

# TKG-DM: Training-free Chroma Key Content Generation Diffusion Model

Ryugo Morita<sup>1,2</sup>, Stanislav Frolov<sup>2</sup>, Brian Bernhard Moser<sup>2</sup>, Takahiro Shirakawa,  
 Ko Watanabe<sup>2</sup>, Andreas Dengel<sup>2</sup>, Jinjia Zhou<sup>1</sup>

<sup>1</sup>Faculty of Science and Engineering, Hosei University, Tokyo, Japan

<sup>2</sup>RPTU Kaiserslautern-Landau & DFKI GmbH, Kaiserslautern, Germany

## Abstract

Diffusion models have enabled the generation of high-quality images with a strong focus on realism and textual fidelity. Yet, large-scale text-to-image models, such as Stable Diffusion, struggle to generate images where foreground objects are placed over a chroma key background, limiting their ability to separate foreground and background elements without fine-tuning. To address this limitation, we present a novel Training-Free Chroma Key Content Generation Diffusion Model (TKG-DM), which optimizes the initial random noise to produce images with foreground objects on a specifiable color background. Our proposed method is the first to explore the manipulation of the color aspects in initial noise for controlled background generation, enabling precise separation of foreground and background without fine-tuning. Extensive experiments demonstrate that our training-free method outperforms existing methods in both qualitative and quantitative evaluations, matching or surpassing fine-tuned models. Finally, we successfully extend it to other tasks (e.g., consistency models and text-to-video), highlighting its transformative potential across various generative applications where independent control of foreground and background is crucial.

## 1. Introduction

Diffusion models have disrupted the landscape of image generation by producing high-quality images with high realism and textual fidelity [8, 11, 31, 38, 44]. Research has expanded into text-to-image generation conditioned on inputs like edges [49, 56], layouts [54, 57], and Virtual Try-ON (VTON) [10, 52]. However, real-world applications such as advertising, design, and game development often require precise control over foreground and background elements. The ability to generate images with transparent or chroma key backgrounds is essential for seamlessly integrating foreground objects into various scenes [55]. Traditional models like Stable Diffusion [37] and DeepFloyd [38] struggle to generate images where foreground and

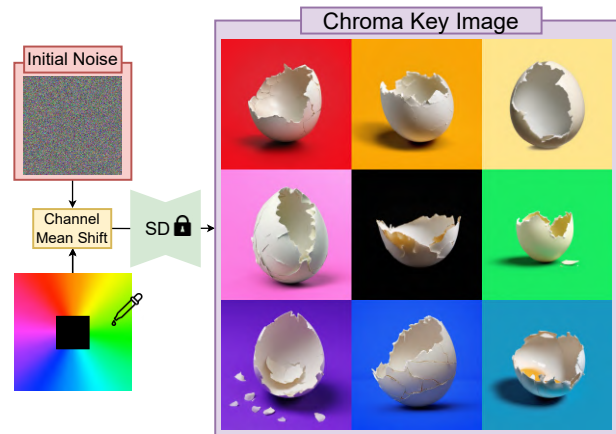


Figure 1. TKG-DM provides a training-free generation of foreground objects on a chroma key background, enabling independent control of foreground and background elements. Additionally, our method extends to conditional text-to-image and video generation tasks, making it suitable for creating foreground assets for various applications.

background elements are separated without fine-tuning [5]. Specifically, these models cannot produce foreground objects over specifiable backgrounds or generate independent foreground and background layers, limiting their applicability in workflows requiring such precise control.

Recent works like MAGICK [5] and LayerDiffuse [55] have significantly addressed this limitation. MAGICK introduces a prompt engineering-based approach for chroma key image generation. However, due to limited chroma key accuracy, it relies on post-processing and manual efforts to improve quality, including the release of foreground alpha image datasets. LayerDiffuse proposes a fine-tuning method for layer-wise generation using a dataset with 1 million images. Unfortunately, this dataset is not publicly available due to licensing restrictions, necessitating further resource-intensive dataset curation and fine-tuning.

To overcome these challenges, we present the Training-Free Chroma Key Content Generation Diffusion Model (TKG-DM), which optimizes the initial noise in the dif-

fusion process to produce images with foreground objects on specifiable color backgrounds without fine-tuning. Our method is the first to explore manipulating the color aspects of initial noise for controlled background generation, enabling precise separation of foreground and background elements. TKG-DM offers high flexibility and provides precise control over the background color, layout, size, and number of foreground objects. Our approach expands possibilities in generative content by maintaining the independence of the foreground and chroma key background. It seamlessly extends to applications like conditional text-to-image generation, consistency models, and text-to-video generation.

Our extensive experiments demonstrate that TKG-DM improves FID and mask-FID scores by 33.7% and 35.9%, respectively. Thus, our training-free model rivals fine-tuned models, offering an efficient and versatile solution for various visual content creation tasks that require precise foreground and background control. Our contributions are summarized as follows:

- We introduce TKG-DM, a training-free diffusion model that eliminates the need for fine-tuning and datasets for controllable foreground and background generation.
- TKG-DM provides precise control over background color, as well as the size, position and number of foreground elements, enabling flexible monochromatic chroma key generation.
- Our method is highly versatile, seamlessly extending to conditional text-to-image generation, consistency models, and text-to-video applications, thus enabling the creation of diverse chroma key content.
- TKG-DM outperforms existing models, delivering results that match or surpass fine-tuned alternatives while maintaining lower computational costs.

## 2. Related Works

### 2.1. Diffusion Models

Diffusion models [19, 33] generate realistic, high-quality images by iteratively refining Gaussian noise to approximate a target distribution [11, 45]. Building upon this, Latent Diffusion Models (LDMs) [37] improve efficiency by exploiting the low-dimensional latent spaces while preserving image quality. In particular, text-to-image models [32, 36–38] align text and image features through cross-attention to achieve diverse, high-quality outputs, with classifier-free guidance [18] and attention control [6, 8] further enhancing text fidelity. Beyond image generation, diffusion models find use in inpainting [24, 53], image editing [4, 15], and layout generation [7, 21]. Their versatility extends across multiple modalities, including audio generation [20, 22], text-to-video [14, 50], and 3D object generation [35, 48].

