

# AnyControl: Create Your Artwork with Versatile Control on Text-to-Image Generation

Yanan Sun<sup>1</sup>, Yanchen Liu<sup>1,2</sup>, Yinhao Tang<sup>1</sup>, Wenjie Pei<sup>2</sup>, and Kai Chen<sup>1\*</sup>

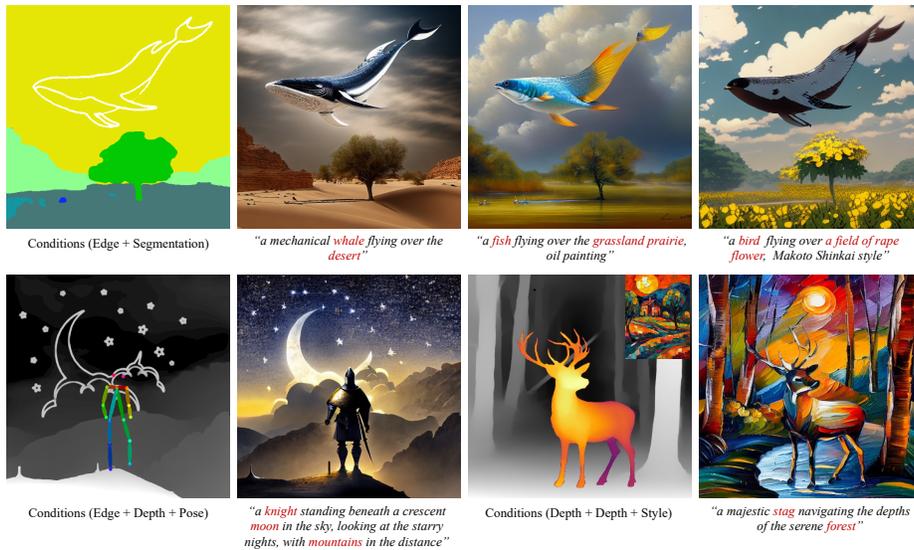
<sup>1</sup> Shanghai AI Laboratory

<sup>2</sup> Harbin Institute of Technology, Shenzhen

now.syn@gmail.com

wenjiecoder@outlook.com

{liuyanchen, tangyinhao, chen kai}@pjlab.org.cn



**Fig. 1:** Multi-control image synthesis of **AnyControl**. Our model supports free combinations of multiple control signals and generates harmonious results that are well-aligned with each input. The input control signals fed into the model are shown in a combined image for better visualization.

**Abstract.** The field of text-to-image (T2I) generation has made significant progress in recent years, largely driven by advancements in diffusion models. Linguistic control enables effective content creation, but struggles with fine-grained control over image generation. This challenge has been explored, to a great extent, by incorporating additional user-supplied spatial conditions, such as depth maps and edge maps, into pre-trained T2I models through extra encoding. However, multi-control image synthesis still faces several challenges. Specifically, current approaches are limited in handling free combinations of diverse input control signals, overlook the complex relationships among multiple spatial conditions, and often fail to maintain semantic alignment with provided

\* Corresponding author.

textual prompts. This can lead to suboptimal user experiences. To address these challenges, we propose **AnyControl**, a multi-control image synthesis framework that supports arbitrary combinations of diverse control signals. AnyControl develops a novel Multi-Control Encoder that extracts a unified multi-modal embedding to guide the generation process. This approach enables a holistic understanding of user inputs, and produces high-quality, faithful results under versatile control signals, as demonstrated by extensive quantitative and qualitative evaluations. Our project page is available in <https://any-control.github.io>.

**Keywords:** Controllable Image Synthesis · Multi-Control · Text-to-image

## 1 Introduction

In recent years, the field of text-to-image (T2I) generation has experienced significant advancements, leading to unprecedented improvements in generated image quality and diversity, primarily attributed to the introduction of diffusion models [13, 14, 33, 38]. While linguistic control allows for effective and engaging content creation, it also presents challenges in achieving fine-grained control over image generation. This challenge is explored in [25, 44], where an additional network is employed to encode and inject the user-supplied control signal into the pre-trained T2I model such as Stable Diffusion [33], so that to exert influence over the image generation process. Built upon [44], subsequent approaches [15, 28, 45] presents unified architecture designs for managing multiple spatial conditions.

However, the task of **multi-control image synthesis** remains challenging in the following aspects: (1) accommodating free combinations of input conditions, (2) modeling complex relationships among multiple spatial conditions, and (3) maintaining compatibility with textual prompts. We refer to these challenges as input flexibility, spatial compatibility, and textual compatibility, respectively.

**Input flexibility.** The first challenge comes from the any combination of available control signals based on user requirements. The amount and modality of control signals provided by users are varying, placing high demands on the input flexibility of the model. However, existing methods [15, 45] typically employ fixed-length input channels, limiting their ability to accommodate diverse inputs. Other approaches [25, 28, 44] adopt MoE design to solve varying-number conditions, which can result in unforeseen artifacts when processing unseen combinations of inputs.

**Spatial compatibility.** Secondly, control signals are not isolated; instead, they collectively influence the composition of a complete image. It is crucial to consider the relationships among these control signals, especially when managing occlusions among multiple spatial conditions. Unfortunately, current algorithms commonly combine multiple conditions through weighted summation with hand-crafted weights, easily leading to undesired blending results, or even causing low-response control signals to disappear when addressing occlusions.

**Textual compatibility.** Ultimately, textual compatibility emerges as an important factor influencing user experience. Typically, the textual descriptions govern the content of images, whereas spatial conditions compensate the structural information. Nevertheless, a lack of communication between the textual

and spatial conditions often leads current algorithms to prioritize accommodating the spatial conditions, thereby disregarding the impact of textual prompts.

