizes the consistency between input conditions and generated images using a pre-trained discriminative reward model in a cycle consistency manner, which is different from existing methods that implicitly achieve controllability through latent diffusion denoising. We also propose a novel and efficient reward strategy that calculates consistency loss by adding noise to input images followed by single-step denoising, thus avoiding the significant computational and memory costs associated with sampling from random Gaussian noise. Experimental results under multiple conditional controls show that ControlNet++ significantly improves controllability without compromising image quality and image-text alignment, offering new insights into controllable visual generation.

## References

1. Black, K., Janner, M., Du, Y., Kostrikov, I., Levine, S.: Training diffusion models with reinforcement learning. arXiv preprint arXiv:2305.13301 (2023)
2. Brooks, T., Holynski, A., Efros, A.A.: Instructpix2pix: Learning to follow image editing instructions. In: CVPR (2023)
3. Caesar, H., Uijlings, J., Ferrari, V.: Coco-stuff: Thing and stuff classes in context. In: CVPR (2018)
4. Chen, L.C., Papandreou, G., Schroff, F., Adam, H.: Rethinking atrous convolution for semantic image segmentation. arXiv preprint arXiv:1706.05587 (2017)
5. Chen, M., Laina, I., Vedaldi, A.: Training-free layout control with cross-attention guidance. arXiv preprint arXiv:2304.03373 (2023)
6. Chen, T., Xu, B., Zhang, C., Guestrin, C.: Training deep nets with sublinear memory cost. arXiv (2016)
7. Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R.: Masked-attention mask transformer for universal image segmentation. In: CVPR (2022)
8. Cheng, B., Schwing, A., Kirillov, A.: Per-pixel classification is not all you need for semantic segmentation. NeurIPS (2021)
9. Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H.W., Sutton, C., Gehrmann, S., et al.: Palm: Scaling language modeling with pathways. arXiv preprint arXiv:2204.02311 (2022)
10. Clark, K., Vicol, P., Swersky, K., Fleet, D.J.: Directly fine-tuning diffusion models on differentiable rewards. arXiv preprint arXiv:2309.17400 (2023)
11. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. NeurIPS (2021)
12. Fan, Y., Watkins, O., Du, Y., Liu, H., Ryu, M., Boutilier, C., Abbeel, P., Ghavamzadeh, M., Lee, K., Lee, K.: Dpok: Reinforcement learning for fine-tuning text-to-image diffusion models. NeurIPS (2023)
13. Gal, R., Alaluf, Y., Atzmon, Y., Patashnik, O., Bermano, A.H., Chechik, G., Cohen-or, D.: An image is worth one word: Personalizing text-to-image generation using textual inversion. In: ICLR (2023)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
15. Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-or, D.: Prompt-to-prompt image editing with cross-attention control. In: ICLR (2023)
16. Hertzmann, A., Jacobs, C.E., Oliver, N., Curless, B., Salesin, D.H.: Image analogies. In: SIGGRAPH (2001)
17. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. NeurIPS (2020)
18. Ho, J., Salimans, T.: Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022)
19. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: LoRA: Low-rank adaptation of large language models. In: ICLR (2022)
20. Hu, M., Zheng, J., Liu, D., Zheng, C., Wang, C., Tao, D., Cham, T.J.: Cocktail: Mixing multi-modality controls for text-conditional image generation. NeurIPS (2023)
21. Huang, L., Chen, D., Liu, Y., Shen, Y., Zhao, D., Zhou, J.: Composer: Creative and controllable image synthesis with composable conditions. In: ICML (2015)
22. Ju, X., Zeng, A., Zhao, C., Wang, J., Zhang, L., Xu, Q.: Humansd: A native skeleton-guided diffusion model for human image generation. In: ICCV (2023)