### 2.2. Initial Noise

Diffusion models typically start with Gaussian noise and iteratively refine it to generate high-quality images that align closely with a given target distribution. In general, the choice and optimization of this initial noise are crucial, as they significantly impact image quality and alignment [40, 41, 51]. More specifically, noise optimization methods, such as human preference models or attention-based scores, can enhance fidelity and text alignment without additional training [12, 13]. Moreover, optimal noise seeds and targeted noise regions significantly impact image quality and object placement [1, 51]. In contrast, noise manipulation enables precise layout control, supporting layout-aware image generation and editing [25, 26, 43]. Building on these insights, our work introduces a novel method for shifting initial noise to control chromatic information. As a result, our method enables accurate chroma key backgrounds while preserving high-quality foreground content.

### 2.3. Foreground and Background Separation in Image Generation

Foreground-background separation is crucial in advertising, design, and game development, allowing creators to manipulate individual image elements independently for tasks such as background replacement, foreground adjustments, and complex compositions [3, 9, 28, 29].

In response to this challenge, MAGICK [5] employs a prompt engineering-based model to solve this task. However, as noted in this work: “Unfortunately, all the publicly available methods we tested were incapable of consistently creating keyable images.” To overcome this limitation, MAGICK combines multiple models and applies manual effort to release a foreground alpha image dataset. Similarly, LayerDiffuse [55] introduces a fine-tuning approach using a 1 million image dataset, which cannot be publicly shared due to licensing restrictions. Although these methods improve foreground-background separation, they require substantial data collection and computational resources, making them less efficient for real-time applications. In contrast, our training-free approach provides precise control over foreground and background elements without the need for extensive datasets and fine-tuning.

## 3. Methodology

As shown in Fig. 2, TKG-DM extends Stable Diffusion [37] by manipulating the initial Gaussian noise  $\mathbf{z}_T \in \mathbb{R}^{h \times w \times 4}$  through *channel mean shift*  $F_c$ . This transformation produces the *init color noise*  $\mathbf{z}_T^* = F_c(\mathbf{z}_T)$ , which guides the vanilla Stable Diffusion to generate a uniform color image  $\mathbf{x}_0^*$  without any text prompt. To generate an image where the foreground aligns with the input text prompt  $p$  and the background has a specified color (e.g., green for chroma keying),

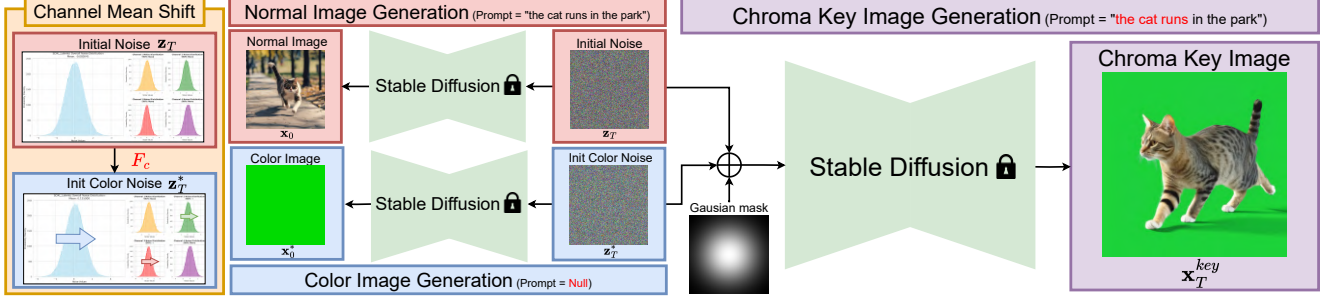


Figure 2. Starting with random noise  $\mathbf{z}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ , init color noise  $\mathbf{z}_T^*$  is generated by applying channel mean shift  $F_c$ . This produces a single-colored image  $\mathbf{x}_0^*$  without a prompt. By combining normal noise with init color noise via a Gaussian mask, TKG-DM generates a chroma key image  $\mathbf{x}_0^{key}$  with the specified foreground (e.g., “the cat runs in the park”) over a uniform background, effectively separating the foreground from the monochromatic background.

we combine the initial noise  $\mathbf{z}_T$  and the init color noise  $\mathbf{z}_T^*$  using a Gaussian mask. This combined noise  $\mathbf{z}_T^{key}$  is input into the vanilla Stable Diffusion, generating the chroma key image  $\mathbf{x}_0^{key}$ .

### 3.1. Channel Mean Shift

Inspired by the relationship between Stable Diffusion’s latent space and generated image color [47], we introduce a novel initial noise optimization technique, *channel mean shift*. It adjusts the mean of each channel in  $\mathbf{z}_T$  while keeping its standard deviation constant, enabling control over the generated image’s color. To determine the optimal shift  $\Delta_c$  for each channel  $c \in \{1, 2, 3, 4\}$  to achieve a target positive ratio in that channel. The values of channel  $c$  in the noise tensor are denoted as  $\mathbf{z}_T^{(c)}$ . The initial positive ratio for channel  $c$  is defined as:

$$\text{InitialRatio}_c = \frac{\sum_{i,j} \mathbb{1}(\mathbf{z}_T^{(c)}(i,j) > 0)}{\text{TotalPixels}_c}, \quad (1)$$

where  $\mathbb{1}(\cdot)$  is the indicator function, and  $\text{TotalPixels}_c$  is the total number of elements in  $\mathbf{z}_T^{(c)}$ . Given a target shift  $\text{TargetShift}_c$ , the target positive ratio is:

$$\text{TargetRatio}_c = \text{InitialRatio}_c + \text{TargetShift}_c. \quad (2)$$

To achieve  $\text{TargetRatio}_c$ , we iteratively adjust the mean shift  $\Delta_c$  for each channel  $\mathbf{z}_T^{(c)}$ . We initialize the shift with  $\Delta_c^{\text{init}} = 0$  and incrementally adjust  $\Delta_c$  until the positive ratio meets or exceeds  $\text{TargetRatio}_c$ . Once the target ratio is reached, we record the final shift as  $\Delta_c^{\text{final}} = \Delta_c$ . The noise tensor obtained through this method is called *init color noise*  $\mathbf{z}_T^* = F_c(\mathbf{z}_T) = \mathbf{z}_T + \Delta_c^{\text{final}}$ .