In summary, generating comprehensive and harmonious results that satisfy both textual prompts and multiple spatial conditions presents a significant challenge for multi-control image synthesis. To tackle the challenges of input flexibility, spatial compatibility, and textual compatibility, we propose **AnyControl**, a controllable image synthesis framework that supports arbitrary combinations of diverse control signals.

At the core of AnyControl is the **Multi-Control Encoder**, which plays a crucial role in ensuring coherent, spatially and semantically aligned multi-modal embeddings. This novel component allows AnyControl to extract a unified representation from various control signals, enabling a truly versatile and high-performing multi-control image synthesis framework.

Specifically, Multi-Control Encoder is driven by **multi-control fusion** block and **multi-control alignment** block in turns, with a set of **query tokens** to unite the two seamlessly.

Multi-control fusion block is employed to aggregate compatible information from multiple spatial conditions through the query tokens. A cross-attention transformer block is employed on the query tokens and the visual tokens of spatial conditions extracted from a pre-trained visual encoder. Therefore, the rich spatial controllable information is passed to the query tokens, which will be utilized in the multi-control alignment block.

Multi-control alignment is used to guarantee the compatibility between all forms of control signals by aligning all other signals to the textual signal. A self-attention transformer block is employed on the query tokens and textual tokens. The query tokens contain spatial controllable information, while the textual tokens carry semantic information. Through information exchange between the query and textual tokens, both types of tokens are able to represent compatible multi-modal information.

With alternating multi-control fusion and alignment blocks in several turns, the query tokens achieve a comprehensive understanding with highly aligned and compatible information from versatile user inputs. This capability empowers our method to handle complex relationships among conditions and uphold strong compatibility with textual prompts. Consequently, this approach fosters a more smooth and harmonious control over the generated images. Furthermore, transformer blocks with attention mechanisms inherently excel in accommodating a variety of control signals, and thus enable free combinations of user inputs.

In summary, our contributions are manifold in the following:

1. AnyControl proposes a novel Multi-Control Encoder comprising a sequence of alternative multi-control fusion and alignment blocks to achieve comprehensive understanding of complex multi-modal user inputs.
2. AnyControl supports flexible combinations of user inputs, regardless of the amount and modality of different control signals.
3. AnyControl produces more harmonious and natural high-quality outcomes, demonstrating state-of-the-art performance in multi-control image synthesis.

## 2 Related Work

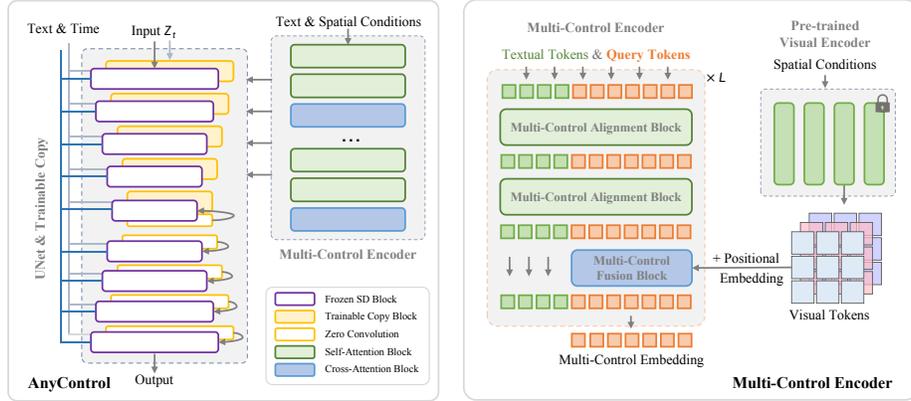
### 2.1 Text-to-image Generation

T2I diffusion models [26, 31, 33, 36] have emerged as a promising approach for generating high-quality images from textual prompts. Diffusion models [13, 38], originally developed for image generation, have been adapted to T2I domain, offering a novel perspective on the problem. These models leverage the concept of iterative denoising, where the generation process unfolds step-by-step, progressively refining the image quality. The diffusion process allows for better control over the generated images by conditioning on both the text input and intermediate representations at each diffusion step. Recent advances in T2I diffusion models have explored various techniques to enhance the generation process, such as introducing attention mechanisms to better align textual and visual features [33] and operating in latent space [31, 33] to achieve complexity reduction and detail preservation. While T2I diffusion models have shown promising results, there is still ongoing research to address challenges such as controllability, also the focus of this paper, in the context of diffusion-based T2I generation.

### 2.2 Controllable Image Synthesis

Text descriptions guide the diffusion model to generate user-desired images but are insufficient in fine-grained control over the generated results. The fine-grained control signals are diverse in modality, for instance, layout constraint is introduced to arrange the location of the objects given; a bundle of works [6, 17, 22, 29, 40–42, 46] thoroughly explore to synthesis images with high layout alignment given the semantic-aware boxes. Besides, segmentation map [1, 2, 7, 8] is another popular control signal for controlling the layout and object shape of generated images. InstanceFusion [39] proposes a method to support location control in more free-form such as point, scribble, box and segmentation map. Highly detailed control can be achieved by structure signal such as sketch [3]. Depth map [20] can provide the control of the depth of field for the generated images. As layout, structure and depth control all outline the generated image in spatial alignment, content control [9, 35, 43] enables the personalization of the generated appearance in semantic level through an additional image input.