23. Kawar, B., Zada, S., Lang, O., Tov, O., Chang, H., Dekel, T., Mosseri, I., Irani, M.: Imagic: Text-based real image editing with diffusion models. In: CVPR (2023)
24. Kingma, D., Salimans, T., Poole, B., Ho, J.: Variational diffusion models. NeurIPS (2021)
25. Kirstain, Y., Polyak, A., Singer, U., Matiana, S., Penna, J., Levy, O.: Pick-a-pic: An open dataset of user preferences for text-to-image generation. arXiv preprint arXiv:2305.01569 (2023)
26. Li, Y., Liu, H., Wu, Q., Mu, F., Yang, J., Gao, J., Li, C., Lee, Y.J.: Gligen: Open-set grounded text-to-image generation. In: CVPR (2023)
27. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR (2015)
28. Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: Sdedit: Guided image synthesis and editing with stochastic differential equations. In: ICLR (2022)
29. Mou, C., Wang, X., Xie, L., Zhang, J., Qi, Z., Shan, Y., Qie, X.: T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. arXiv preprint arXiv:2302.08453 (2023)
30. Nichol, A.Q., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In: ICML (2022)
31. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al.: Training language models to follow instructions with human feedback. NeurIPS (2022)
32. Parmar, G., Zhang, R., Zhu, J.Y.: On aliased resizing and surprising subtleties in gan evaluation. In: CVPR (2022)
33. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. arXiv preprint arXiv:2307.01952 (2023)
34. Prabhudesai, M., Goyal, A., Pathak, D., Fragkiadaki, K.: Aligning text-to-image diffusion models with reward backpropagation. arXiv preprint arXiv:2310.03739 (2023)
35. Qin, C., Zhang, S., Yu, N., Feng, Y., Yang, X., Zhou, Y., Wang, H., Niebles, J.C., Xiong, C., Savarese, S., et al.: Unicontrol: A unified diffusion model for controllable visual generation in the wild. NeurIPS (2023)
36. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML (2021)
37. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J.: Exploring the limits of transfer learning with a unified text-to-text transformer. JMLR (2020)
38. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125 (2022)
39. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: ICML (2021)
40. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022)
41. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI (2015)
42. Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dream-booth: Fine tuning text-to-image diffusion models for subject-driven generation. In: CVPR (2023)



43. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. *NeurIPS* (2022)
44. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: *CVPR* (2018)
45. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., Schramowski, P., Kundurthy, S., Crowson, K., Schmidt, L., Kaczmarczyk, R., Jitsev, J.: Laion-5b: An open large-scale dataset for training next generation image-text models. *ArXiv* (2022)
46. Schuhmann, C., Vencu, R., Beaumont, R., Kaczmarczyk, R., Mullis, C., Katta, A., Coombes, T., Jitsev, J., Komatsuzaki, A.: Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *ArXiv* (2021)
47. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: *ICML* (2015)
48. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: *ICLR* (2021)
49. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: *ICLR* (2021)
50. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al.: Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023)
51. Wang, X., Darrell, T., Rambhatla, S.S., Girdhar, R., Misra, I.: Instancediffusion: Instance-level control for image generation (2024)
52. Wu, X., Hao, Y., Sun, K., Chen, Y., Zhu, F., Zhao, R., Li, H.: Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *arXiv preprint arXiv:2306.09341* (2023)
53. Wu, X., Sun, K., Zhu, F., Zhao, R., Li, H.: Better aligning text-to-image models with human preference. In: *ICCV* (2023)
54. Xiao, T., Liu, Y., Zhou, B., Jiang, Y., Sun, J.: Unified perceptual parsing for scene understanding. In: *ECCV* (2018)
55. Xie, J., Li, Y., Huang, Y., Liu, H., Zhang, W., Zheng, Y., Shou, M.Z.: Boxdiff: Text-to-image synthesis with training-free box-constrained diffusion. In: *ICCV* (2023)
56. Xie, S., Tu, Z.: Holistically-nested edge detection. In: *ICCV* (2015)
57. Xu, J., Liu, X., Wu, Y., Tong, Y., Li, Q., Ding, M., Tang, J., Dong, Y.: Imagereward: Learning and evaluating human preferences for text-to-image generation. *NeurIPS* (2023)
58. Yang, Z., Wang, J., Gan, Z., Li, L., Lin, K., Wu, C., Duan, N., Liu, Z., Liu, C., Zeng, M., et al.: Reco: Region-controlled text-to-image generation. In: *CVPR* (2023)
59. Ye, H., Zhang, J., Liu, S., Han, X., Yang, W.: Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721* (2023)
60. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: *ICCV* (2023)
61. Zhang, T., Zhang, Y., Vineet, V., Joshi, N., Wang, X.: Controllable text-to-image generation with gpt-4. *arXiv preprint arXiv:2305.18583* (2023)
62. Zhao, S., Chen, D., Chen, Y.C., Bao, J., Hao, S., Yuan, L., Wong, K.Y.K.: Uni-controlnet: All-in-one control to text-to-image diffusion models. *NeurIPS* (2023)
63. Zhao, W., Bai, L., Rao, Y., Zhou, J., Lu, J.: Unipc: A unified predictor-corrector framework for fast sampling of diffusion models. *arXiv preprint arXiv:2302.04867* (2023)

- 64. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: CVPR (2017)
- 65. Zhou, B., Zhao, H., Puig, X., Xiao, T., Fidler, S., Barriuso, A., Torralba, A.: Semantic understanding of scenes through the ade20k dataset. IJCV (2019)
- 66. Zhou, D., Li, Y., Ma, F., Yang, Z., Yang, Y.: Migc: Multi-instance generation controller for text-to-image synthesis. In: CVPR (2024)
- 67. Zhu, D., Chen, J., Shen, X., Li, X., Elhoseiny, M.: Minigpt-4: Enhancing vision-language understanding with advanced large language models. arXiv preprint arXiv:2304.10592 (2023)
- 68. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV (2017)