### 3.2. Init Noise Selection

To generate the foreground object on the chroma key background, we apply an init noise selection strategy that selectively combines the initial noise  $\mathbf{z}_T$  and the init color noise

$\mathbf{z}_T^*$  using a 2D Gaussian mask  $\mathbf{A}(i, j)$ . This mask creates a gradual transition by preserving the original noise in the foreground region and applying the color-shifted noise to the background region. Formally, the masked initial noise is computed as:

$$\mathbf{z}_T^{key}(i, j) = \mathbf{A}(i, j) \cdot \mathbf{z}_T(i, j) + (1 - \mathbf{A}(i, j)) \cdot \mathbf{z}_T^*(i, j), \quad (3)$$

where  $\mathbf{A}(i, j) = e^{-\frac{(i-\mu_i)^2 + (j-\mu_j)^2}{2\sigma^2}}$  controls the blend between the original noise and init color noise and  $\mu_i$  and  $\mu_j$  specify the center and  $\sigma$  the spread.

Adjusting the Gaussian parameters  $\mu_i$ ,  $\mu_j$ , and  $\sigma$  controls the position and size of the foreground: shifting  $\mu_i$  and  $\mu_j$  adjusts its position while increasing  $\sigma$  enlarges the foreground element. Interestingly, TKG-DM also generates multiple foreground elements by applying multiple Gaussian masks, each independently configured to produce varied positions and sizes for flexible compositions.

## 4. Mechanism of TKG-DM

In this section, we explain how TKG-DM specifies the background color and generates the foreground and background separately to generate chroma key content.

### 4.1. Color Specify via Channel Mean Shift

Inspired by previous research [47], we control the chroma key background color by applying channel mean shift to specific channels of the initial noise  $\mathbf{z}_T \in \mathbb{R}^{h \times w \times 4}$ . Specifically, we adjust the mean of each channel  $\mathbf{z}_T^{(c)}$ , where  $c \in \{1, 2, 3, 4\}$ , to influence the color composition of the generated images. In this experiment, we set  $\text{TargetShift}_c = \pm 7\%$ , adjusting the positive ratio of each channel by adding or subtracting 7% relative to  $\text{InitialRatio}_c$ .

As shown in Fig. 3a, positive shifts in  $\mathbf{z}_T^{(2)}$  and  $\mathbf{z}_T^{(3)}$  intensify cyan and yellow tones, respectively, while negative shifts emphasize red and blue-purple hues. Shifts in  $\mathbf{z}_T^{(1)}$  and  $\mathbf{z}_T^{(4)}$  primarily affect luminance and shades of white and

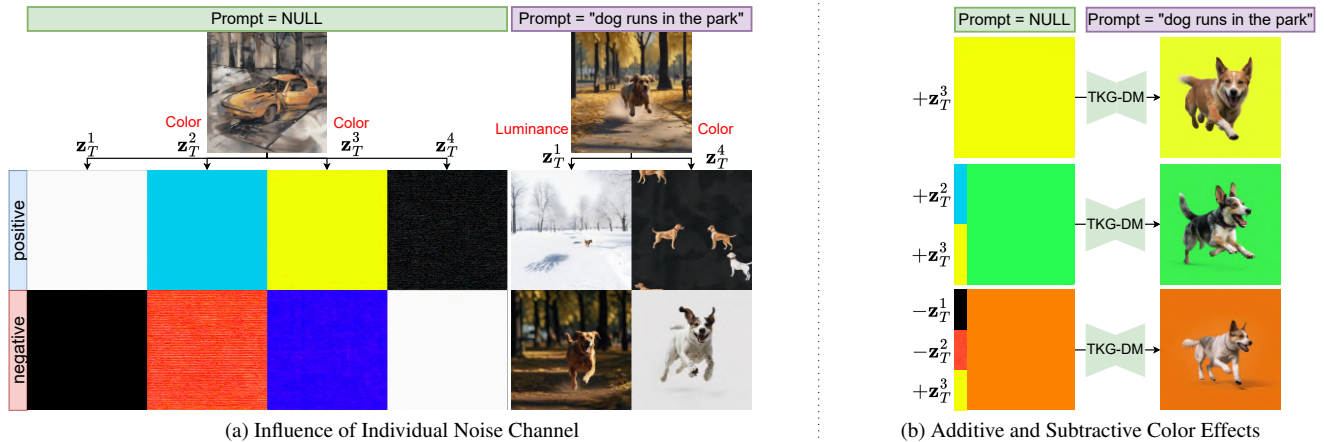


Figure 3. Fig. 3a illustrates the relationship between positive/negative channel mean shift with initial noise and color variations in the generated images. Different channel shifts across noise channels result in changes in the hue of the generated image. Fig. 3b shows how simultaneous shifts across multiple channels facilitate additive and subtractive color mixing, providing intuitive and flexible color control.

black, respectively. These effects are confirmed by comparing generated images conditioned on specific text prompts.

When we apply simultaneous shifts to multiple channels, we observe intuitive color mixing that aligns with additive and subtractive color theories. For example, positive shifts in both  $z_T^{(2)}$  and  $z_T^{(3)}$  enhance green tones, while combining red and yellow shifts with reduced luminance yields orange hues (see Fig. 3b). This channel-based manipulation enables flexible and natural color control in the image generation process.

#### 4.2. Content Separation via Attention Mechanism

As illustrated in Fig. 4, self-attention and cross-attention mechanisms play distinct roles in generating the foreground and background elements, which is crucial for achieving effective foreground-background separation in TKG-DM.