Studies [24, 25, 44] propose general framework designs to process diverse spatial conditions instead of control-specific design, while both spatial and content control are jointly taken into consideration in works [16, 18]. Specifically, considering the powerful generation ability of T2I model, ControlNet [44] proposes to utilize the trainable copy of the UNet encoder in the T2I diffusion model to encode extra condition signals into latent representations and then apply zero convolution to inject into the backbone of the UNet in diffusion modal. The simple but effective design shows generalized and stable performance in spatial control, and thus are widely adopted in various downstream applications. However, ControlNet is a single-modality framework and requires separate model for each modality. To address this, unified ControlNet-like models [15, 28, 45] are proposed to handle diverse control signals with only one multi-modality



**Fig. 2: AnyControl and Multi-Control Encoder.** Left shows the overall framework of nyControl, which develops a Multi-Control Encoder for extracting comprehensive multi-control embeddings based on the textual prompts and multiple spatial conditions. The multi-control embeddings are then utilized to guide the generation process. Right shows the detailed design of Multi-Control Encoder driven by alternating multi-control fusion and alignment blocks, with query tokens defined to distill the compatible information from textual tokens and visual tokens of the spatial conditions.

model. Another advantage of these methods is that they can support multi-control image synthesis. They adopt fixed length input channels or MoE design with hand-crafted weighted summation to aggregate conditions. Nevertheless, these methods are short in handling conditions with complex relations and hard to generate harmonious, natural results under various control signals.

### 3 Method

In this section, we first give a preliminary overview of Stable Diffusion [33] and ControlNet [44]. Subsequently, we introduce AnyControl, featuring a pioneering Multi-Control Encoder crafted for extracting a unified representation with compatible information for multiple control signals. Finally, we expound on our training dataset and strategy. Figure 2 depicts the architecture of AnyControl and the Multi-Control Encoder.

#### 3.1 Preliminary

**Stable Diffusion.** T2I generation introduces text as conditions in diffusion models. In the forward pass, Gaussian noises are gradually added to the sample over a series of steps; while the backward process learns to recover the image by estimating and eliminating the noise with the text guidance. In this paper, we base Stable Diffusion [33], one of the most popular T2I diffusion model, to develop AnyControl for multi-control image synthesis.

Stable Diffusion model operates the diffusion and denoising process in latent space rather than pixels to reduce computation cost. It adopts UNet-like [34] structure as its backbone, comprising downsampling blocks, middle block and

upsampling blocks. The text guidance are encoded through CLIP [30] text encoder and integrated into the UNet through a *CrossAttention* block after each ResBlock [10]. If we use  $Z$  to denote the noise features derived from the last ResBlock and  $Y$  to denote the embeddings encoded by the text encoder, the output noise features  $\tilde{Z}$  from *CrossAttention* block can be obtained by

$$Q = W_q(Z), K = W_k(Y), V = W_v(Y), \quad (1)$$

$$\tilde{Z} = \text{Softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V, \quad (2)$$

where  $W_q$ ,  $W_k$  and  $W_v$  are projection layer and  $d$  is the dimension of the embedding space.

**ControlNet.** ControlNet [44] is developed to adapt Stable Diffusion model for spatial conditions. To be specific, it locks the parameters of Stable Diffusion, and makes a trainable copy of the encoding layers in the UNet. The two parts are connected by zero convolution layers with zero-initialized weights to progressively increase spatial control influence as the training goes. This design empowers ControlNet to achieve robust controllable image generation while preserve the quality and capabilities of Stable Diffusion model.

### 3.2 AnyControl

**Multi-Control Encoder.** Similar to ControlNet, in AnyControl, we also lock the pre-trained stable diffusion model, and instead design a Multi-Control Encoder for understanding complex control signals. We first obtain three types of tokens, i.e., textual tokens  $\mathcal{T}$ , visual tokens  $\mathcal{V}$  and query tokens  $\mathcal{Q}$ . Textual tokens are extracted from CLIP text encoder on textual prompts, while visual tokens are obtained from a pre-trained visual encoder (e.g., CLIP image encoder) on all of the user-provided spatial conditions in image form. The query tokens are defined as a set of learnable parameters. To address the three challenges discussed in the introduction, i.e., input flexibility, spatial compatibility and textual compatibility, we develop the multi-control encoder via alternating multi-control fusion blocks and multi-control alignment blocks united by the query tokens.

**Multi-Control Fusion.** Multi-control fusion block aims to extract compatible information from various spatial conditions. This is accomplished by utilizing a cross-attention transformer block to facilitate interactions between the query tokens and the visual tokens of all spatial conditions.

Specifically, suppose that there are  $n$  spatial conditions in image form of various modalities including depth, segmentation, etc. We can obtain the visual tokens  $\mathcal{V}_{i,j}$  for condition  $C_i$  from the  $j$ -th block of the pre-trained visual encoder. Here, we use  $[\mathcal{V}_{1,j}, \mathcal{V}_{2,j}, \dots, \mathcal{V}_{n,j}]$  to represent the visual tokens for all the spatial conditions from the  $j$ -th block. Then the interactions in the multi-control fusion block can be formulated as

$$\mathcal{Q}_j = \text{CrossAttention}(\mathcal{Q}_j, [\mathcal{V}_{1,j} + P, \mathcal{V}_{2,j} + P, \dots, \mathcal{V}_{n,j} + P]), \quad (3)$$

where  $P$  denotes a shared learnable positional embedding additive to each  $\mathcal{V}_{i,j}$  for better alignment between the query tokens and the visual tokens.

After this process, the spatial controllable information encoded in the visual tokens is passed on to the query tokens.

**Multi-Control Alignment.** Although the various controllable information is integrated into the query tokens, it is challenging to infer the priority of spatial control signals within the overlapping region due to the absence of a global condition that indicates the relationships among spatial conditions. Fortunately, textual prompts can serve as a global control that regulates the content of generated image. Therefore, in the multi-control alignment block, we facilitate the interactions between the query tokens and the textual tokens with a self-attention transformer block. Before we encode the textual prompts to tokens, we append a textual task prompt at the tail of the user-provided text to solve the modality discrepancy among diverse spatial conditions. Then we concatenate the query tokens  $\mathcal{Q}$  and textual tokens  $\mathcal{T}$  together and perform the self-attention as

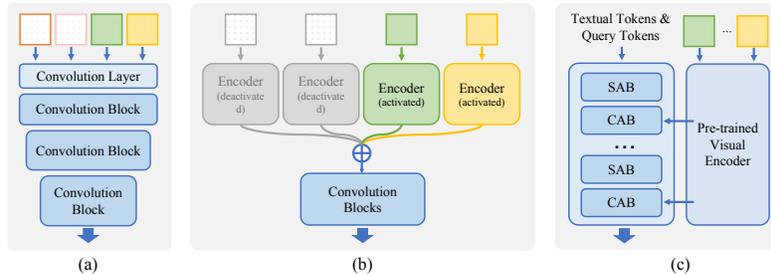
$$[\mathcal{Q}_{j+1}, \mathcal{T}_{j+1}] = \text{SelfAttention}([\mathcal{Q}_j, \mathcal{T}_j]). \quad (4)$$

With self-attention, the query tokens, which carry the mixed controllable information, will exchange information with textual tokens and thus can achieve semantic alignment with user prompts.

**Alternating Fusion and Alignment.** To ensure the information aligned and compatible of all the control signals, we employ the multi-control fusion and alignment blocks alternately for multiple turns. Notably, we utilize multi-level visual tokens for fine-grained spatial control. Specifically, in each turn, the visual tokens consumed in the cross-attention transformer block are extracted from different levels of the pre-trained visual encoder, considering the spatial conditions are diverse in controlling level, *i.e.*, layout control such as segmentation map and structural control such as edge map. Therefore, multi-level visual tokens are necessary for multi-control fusion blocks in different depth.

**Advantages of AnyControl.** The query tokens work as a bridge, uniting the two types of blocks seamlessly. After several turns, the query tokens retain well-aligned compositional information, served as a unified multi-modal representations for user inputs. This design empowers AnyControl in multi-control image synthesis even with occlusion, generating high-quality harmonious results with high spatial and textual compatibility. Our multi-control encoder shares a similar idea to Q-Former [21], however, AnyControl incorporates many dedicated design for the multi-control image synthesis such as the appended textual task prompt, additional shared position embeddings across all the conditions and the usage of multi-level visual tokens. In implementation, to save computation cost, we insert the cross-attention block after every two self-attention blocks.

Another natural advantage of our AnyControl lies in the input flexibility. AnyControl, utilizing the transformer blocks with attention mechanism, has natural advantage in accommodating free combinations of user inputs. Previous



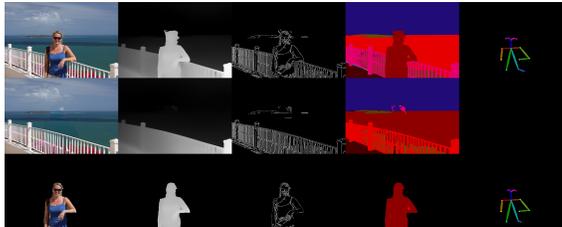
**Fig. 3: Three types of multi-control methods.** Square in different color denotes different condition type while dotted square denotes zero tensor. **(a)** Some methods [15, 45] adopt fixed-length channels of the input convolution layer followed by convolution blocks to serve as the Multi-Control Encoder. **(b)** Other methods [25, 28, 44] utilize MoE design, that is, construct separate encoder for each type of control signal and then obtain the embeddings through weighted sum. **(c)** Different from them, AnyControl adopts attention mechanism to accommodate varying-number and varying-modality of conditions. “SAB” and “CAB” denotes self- and cross-attention block, respectively.

methods either adopt the design of fixed-length input channels or MoE structure as illustrated in Figure 3. The former limits the freedom of user inputs, while the latter, MoE design, supports combining flexible inputs with hand-crafted weights, leading to laborious adjustments to the combination weights.

### 3.3 Training

#### Datasets.

We adopt the training dataset, MultiGen, for multi-control image synthesis presented in [28]. This dataset is built from LAION [37] with aesthetics score above 6. Low-resolution images are removed and finally 2.8M images are kept. Different methods are utilized to extract the control signals. Unfortunately, there is a domain gap between



**Fig. 4: Visualization of aligned and unaligned conditions.** The first row shows the aligned case where pixels at the same location of all the control signals describe the same object. Conditions in the second and third rows describe the foreground and background respectively, contributing to a complete image together, constructing the unaligned case.

the combinations of the spatial conditions during training and inference time, *i.e.*, during training, all the spatial conditions extracted from the same image are fully aligned while the multiple spatial conditions accepted from users are not the case. User-provided conditions usually have multiple image sources, thus the extracted spatial conditions are not always aligned and sometimes have occlusion in the overlapping region, which requires the model to handle the spatial conditions in right arrangement according to the depth of the target scene. To relieve the discrepancy, we collect a subset of unaligned data as shown in Figure 4. To be specific, we utilize the images in Open Images dataset [19] and MSCOCO [23]

dataset which are rich in objects to make the synthetic data. Given an image and the mask of a foreground object, we recover the background image with the masked region using the inpainting tool [47]. We discard images with too small or too large objects, and finally produce 0.44M images as supplementary unaligned training data.

**Training Strategy.** When utilizing the unaligned data for training, we take the combination of spatial conditions for the foreground object and the inpainted background image together while treating the original image as target. During training, for the data with fully aligned spatial conditions, we randomly pick two conditions for each training sample; for the synthetic unaligned data, we randomly pick a condition for the foreground object and the background inpainted image respectively. We randomly drop all the conditions at a rate of 0.05 to enable classifier free guidance, and also randomly drop the textual prompts at a rate of 0.05 to let the model learn from pure spatial conditions only.