For the foreground, self-attention ensures internal consistency and coherence within the object, while cross-attention aligns the generated content with the text prompt. Due to the inherent bias in the training dataset [42], where foreground objects are more prominently described in captions, cross-attention forms a strong link between the foreground and the text prompt.

For the background, the init color noise introduced by channel mean shift dominates the generation process. The self-attention mechanism synergizes with this modified initial noise to guide the background toward the specified chroma key color. Since background elements are often less detailed or vaguely described in dataset captions, the cross-attention has a weaker influence on the background, allowing the init color noise to take precedence.

By exploiting this bias and manipulating the initial noise,

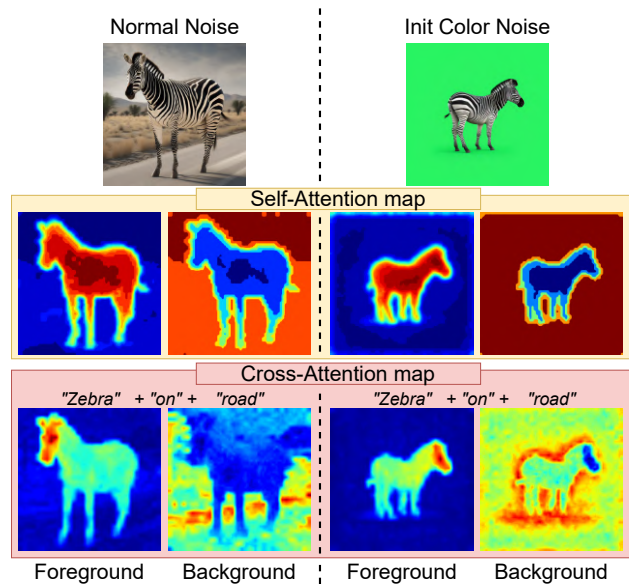


Figure 4. **Visualization of self- and cross-attention maps.** The self-attention maps, segmented using a foreground segmentation model [58], show the regions attended to during generation. The cross-attention maps illustrate how the model attends to relevant regions based on the foreground and background prompts. This analysis shows how self-attention and cross-attention interact to influence foreground and background content generation.

TKG-DM effectively decouples the background from the text prompt. This results in a uniform chroma key background and enables the isolated generation of foreground content without interference from undesired background details.



Figure 5. **Qualitative comparison in SD1.5.** Existing methods fail to produce an accurate chroma key background. In contrast, ours produces a highly accurate chroma key background while generating high-quality foregrounds without Green Background Prompt (GBP).

## 5. Experiments

### 5.1. Experimental Setup

We adopt SD1.5 and SDXL in TKG-DM to generate images at  $512 \times 512$  and  $1024 \times 1024$  resolutions, respectively, with DDIMScheduler, a guidance scale of 7.5, and 50 denoising steps. To obtain the init color noise for a green background, we apply channel mean shift with channels 2 and 3 in the positive direction for  $\text{TargetShift}_c = +7\%$ . Gaussian mask parameters are set to  $(\sigma, \mu_i, \mu_j) = (0.5, w/2, h/2)$ .

We compare our method against the naive usage of SD1.5 [37], SDXL [34], and DeepFloyd [38] using a Green Background Prompt (GBP) (“isolated on a solid green background”) following the setup of MAGICK [5], as well as against fine-tuned models, namely GreenBack LoRA from CivitAI<sup>1</sup> and LayerDiffuse [55].

### 5.2. Dataset and Metrics

We collect 3,000 images from the MAGICK dataset [5], tailored for foreground image generation from text prompts. The MAGICK dataset is constructed by first generating monochromatic backgrounds using Deepfloyd [38] with GBP, followed by quality enhancement using SDEdit [27] and SDXL’s img2img [34]. Foreground images with alpha values are curated using pixel-, deep learning-, and human-based methods. Our experiment utilizes alpha images, alpha masks, and corresponding text prompts.

We construct ground truth images for evaluation by overlaying the alpha images onto a lime green background images (RGB = 50, 205, 50). We use the Fréchet Inception

Distance (FID) to assess foreground quality [17]. Additionally, we introduce a novel m-FID metric, which compares the mask extracted from generated chroma key content using BiRefNet [58] to the ground truth mask, providing a comprehensive assessment of foreground accuracy. For semantic alignment, we employ CLIP-Sentence (CLIP-S) and CLIP-Image (CLIP-I) [16]: CLIP-S measures alignment with the text, while CLIP-I assesses visual similarity to the ground truth.

### 5.3. Qualitative Result

Fig. 5 and Fig. 6 show the qualitative results of our model and existing methods with SD1.5 and SDXL, respectively. They demonstrate that our model generates high-precision chroma key images without prompt engineering.

With the Green Background Prompt (GBP), SD1.5 struggles to generate a clean green background, while SDXL, while slightly better, generates unstable light green tints on the background and often produces multiple foreground objects, complicating foreground-background separation. Additionally, the GBP’s “green” element unintentionally adds green tones to the foreground, such as on clothing or in a cat’s eyes (Fig. 6, first and third columns). DeepFloyd performs the worst foreground quality and text alignment.

In the fine-tuned model, while LoRA (SD1.5) shows minor improvements, it still struggles with chroma key accuracy. LayerDiffuse produces well-separated foregrounds but occasionally loses details, like precise numbers or letters, due to dataset limitations in fine-tuning. Mask generation also occasionally fails, resulting in uncut images.