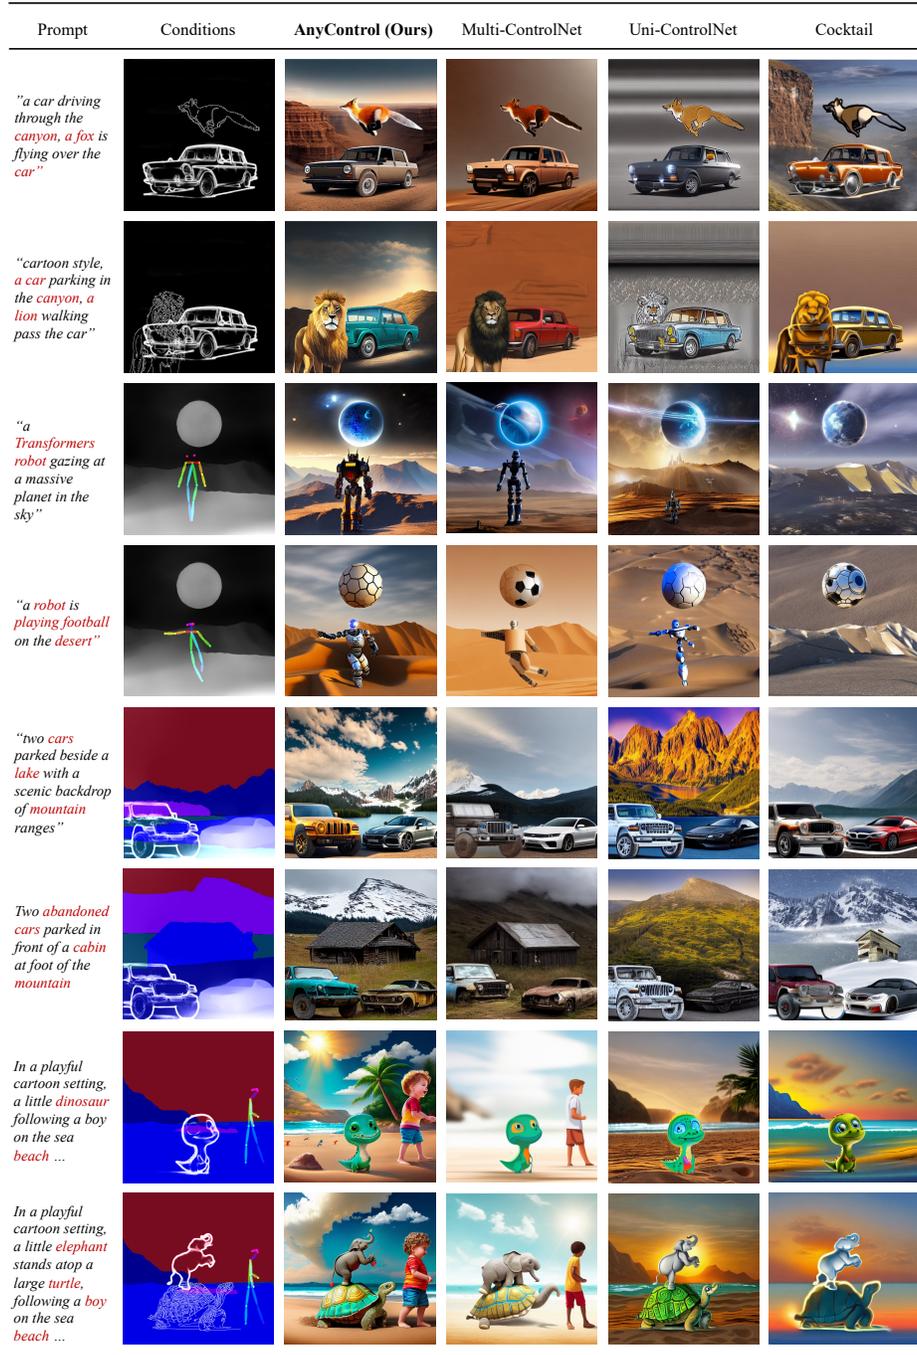
## 4 Experiments

We validate the effectiveness of AnyControl with Stable Diffusion [33] of version 1.5 on four types of conditions, including Edge [4], Depth [32], Segmentation [27], and Human Pose [5]. We compare our AnyControl with state-of-the-art methods including Multi-ControlNet and Multi-Adapter, the versions of ControlNet [44] and T2I-Adapter [25] which support multiple spatial conditions, as well as Uni-ControlNet [45], Cocktail [15], UniControl [28], DiffBlender [18] and CnC [20] with extensive qualitative and quantitative results. Implementation details including network structure, hyper-parameters of training and inference can be found in the supplementary material.

### 4.1 Qualitative Results

In this section, we analyze input flexibility, spatial compatibility, and also the compatibility with text, style and color control.

**Input Flexibility.** There are three ways to process free combinations of spatial conditions from users: 1) MoE design (*i.e.*, UniControl [28], Multi-ControlNet [44]); 2) Attention design (AnyControl); 3) Composition design, which merges spatial conditions of the same type into one image so that methods with fixed-length input channels can work smoothly. In MoE-based methods, the composition of different conditions is achieved by hand-crafted weighted summation. Instead, our AnyControl adopts attention mechanism to learn the composition weights dynamically, achieving superior performance on multi-control image synthesis as shown in Figure 5 and Figure 6. In addition, “sticker” artifact is observed in the results of Cocktail [15] with composition design.



**Fig. 5: Comparison on multi-control image synthesis.** Multi-ControlNet adopts MoE design to process diverse conditions while Cocktail adopts the composition design by combining multiple conditions of the same type into one.



Fig. 6: Varied-number and varied-type of input spatial conditions.

**Spatial Compatibility.** In Figure 5, we provide comparisons given various conditions with occlusion, demonstrating the superiority of AnyControl in handling complex multi-control synthesis. Blending issues are difficult to avoid in previous methods [15, 44, 45] with trivial design on multi-control combination.