Moreover, existing models struggle to generate accurate chroma key images when background information is in-

<sup>1</sup><https://civitai.com/>

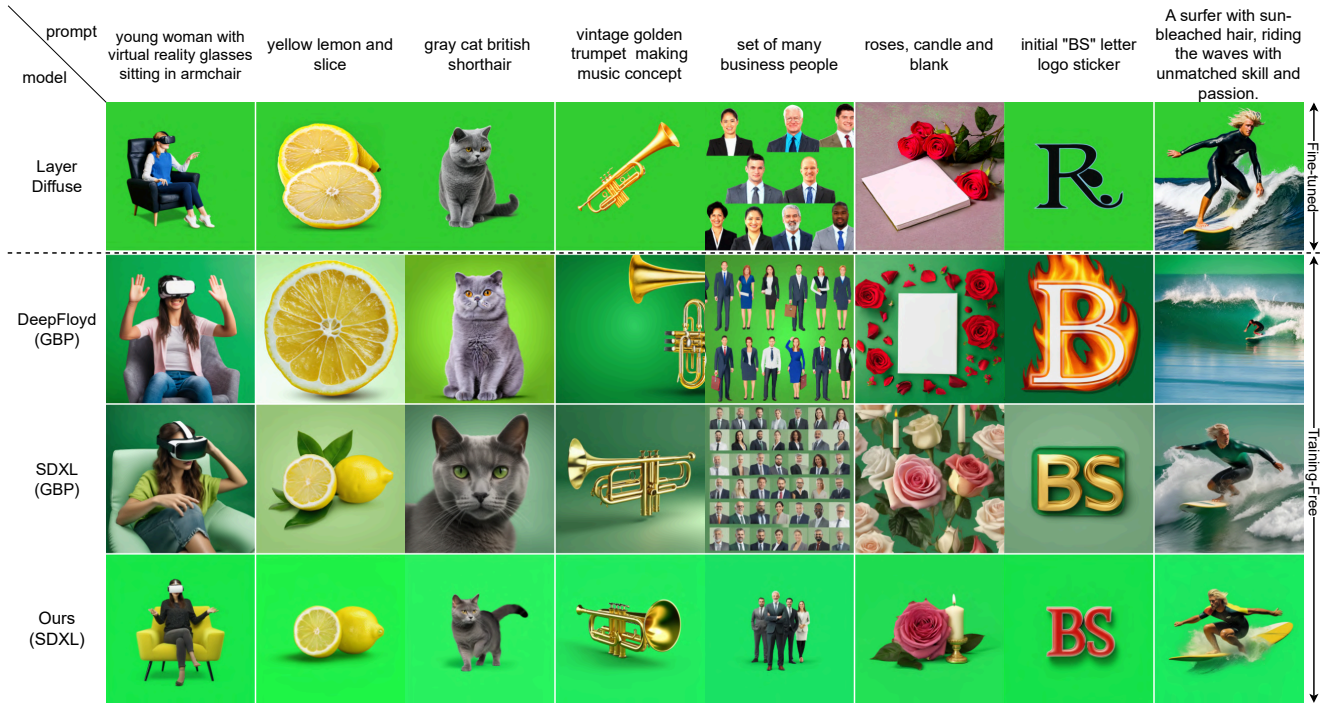


Figure 6. **Qualitative comparison in SDXL.** Existing training-free methods fail to isolate foreground objects from the background. While LayerDiffuse, a fine-tuned model, successfully generates the transparent foregrounds on the lime green background, it may struggle with accurate text generation and background handling. In contrast, our model efficiently generates highly accurate chroma key images with foreground content placed over the background without fine-tuning and needing Green Background Prompt (GBP).



Figure 7. **Qualitative various color result in SDXL.** By adjusting the channel mean shift, TKG-DM controls the background color while generating a high-quality foreground.

cluded in the text prompt (Fig. 6, last columns). In contrast, our training-free model consistently produces high-quality chroma key backgrounds with only the foreground content, ensuring precise separation.

Furthermore, as shown in Fig. 7, TKG-DM generates the various color background images through channel mean shift without prompt engineering. Additional results are provided in the supplementary material.

Table 1. **Quantitative results.** It indicates that our training-free model outperforms the existing model with Green Background Prompt (GBP) and rivals the fine-tuned model, LayerDiffuse [55].

SD1.5				
	FID ↓	m-FID ↓	CLIP-I ↑	CLIP-S ↑
SD1.5 (GBP) [37]	85.00	63.54	0.710	0.256
GB LoRA (GBP)	60.29	49.03	0.704	0.243
<b>Ours</b>	<b>56.32</b>	<b>40.75</b>	<b>0.737</b>	<b>0.261</b>
SDXL				
DeepFloyd (GBP) [38]	31.57	20.31	0.781	0.270
SDXL (GBP) [34]	45.32	39.17	0.759	0.272
LayerDiffuse [55]	29.34	29.82	0.778	0.276
<b>Ours</b>	<b>41.81</b>	<b>31.43</b>	<b>0.763</b>	<b>0.273</b>




Figure 8. Results from our user study show the preference rates for foreground objects generated by our model and existing models based on their image quality and alignment with the text prompt. GBP means a model with a Green Background Prompt.