Figure 7 further shows examples on AnyControl in dealing with relative spatial positions of two conditions, which generates high-quality results with correct spatial relation between conditions. For example, given different layouts, the generated cat and teddy bear from AnyControl are always positioned in a natural way, rather than mixed within the occluded region. Another notable thing is that AnyControl shows a surprising ability of dealing with the interaction of generated objects and the corresponding environment. When placed at the same horizontal line, the cat and teddy bear are seated at the same plane, while the teddy bear sitting on a stage when its vertical position axis is raised up. These advantages contribute to the introduction of Multi-Control Encoder, which strengthens the interactions between all control signals and consequently achieves a comprehensive understanding of complex inputs.

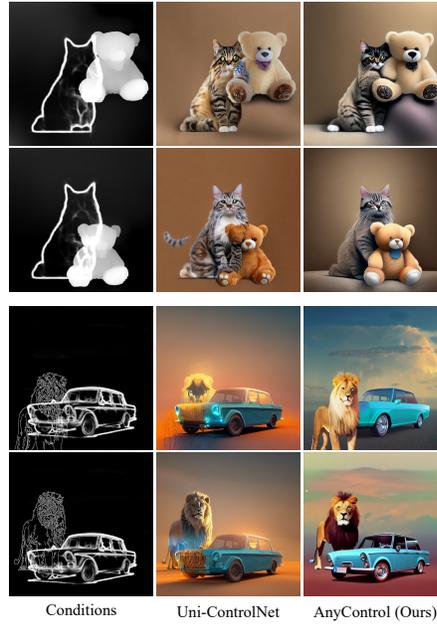


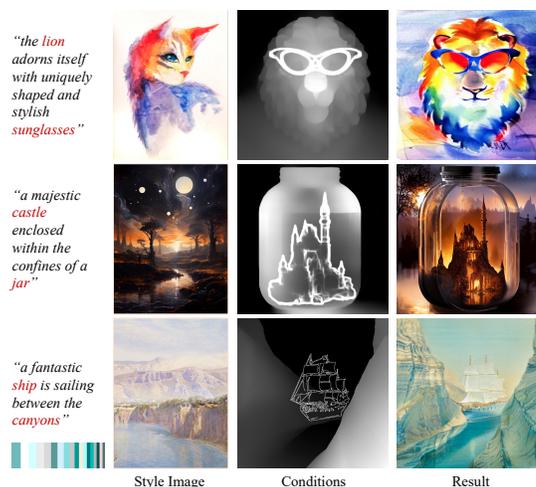
Fig. 7: **Spatial compatibility.** AnyControl is capable of inferring the relationships not only between conditions but also between the generated objects and the environment.

**Text Compatibility.** Textual prompts play an important role in controllable image synthesis as text is typically the primary communication means between

human and T2I models, while spatial conditions work as auxiliary roles in providing the fine-grained information. Therefore, while guided by spatial conditions, maintaining the compatibility with textual prompts is essential. However, existing methods commonly prioritize the response to spatial conditions and miss important message in textual prompts. For instance, in the third row of Figure 5, other methods either totally neglect the “Transformers Robot” information so that fail to bind the concept with the input human pose, or only partially respond to “Robot” but drop “Transformers” information. On the contrary, our method responds to all important information in multi-modality user inputs and produces harmonious results.

**Compatibility with Style and Color Control.** As a plug-and-play model,

AnyControl can be integrated with existing conditional generation methods in a convenient way. We take style and color controls as examples to demonstrate the effectiveness of AnyControl in the collaboration with other plug-and-play modules for wider applications. Specifically, we enhance AnyControl with decoupled cross-attention [43] in UNet to employ style and color controls. The compositional outcomes are visually depicted in Figure 8, revealing the generation of high-quality results that adhere to style, color and spatial constraints.



**Fig. 8: AnyControl with style and color controls.** The first two cases take style, depth and edge controls, while the last further takes color control.

## 4.2 Multi-Control Benchmark: COCO-UM

Most existing methods evaluate multi-control image synthesis on MSCOCO validation set with totally spatio-aligned conditions extracted from different algorithms. However, we argue that evaluation on the dataset with well-aligned multi-control conditions cannot reflect the ability of methods to handle occluded multiple conditions in practical applications, given that the user provided conditions are typically collected from diverse sources which are not aligned.

Therefore, we construct an **Unaligned Multi-control** benchmark based on MSCOCO validation set, short for **COCO-UM**, for a more effective evaluation on multi-control image synthesis. The construction pipeline is similar to that used in unaligned data synthesis described in Section 3.3. That is, we decompose an image into the background image and foreground image outlined by an object mask, and then recover the background image through inpainting tools [47].

**Table 1: Comparisons of multi-control image synthesis on COCO-UM with existing multi-control methods.** Results with the best and the second best performance are highlighted in **red** and **blue**, respectively.

Methods	FID↓	CLIP↑	Depth (RMSE↓)	Seg. (mPA↑)	Edge (RMSE↓)	Pose (mAP↑)
Multi-ControlNet [44]	55.95	24.80	<b>17.81</b>	<b>42.78</b>	47.35	15.69
Multi-Adapter [25]	51.67	<b>25.73</b>	24.08	32.62	<b>44.86</b>	6.28
Uni-ControlNet [45]	55.28	24.48	20.57	41.10	<b>43.28</b>	<b>18.40</b>
UniControl [28]	57.46	24.72	23.69	34.58	45.48	4.06
Cocktail [15]	<b>47.39</b>	25.33	-	31.74	-	12.16
AnyControl (Ours)	<b>44.28</b>	<b>26.41</b>	<b>18.00</b>	<b>43.34</b>	45.25	<b>18.81</b>

**Table 2: Comparisons of single control with existing single- and multi-control methods on COCO-5K.** Results with the best and the second best performance are highlighted in **red** and **blue**, respectively.