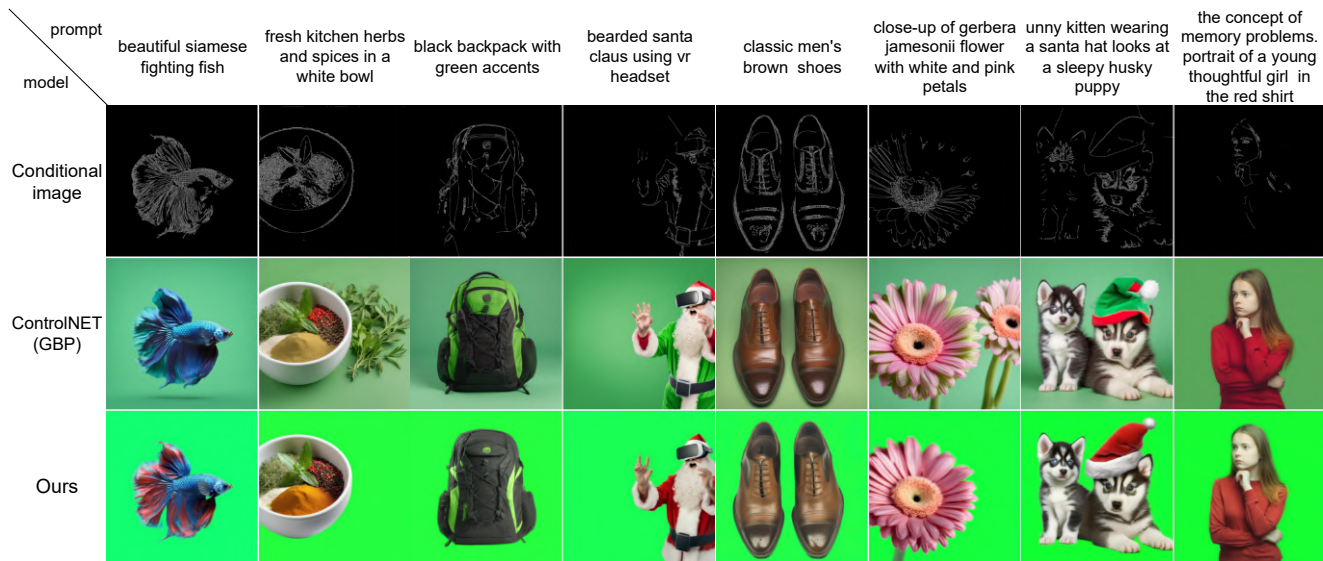


Figure 9. **Qualitative comparison in ControlNet.** Existing methods typically generate images with issues such as color erosion in the foreground, uneven backgrounds, or unintended elements from the conditioning input. In contrast, our model generates images with the foreground isolated from the background, maintaining a consistent chroma key background without the Green Background Prompt (GBP).

## 5.4. Quantitative Result

As shown in Table 1, our method consistently outperforms other training-free models across all metrics, achieving better image realism and mask accuracy, as indicated by lower FID and m-FID scores. It also aligns well with input text prompts and has closer visual similarity to the ground truth, as indicated by higher CLIP-S and CLIP-I scores.

In the SDXL setting, LayerDiffuse achieves the best FID, while our model remains competitive across all other metrics. The FID advantage for LayerDiffuse likely results from its setup, which overlays foregrounds on a lime green background matching the ground truth.

DeepFloyd’s high FID, m-FID, and CLIP-I scores reflect its similarity to the ground truth based on DeepFloyd’s outputs. However, this alignment gives it an inherent advantage, making it unsuitable as a fair benchmark for image quality. Its lower CLIP-S score further indicates weaker text alignment compared to other models.

Overall, these results underscore our model’s ability to generate high-quality, text-aligned foregrounds without fine-tuning, offering an efficient chroma key content generation solution.

## 5.5. User Study

Besides objective metrics, we conducted a user study to assess subjective human preferences for prompt adherence and image quality, comparing our method against SDXL with Green Background Prompt (GBP) and LayerDiffuse. We generated 30 image pairs per method, totaling 60 images. 100 participants ranked images via two alternative

Table 2. Quantitative results comparing our method with existing methods for ControlNet [56]. GBP means a model with a Green Background Prompt.

	ControlNet [56]			
	FID ↓	m-FID ↓	CLIP-I ↑	CLIP-S ↑
SDXL (GBP) [34]	22.04	18.62	0.819	0.279
<b>Ours</b>	<b>17.09</b>	<b>17.22</b>	<b>0.834</b>	<b>0.284</b>

forced-choice [30, 39] and based on foreground quality and text alignment. To ensure focus on the foreground, we used a foreground extraction model [58] and manually refined areas as needed. As shown in Fig. 8, our training-free approach improves object realism and text alignment over the existing model.

## 6. Applications

We demonstrate the flexibility of TKG-DM in various tasks, implemented with the SD1.5 and SDXL, as presented in Table 2, Fig. 9 and Fig. 10. Further details and additional results are provided in the supplementary material.

**Application to ControlNet.** As depicted in Fig. 9 and Table 2, our method integrates seamlessly with ControlNet [56], enabling precise generation control using conditioning inputs like canny edges. By applying a Gaussian mask to the init color noise, we preserve normal initial noise in the foreground regions, resulting in distinct foregrounds on uniform chroma key backgrounds.

Compared to the standard ControlNet, our method demonstrates superior qualitative and quantitative perfor-

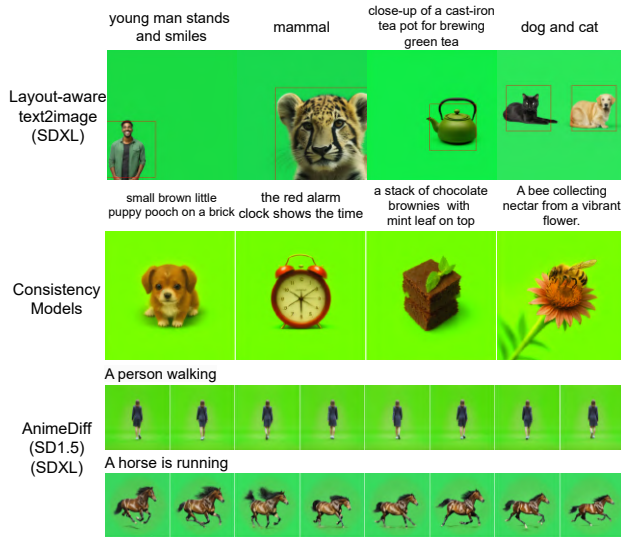


Figure 10. **Application results of TKG-DM across various tasks.** TKG-DM effectively supports layout-aware text-to-image generation, consistency models and text-to-video. Each row demonstrates its versatility in adapting to different domains, from realistic object and character placements to animation.

mance. Existing methods often face issues like color erosion in the foreground, uneven backgrounds, or unintended elements from the conditioning input. By isolating the foreground from background influence and maintaining a consistent chroma key background, our approach is highly effective for compositing tasks without prompt engineering.

**Layout-Aware Text-to-Image Generation.** As discussed in Section 3.2, our model enables control over foreground objects’ location and size by adjusting the Gaussian mask’s center and spread. This control allows users to place objects precisely within specific areas of the generated image, making it practical for layout-aware tasks. Additionally, using multiple Gaussian masks enables the generation of multiple foreground objects within a single scene (Fig. 10, top-right). Unlike existing models that lack targeted control over object placement and size, our approach integrates layout information seamlessly without fine-tuning.