Methods	Depth			Segmentation			Edge		Pose	
	FID ↓	CLIP ↑	RMSE ↓	FID ↓	CLIP ↑	mPA ↑	FID ↓	CLIP ↑	FID ↓	CLIP ↑
<b>Single-Control Methods</b>										
ControlNet [44]	<b>19.80</b>	25.30	<b>13.86</b>	20.39	25.46	45.68	<b>16.16</b>	25.34	26.15	25.39
T2I-Adapter [25]	20.08	25.67	15.62	20.95	24.91	33.67	18.76	25.25	<b>24.94</b>	<b>25.76</b>
<b>Multi-Control Methods</b>										
Uni-ControlNet [45]	20.09	25.25	15.93	22.96	25.11	35.56	17.51	25.18	26.61	24.86
UniControl [28]	20.67	25.51	<b>14.07</b>	<b>19.73</b>	<b>26.11</b>	<b>45.77</b>	16.69	25.15	27.90	25.27
Cocktail [15]	-	-	-	26.02	25.28	33.10	<b>15.70</b>	24.91	28.36	25.27
DiffBlender [18]	21.32	25.59	18.07	-	-	-	21.48	<b>25.57</b>	30.33	25.30
CnC [20]	20.42	<b>26.95</b>	17.57	-	-	-	-	-	-	-
AnyControl (Ours)	<b>18.04</b>	<b>25.98</b>	15.48	<b>18.89</b>	<b>26.02</b>	<b>48.73</b>	18.89	<b>25.88</b>	<b>24.12</b>	<b>25.99</b>

Additionally, after obtaining the recovered background image, we remove the bad cases with low image quality, such as cases with new generated object within the hole region rather than filled by background scene. Finally, we construct an occluded multi-control dataset with 1480 samples.

### 4.3 Quantitative Evaluation

For thorough quantitative evaluation, we employ various evaluation metrics, including FID [12] for the image quality, CLIP-Score [11] for the alignment with textual prompt. For the condition fidelity, we adopt RMSE for depth map and edge map, mPA and mAP for segmentation map and pose map, respectively.

**Multi-Control Synthesis Evaluation.** We evaluate multi-control synthesis on COCO-UM. As depicted in Table 1, our method outperforms other multi-control methods on FID and CLIP-Score by a large margin. This remarkable achievement signifies that AnyControl is capable of processing complex combinations of multiple spatial conditions, and generates high-quality harmonious results well aligned with the textual prompts and spatial conditions.

**Single-Control Synthesis Evaluation.** For comprehensive evaluation, we also conduct the comparisons on each single condition as tabulated in Table 2. To be specific, we evaluate the single control synthesis on the full validation set of MSCOCO, that is, COCO-5K since there is no occlusion in single control scenario. Overall, AnyControl outperforms existing single- and multi-control methods on most metrics, illustrating the superiority to existing methods.

**Ablation Study on Unaligned Data.** Unaligned data is provided to solve the gap between the alignment of input conditions during training and inference. That is, in training, conditions are totally aligned; while, in testing, the control signals from users contributing to a whole image are almost not aligned at all. In Table 3, we provide comparisons on FID and CLIP-Score of AnyControl trained with and without unaligned data. As illustrated, a large improvement on FID and CLIP-Score is observed through the data expansion on occluded cases. The introduction of unaligned data during training strengthens AnyControl in modeling complex multi-control synthesis, especially the occluded cases.

**Table 3: Ablation study of training with and without unaligned data on COCO-UM.**

Metric	FID↓	CLIP↑
w/o unaligned data	52.10	25.62
w unaligned data	44.28	26.40

#### 4.4 Discussion

Although the input number of spatial conditions is not limited in AnyControl, we observe the miss-blending issue as shown in Figure 9, when the number of spatial conditions is overlarge, such as 8 in this case. The possible reasons are as follows: 1) the limited ability of CLIP text encoder in understanding complex textual prompts with numerous concepts; 2) Too many visual tokens in cross-attention transformer block results in a decrease in the accuracy of softmax, and thus weaken AnyControl in precise multi-control understanding. We leave this issue as future work.



**Fig. 9: Miss-blending issue under too many spatial conditions.**

## 5 Conclusion

In conclusion, we propose a multi-control image synthesis framework based on the public T2I model to address the limitations of existing methods in accommodating diverse inputs, handling relationships among spatial conditions, and maintaining the compatibility with textual prompts. AnyControl supports free combination of versatile control signals, which develops a Multi-Control Encoder that enables holistic understanding of multi-modal user inputs. We achieve this through employing alternating multi-control fusion and alignment blocks united by a set of query tokens. This approach enables AnyControl model complex relationships among diverse control signals and extract a multi-control embedding with compatible information. Our method produces high-quality natural outcomes, positioning it as a state-of-the-art solution for multi-condition image generation. The advancements introduced by AnyControl contribute to the broader goal of enhancing controllable image synthesis and pushing the boundaries of T2I generation.

**Acknowledgement.** This project is supported by the National Key R&D Program of China (No. 2022ZD0161600)

## References

1. Avrahami, O., Hayes, T., Gafni, O., Gupta, S., Taigman, Y., Parikh, D., Lischinski, D., Fried, O., Yin, X.: Spatext: Spatio-textual representation for controllable image generation. In: CVPR (2023) [4](#)
2. Bar-Tal, O., Yariv, L., Lipman, Y., Dekel, T.: Multidiffusion: Fusing diffusion paths for controlled image generation (2023) [4](#)
3. Bashkirova, D., Lezama, J., Sohn, K., Saenko, K., Essa, I.: Masksketch: Unpaired structure-guided masked image generation. In: CVPR (2023) [4](#)
4. Canny, J.: A computational approach to edge detection. TPAMI **PAMI-8**(6), 679–698 (1986) [9](#)
5. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. In: CVPR (2017) [9](#)
6. Cheng, J., Liang, X., Shi, X., He, T., Xiao, T., Li, M.: Layoutdiffuse: Adapting foundational diffusion models for layout-to-image generation. arXiv preprint arXiv:2302.08908 (2023) [4](#)
7. Couairon, G., Careil, M., Cord, M., Lathuilière, S., Verbeek, J.: Zero-shot spatial layout conditioning for text-to-image diffusion models. In: ICCV. pp. 2174–2183 (2023) [4](#)
8. Gafni, O., Polyak, A., Ashual, O., Sheynin, S., Parikh, D., Taigman, Y.: Make-a-scene: Scene-based text-to-image generation with human priors. In: ECCV (2022) [4](#)
9. 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. arXiv preprint arXiv:2208.01618 (2022) [4](#)
10. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016) [6](#)
11. Hessel, J., Holtzman, A., Forbes, M., Bras, R.L., Choi, Y.: Clipscore: A reference-free evaluation metric for image captioning. arXiv preprint arXiv:2104.08718 (2021) [13](#)
12. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. NeurIPS **30** (2017) [13](#)
13. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. NeurIPS **33**, 6840–6851 (2020) [2](#), [4](#)
14. Ho, J., Salimans, T.: Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022) [2](#)
15. Hu, M., Zheng, J., Liu, D., Zheng, C., Wang, C., Tao, D., Cham, T.J.: Cocktail: Mixing multi-modality control for text-conditional image generation. In: NeurIPS (2023) [2](#), [4](#), [8](#), [9](#), [11](#), [13](#)
16. Huang, L., Chen, D., Liu, Y., Shen, Y., Zhao, D., Zhou, J.: Composer: Creative and controllable image synthesis with composable conditions. arXiv preprint arXiv:2302.09778 (2023) [4](#)
17. Jia, C., Luo, M., Dang, Z., Dai, G., Chang, X., Wang, M., Wang, J.: SSMG: spatial-semantic map guided diffusion model for free-form layout-to-image generation. arXiv preprint arXiv:2308.10156 (2023) [4](#)

18. Kim, S., Lee, J., Hong, K., Kim, D., Ahn, N.: Diffblender: Scalable and composable multimodal text-to-image diffusion models. arXiv preprint arXiv:2305.15194 (2023) [4](#), [9](#), [13](#)
19. Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Mallocci, M., Kolesnikov, A., et al.: The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. IJCV **128**(7), 1956–1981 (2020) [8](#)
20. Lee, J., Cho, H., Yoo, Y., Kim, S.B., Jeong, Y.: Compose and conquer: Diffusion-based 3d depth aware composable image synthesis. In: ICLR (2024) [4](#), [9](#), [13](#)
21. Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. arXiv preprint arXiv:2301.12597 (2023) [7](#)
22. 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) [4](#)
23. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV. pp. 740–755 (2014) [8](#)
24. Mo, S., Mu, F., Lin, K.H., Liu, Y., Guan, B., Li, Y., Zhou, B.: Freecontrol: Training-free spatial control of any text-to-image diffusion model with any condition. arXiv preprint arXiv:2312.07536 (2023) [4](#)
25. Mou, C., Wang, X., Xie, L., Wu, Y., 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) [2](#), [4](#), [8](#), [9](#), [13](#)
26. Nichol, A., 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. arXiv preprint arXiv:2112.10741 (2021) [4](#)
27. Qi, L., Kuen, J., Shen, T., Gu, J., Guo, W., Jia, J., Lin, Z., Yang, M.H.: High quality entity segmentation. In: ICCV (2023) [9](#)
28. 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. arXiv preprint arXiv:2305.11147 (2023) [2](#), [4](#), [8](#), [9](#), [13](#)
29. Qu, L., Wu, S., Fei, H., Nie, L., Chua, T.S.: Layoutllm-t2i: Eliciting layout guidance from llm for text-to-image generation. In: MM (2023) [4](#)
30. 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. pp. 8748–8763 (2021) [6](#)
31. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125 **1**(2), 3 (2022) [4](#)
32. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. TPAMI **44**(3) (2022) [9](#)
33. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022) [2](#), [4](#), [5](#), [9](#)
34. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI. pp. 234–241 (2015) [5](#)
35. 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) [4](#)

36. 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* **35**, 36479–36494 (2022) [4](#)
37. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al.: Laion-5b: An open large-scale dataset for training next generation image-text models. *NeurIPS* **35**, 25278–25294 (2022) [8](#)
38. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502* (2020) [2](#), [4](#)
39. Wang, X., Darrell, T., Rambhatla, S.S., Girdhar, R., Misra, I.: Instancediffusion: Instance-level control for image generation (2024) [4](#)
40. 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) [4](#)
41. Yang, B., Luo, Y., Chen, Z., Wang, G., Liang, X., Lin, L.: Law-diffusion: Complex scene generation by diffusion with layouts. In: *ICCV* (2023) [4](#)
42. 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) [4](#)
43. 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) [4](#), [12](#)
44. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: *ICCV* (2023) [2](#), [4](#), [5](#), [6](#), [8](#), [9](#), [11](#), [13](#)
45. 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* **36** (2024) [2](#), [4](#), [8](#), [9](#), [11](#), [13](#)
46. Zheng, G., Zhou, X., Li, X., Qi, Z., Shan, Y., Li, X.: Layoutdiffusion: Controllable diffusion model for layout-to-image generation. In: *CVPR* (2023) [4](#)
47. Zhuang, J., Zeng, Y., Liu, W., Yuan, C., Chen, K.: A task is worth one word: Learning with task prompts for high-quality versatile image inpainting. *arXiv preprint arXiv:2312.03594* (2023) [9](#), [12](#)