**Application to Consistency Models.** Consistency Models [46] are generative models that produce high-quality outputs in a few steps, unlike traditional diffusion models. Operating in latent space, Consistency Models demonstrate the robustness of TKG-DM to generative models that handle images in the latent diffusion-based model, enabling significantly faster generation while maintaining high output quality without modifications.

**Application to Text-to-Video.** Our TKG-DM extends seamlessly to text-to-video tasks, ensuring consistent foreground content across frames. We align foreground objects with the input prompt by applying TKG-DM with

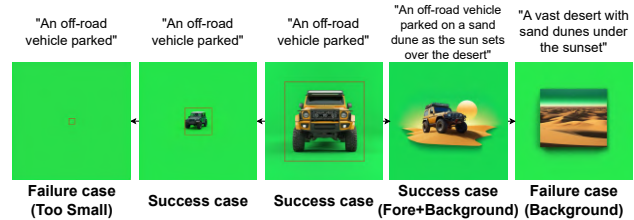


Figure 11. **Failure Cases.** TKG-DM fails to generate a foreground object when the size is too small or when the text prompt lacks object information, resulting in only the background being generated.

AnimateDiff [14] on each frame while maintaining a uniform background for chroma key video generation. This approach, the first to generate foreground video with controlled chroma key backgrounds, facilitates efficient and flexible video editing workflows involving background removal or replacement.

## 7. Limitation and Future Work

As shown in Fig. 11, TKG-DM primarily focuses on generating foreground content (e.g., objects and people), which limits its ability to generate backgrounds like landscapes. Additionally, if the size parameter is too small, ignore generating the foreground object [26, 43].

Future enhancements could expand TKG-DM’s capabilities to include controlled background generation, allowing separate and adjustable generation of both foreground and background content. This would greatly benefit applications in video and 3D content creation where isolated and reusable background assets are essential.

## 8. Conclusion

This work presents the first exploration of initial noise color manipulation in diffusion models, demonstrating its effectiveness in controlling background chromaticity. Leveraging this insight, we introduce the Training-Free Chroma Key Content Generation Diffusion Model (TKG-DM), which generates prompt-based foregrounds on a chroma key background without fine-tuning. TKG-DM extends seamlessly to tasks like consistency models and text-to-video, highlighting its versatility beyond static images. Experimental results show that our method matches or surpasses fine-tuned models in both qualitative and quantitative evaluations without additional training. This approach not only advances theoretical understanding but also provides practical benefits for real-world applications where isolating foreground content is essential.

## Acknowledgements

This research was supported by BMBF project SustainML (Grant 101070408) and the BMDV NITD (45KI22B021).

## References

- [1] Yuanhao Ban, Ruochen Wang, Tianyi Zhou, Boqing Gong, Cho-Jui Hsieh, and Minhao Cheng. The crystal ball hypothesis in diffusion models: Anticipating object positions from initial noise. *arXiv preprint arXiv:2406.01970*, 2024. 2
- [2] Omer Bar-Tal, Lior Yariv, Yaron Lipman, and Tali Dekel. Multidiffusion: Fusing diffusion paths for controlled image generation, 2023. 2
- [3] Yaniv Benny and Lior Wolf. Onegan: Simultaneous unsupervised learning of conditional image generation, foreground segmentation, and fine-grained clustering. In *ECCV*, pages 514–530. Springer, 2020. 2
- [4] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, pages 18392–18402, 2023. 2
- [5] Ryan D Burgert, Brian L Price, Jason Kuen, Yijun Li, and Michael S Ryoo. Magick: A large-scale captioned dataset from matting generated images using chroma keying. In *CVPR*, pages 22595–22604, 2024. 1, 2, 5
- [6] Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiaohu Qie, and Yinqiang Zheng. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *ICCV*, pages 22560–22570, 2023. 2
- [7] Shang Chai, Liansheng Zhuang, and Fengying Yan. Layoutdm: Transformer-based diffusion model for layout generation. In *CVPR*, pages 18349–18358, 2023. 2
- [8] Hila Chefer, Yuval Alaluf, Yael Vinker, Lior Wolf, and Daniel Cohen-Or. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4):1–10, 2023. 1, 2
- [9] Ruidong Chen, Lanjun Wang, Weizhi Nie, Yongdong Zhang, and An-An Liu. Anyscene: Customized image synthesis with composited foreground. In *CVPR*, pages 8724–8733, 2024. 2
- [10] Yisol Choi, Sangkyung Kwak, Kyungmin Lee, Hyungwon Choi, and Jinwoo Shin. Improving diffusion models for virtual try-on. *arXiv preprint arXiv:2403.05139*, 2024. 1
- [11] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *NeurIPS*, 34:8780–8794, 2021. 1, 2
- [12] Luca Eyring, Shyamgopal Karthik, Karsten Roth, Alexey Dosovitskiy, and Zeynep Akata. Reno: Enhancing one-step text-to-image models through reward-based noise optimization. *arXiv preprint arXiv:2406.04312*, 2024. 2
- [13] Xiefan Guo, Jinlin Liu, Miaomiao Cui, Jiankai Li, Hongyu Yang, and Di Huang. Initno: Boosting text-to-image diffusion models via initial noise optimization. In *CVPR*, pages 9380–9389, 2024. 2
- [14] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. *arXiv preprint arXiv:2307.04725*, 2023. 2, 8
- [15] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 2
- [16] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. *arXiv preprint arXiv:2104.08718*, 2021. 5
- [17] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS*, 30, 2017. 5
- [18] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 2
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 33:6840–6851, 2020. 2
- [20] Rongjie Huang, Jiawei Huang, Dongchao Yang, Yi Ren, Luping Liu, Mingze Li, Zhenhui Ye, Jinglin Liu, Xiang Yin, and Zhou Zhao. Make-an-audio: Text-to-audio generation with prompt-enhanced diffusion models. In *ICML*, pages 13916–13932. PMLR, 2023. 2
- [21] Naoto Inoue, Kotaro Kikuchi, Edgar Simo-Serra, Mayu Otani, and Kota Yamaguchi. Layoutdm: Discrete diffusion model for controllable layout generation. In *CVPR*, pages 10167–10176, 2023. 2
- [22] Zhifeng Kong, Wei Ping, Jiayi Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*, 2020. 2
- [23] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022. 2
- [24] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *CVPR*, pages 11461–11471, 2022. 2
- [25] Jiafeng Mao, Xueting Wang, and Kiyoharu Aizawa. Guided image synthesis via initial image editing in diffusion model. *arXiv preprint arXiv:2305.03382*, 2023. 2
- [26] Jiafeng Mao, Xueting Wang, and Kiyoharu Aizawa. Semantic-driven initial image construction for guided image synthesis in diffusion model. *arXiv preprint arXiv:2312.08872*, 2023. 2, 8
- [27] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jianjun Wu, Jun-Yan Zhu, and Stefano Ermon. SDEdit: Guided image synthesis and editing with stochastic differential equations. In *ICLR*, 2022. 5
- [28] Ryugo Morita, Zhiqiang Zhang, Man M Ho, and Jinjia Zhou. Interactive image manipulation with complex text instructions. In *WACV*, pages 1053–1062, 2023. 2
- [29] Ryugo Morita, Zhiqiang Zhang, and Jinjia Zhou. Batinet: Background-aware text to image synthesis and manipulation network. In *ICIP*, pages 765–769. IEEE, 2023. 2
- [30] Brian B Moser, Stanislav Frolov, Federico Raue, Sebastian Palacio, and Andreas Dengel. Yoda: You only diffuse areas. an area-masked diffusion approach for image super-resolution. *arXiv preprint arXiv:2308.07977*, 2023. 7
- [31] Brian B Moser, Arundhati S Shanbhag, Federico Raue, Stanislav Frolov, Sebastian Palacio, and Andreas Dengel. Diffusion models, image super-resolution, and everything: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 2024. 1

- [32] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. 2
- [33] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *ICML*, pages 8162–8171. PMLR, 2021. 2
- [34] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 5, 6, 7
- [35] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. *arXiv preprint arXiv:2306.17843*, 2023. 2
- [36] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, pages 8821–8831. Pmlr, 2021. 2
- [37] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695, 2022. 1, 2, 5, 6
- [38] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *NeurIPS*, 35:36479–36494, 2022. 1, 2, 5, 6
- [39] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE TPAMI*, 45(4): 4713–4726, 2022. 7
- [40] Dvir Samuel, Rami Ben-Ari, Nir Darshan, Haggai Maron, and Gal Chechik. Norm-guided latent space exploration for text-to-image generation. *NeurIPS*, 36, 2024. 2
- [41] Dvir Samuel, Rami Ben-Ari, Simon Raviv, Nir Darshan, and Gal Chechik. Generating images of rare concepts using pre-trained diffusion models. In *AAAI*, pages 4695–4703, 2024. 2
- [42] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 4
- [43] Takahiro Shirakawa and Seiichi Uchida. Noisecollage: A layout-aware text-to-image diffusion model based on noise cropping and merging. In *CVPR*, pages 8921–8930, 2024. 2, 8
- [44] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, pages 2256–2265. PMLR, 2015. 1
- [45] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2
- [46] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *arXiv preprint arXiv:2303.01469*, 2023. 8
- [47] Timothy Alexis Vass. Explaining the sdxl latent space. <https://huggingface.co/blog/TimothyAlexisVass/explaining-the-sdxl-latent-space>, 2024. 3
- [48] Vikram Voleti, Chun-Han Yao, Mark Boss, Adam Letts, David Pankratz, Dmitry Tochilkin, Christian Laforte, Robin Rombach, and Varun Jampani. Sv3d: Novel multi-view synthesis and 3d generation from a single image using latent video diffusion. *arXiv preprint arXiv:2403.12008*, 2024. 2
- [49] Andrey Voynov, Kfir Aberman, and Daniel Cohen-Or. Sketch-guided text-to-image diffusion models. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023. 1
- [50] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaoju Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *ICCV*, pages 7623–7633, 2023. 2
- [51] Katherine Xu, Lingzhi Zhang, and Jianbo Shi. Good seed makes a good crop: Discovering secret seeds in text-to-image diffusion models. *arXiv preprint arXiv:2405.14828*, 2024. 2
- [52] Yuhao Xu, Tao Gu, Weifeng Chen, and Chengcai Chen. Oot-diffusion: Outfitting fusion based latent diffusion for controllable virtual try-on. *arXiv preprint arXiv:2403.01779*, 2024. 1
- [53] Binxin Yang, Shuyang Gu, Bo Zhang, Ting Zhang, Xuejin Chen, Xiaoyan Sun, Dong Chen, and Fang Wen. Paint by example: Exemplar-based image editing with diffusion models. In *CVPR*, pages 18381–18391, 2023. 2
- [54] Zhengyuan Yang, Jianfeng Wang, Zhe Gan, Linjie Li, Kevin Lin, Chenfei Wu, Nan Duan, Zicheng Liu, Ce Liu, Michael Zeng, et al. Reco: Region-controlled text-to-image generation. In *CVPR*, pages 14246–14255, 2023. 1
- [55] Lvmin Zhang and Maneesh Agrawala. Transparent image layer diffusion using latent transparency. *arXiv preprint arXiv:2402.17113*, 2024. 1, 2, 5, 6
- [56] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, pages 3836–3847, 2023. 1, 7
- [57] Guangcong Zheng, Xianpan Zhou, Xuewei Li, Zhongang Qi, Ying Shan, and Xi Li. Layoutdiffusion: Controllable diffusion model for layout-to-image generation. In *CVPR*, pages 22490–22499, 2023. 1
- [58] Peng Zheng, Dehong Gao, Deng-Ping Fan, Li Liu, Jorma Laaksonen, Wanli Ouyang, and Nicu Sebe. Bilateral reference for high-resolution dichotomous image segmentation. *arXiv preprint arXiv:2401.03407*, 2024. 4, 5, 7